

Dell PowerStore: VMware vSphere Best Practices

October 2023

H18116.6

White Paper

Abstract

This document provides best practices for integrating VMware vSphere hosts with Dell PowerStore.

Copyright

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

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

Copyright © 2020–2023 Dell Inc. or its subsidiaries. All Rights Reserved. Published in the USA October 2023 H18116.6.

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

Contents

Executive summary	4
Introduction	5
Host configuration	6
NVMe over Fabric (NVMe-oF)	9
Sizing and performance optimization	16
Management and monitoring	21
Data protection and disaster recovery	27
References	33

Executive summary

Introduction This document provides recommendations, tips, and other helpful guidelines for integrating external VMware vSphere hosts with the Dell PowerStore platform.

Audience This document is intended for IT administrators, storage architects, partners, and Dell Technologies employees. This audience also includes any individuals who may evaluate, acquire, manage, operate, or design a Dell Technologies networked storage environment using PowerStore systems.

Revisions

Date	Part number/ revision	Description
April 2020	H18116	Initial release PowerStoreOS 1.0
April 2021	H18116.1	Updates for PowerStoreOS 2.0
January 2022	H18116.2	Updates for PowerStoreOS 2.1. Document template update.
July 2022	H18116.3	Updates for PowerStoreOS 3.0
December 2022	H18116.4	Updates for PowerStoreOS 3.2 and vSphere 8
May 2023	H18116.5	Updates for PowerStoreOS 3.5: Secure snapshots and PowerStore backup to PowerProtect DD series appliances
October 2023	H18116.6	Updates for PowerStoreOS 3.6

We value your feedback

Dell Technologies and the authors of this document welcome your feedback on this document. Contact the Dell Technologies team by [email](#).

Author: Darin Schmitz

Contributor: Jason Boche

Note: For links to other documentation for this topic, see the [PowerStore Info Hub](#).

Introduction

PowerStore overview

PowerStore is a robust and flexible storage option that is ideal for use with VMware vSphere.

PowerStore achieves new levels of operational simplicity and agility. It uses a container-based microservices architecture, advanced storage technologies, and integrated machine learning to unlock the power of your data. PowerStore is a versatile platform with a performance-centric design that delivers multidimensional scale, always-on data reduction, and support for next-generation media.

PowerStore brings the simplicity of public cloud to on-premises infrastructure, streamlining operations with an integrated machine-learning engine and seamless automation. It also offers predictive analytics to easily monitor, analyze, and troubleshoot the environment. PowerStore is highly adaptable, providing the flexibility to host specialized workloads directly on the appliance and modernize infrastructure without disruption. It also offers investment protection through flexible payment solutions and data-in-place upgrades.

vSphere overview

VMware vSphere is the industry-leading virtualization platform and a core building block to the software-defined data center (SDDC). VMware vSphere is primarily composed of vCenter for management and ESXi hosts that provide the hypervisor for compute and memory virtualization.

Prerequisite reading

Before implementing the best practices in this document, we recommend reviewing and implementing the recommended configurations in the *PowerStore Host Configuration Guide* and reviewing other resources available at Dell.com/powerstoredocs.

Terminology

The following table provides definitions for some of the terms that are used in this document.

Table 1. Terminology

Term	Definition
Appliance	The solution containing a base enclosure and attached expansion enclosures. The size of an appliance could include only the base enclosure or the base enclosure plus expansion enclosures.
Base enclosure	The enclosure containing both nodes (node A and node B) and the NVMe drive slots.
Cluster	Multiple PowerStore appliances in a single grouping.
Expansion enclosure	An enclosure that can be attached to a base enclosure to provide additional SAS-based drive slots.
Node	The component within the base enclosure that contains processors and memory. Each appliance consists of two nodes.
NVM Express over Fabrics (NVMe-oF)	NVMe command fabric which includes Fibre Channel and TCP/IP transport protocols, among others.
NVMe over Fibre Channel (NVMe/FC)	Allows hosts to access storage systems across a network fabric with the NVMe protocol using Fibre Channel as the underlying transport.

Term	Definition
NVMe over TCP (NVMe/TCP)	Allows hosts to access storage systems across a network fabric with the NVMe protocol using TCP as the underlying transport.
PowerStore Manager	The web-based user interface (UI) for storage management.

Host configuration

Introduction

While most settings for stand-alone ESXi hosts that are connected to PowerStore appliances can remain at the default values, some changes are required for PowerStore stability, performance, and efficiency. The recommended changes and instructions about how to set them are specified in the document *PowerStore Host Configuration Guide* on Dell.com/powerstoredocs. While administrators can use this section for high-level explanations and reasoning behind the recommendations, administrators should always consult the Host Configuration Guide for the current settings.

Caution: These recommended settings are for external ESXi hosts only and do not apply to the ESXi instances running within PowerStore X model appliances.

Note: The Virtual Storage Integrator (VSI) allows administrators to easily configure the ESXi host best-practice settings with PowerStore. See [Virtual Storage Integrator](#) for more details.

Queue depth

There are multiple locations in ESXi and the guest operating systems to modify queue depth. While increasing the queue depth in an application, vSCSI device, or ESXi driver module can potentially increase performance, modifying or increasing queue depths can potentially overwhelm the array. For details about queue-depth settings, see the document *PowerStore Host Configuration Guide*.

Timeouts

Setting disk timeouts is an important factor for applications to survive both unexpected and expected node outages, such as failures or rebooting for updates. While the default SCSI timeout in most applications and operating systems is 30 seconds, storage vendors (including Dell Technologies) and application vendors typically recommend increasing these timeouts to 60 seconds or more to help ensure uptime. Two of the main locations to change the timeouts are at the ESXi host level and at the virtual-machine-guest-OS level. For details about setting timeouts for ESXi, see the *PowerStore Host Configuration Guide*.

ESXi host timeouts

The timeout values set at the ESXi-host-driver level help ensure that the hosts and virtual machines can survive a storage node failover event. For details about setting timeouts for ESXi, see the *PowerStore Host Configuration Guide*.

Guest operating system timeouts

If VMware Tools are installed into a guest operating system, they automatically set the timeout values. However, if the guest operating system does not have VMware Tools installed, the administrator can set these values manually. While [VMware documentation](#) has examples for setting the disk timeouts in Microsoft Windows guest operating systems,

consult the knowledge bases from operating system vendors to obtain specific guest settings.

Multipathing

With the vSphere Pluggable Storage Architecture (PSA), the storage protocol determines which Multipathing Plugin (MPP) is assigned to volumes mapped from the PowerStore array. With SCSI-based protocols such as Fibre Channel and iSCSI, the Native Multipathing Plug-in (NMP) is used, whereas with NVMe-oF, the VMware High Performance Plug-in (HPP) is used.

Native Multipathing Plug-in

SCSI-based volumes using Fibre Channel and iSCSI are automatically assigned the Native Multipathing Plug-in (NMP). However, the ESXi Storage Array Type Plug-in (SATP) module and its corresponding path selection policy (PSP) may require you to configure claim rules to use Round Robin (RR) with PowerStore appliances. Applying the settings in the *PowerStore Host Configuration Guide* ensures that all volumes presented to the host use Round Robin as the default pathing policy.

Also, the recommended `esxcli` command sets the IOPS path-change condition to one I/O per path. While the default setting in the RR PSP sends 1,000 IOPS down each path before switching to the next path, this recommended setting instructs ESXi to send one command down each path. This setting results in better utilization of each path's bandwidth, which is useful for applications that send large I/O block sizes to the array.

According to the *PowerStore Host Configuration Guide*, SSH to each ESXi host using root credentials to issue the following command (reboot required):

```
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
```

The claim rule can also be added to discovered ESXi hosts using VMware PowerCLI:

Note: The following commands are for vSphere 7 and 8 ESXi hosts. ESXi 6.7 hosts should also include the `disable_action_OnRetryErrors` option. See the *PowerStore Host Configuration Guide* for more information.

```
# Add or remove a claim rule on each vSphere host
$esxlist | ForEach-Object {
$esxcli = Get-EsxCli -VMHost $_ -V2

# Fill the hash table (optional params are not required)
$sRule = @{
    satp = 'VMW_SATP_ALUA' #esxcli: -s
    psp = 'VMW_PSP_RR' #esxcli: -P
    pspoption = 'iops=1' #esxcli: -O
    claimoption = 'tpgs_on' #esxcli: -c
    #option = 'disable_action_OnRetryErrors' #esxcli: -o
    vendor = 'DellEMC' #esxcli: -V
    model = 'PowerStore' #esxcli: -M
    description = 'PowerStore' #esxcli: -e
```

```

}

# Call the esxcli command to add/remove the rule
Write-Host $selection "rule on" $_
$esxcli.storage.nmp.satp.rule.$selection.Invoke($sRule)

```

High Performance Plug-in

For NVMe-oF targets, the High Performance Plug-in (HPP) replaces the NMP. The HPP will claim NVMe devices and is designed to improve storage performance for modern high-speed interfaces.

The HPP has multiple Path Selection Schemes (PSS) available to determine which physical paths are used for I/O requests. Load Balance – IOPs (LB-IOPS) is the preferred Path Selection Scheme as recommended by the *PowerStore Host Configuration Guide*. In addition, the LB-IOPS path switching frequency should be changed from the default value of 1,000 to 1.

According to the *PowerStore Host Configuration Guide*, SSH to each ESXi host using root credentials to issue the following command (reboot required):

```

esxcli storage core claimrule add -u -t vendor --nvme-controller-
model "dell EMC powerstore" -P HPP -g "pss=LB-IOPS,iops=1"

```

For more information about NVMe-oF and the High Performance Plug-in, see the following resources on the VMware website:

- [VMware NVMe Concepts](#)
- [VMware High Performance Plug-In and Path Selection Schemes](#)
- [Requirements and Limitations of VMware NVMe Storage](#)

Operating system disk formats

While most versions of VMFS are backwards-compatible, it is a best practice to verify and use the latest version of VMFS recommended by VMware. Typically, new VMFS versions are bundled with an ESXi upgrade. As a migration path, VMware vCenter allows administrators to use VMware vSphere Storage vMotion to migrate virtual machines to new VMFS datastores formatted with the latest version.

Properties	
Name	Datastore01
File system	
Type	VMFS 6.82
Maximum file size	64 TB
Maximum virtual disk size	62 TB
Block size	1 MB
Space Reclamation Granularity	1 MB
Drive type	Flash

Figure 1. Example of datastore properties showing the VMFS version

Note: PowerStore X model appliances use VMware vSphere Virtual Volumes (vVols) storage containers for storage of virtual machines that run on internal PowerStore X model nodes. While it is not a best practice, block volumes can only be presented to the internal ESXi hosts in the PowerStore X model using the REST API.

NVMe over Fabric (NVMe-oF)

Introduction

NVMe over Fibre Channel support was introduced in vSphere 7.0 and PowerStoreOS 2.0. NVMe over TCP support was introduced with vSphere 7.0 Update 3 and PowerStoreOS 2.1.

NVMe-oF vVols

PowerStoreOS 3.0 introduced NVMe-vVol host connectivity supporting NVMe/FC vVols. NVMe-oF vVols is a new specification which introduces VASA 4.0 and vVols 3.0. This new specification requires HBAs and fabric switches that are NVMe capable, to extend the volumes from the array to the host. VMware added the corresponding NVMe-oF vVol support in vSphere 8.0.

Support for NVMe/TCP vVols was introduced with PowerStoreOS 3.6 and VMware vSphere 8.0 Update 1. NVMe/TCP yields high fibre channel-like performance at Ethernet prices. Customers of any size can benefit from this feature by delivering enterprise level performance for demanding application requirements.

Note: TCP ports 8009 and 4420 must be open for discovery and data respectively. In addition, the ESXi host time must be synchronized with PowerStore. Time synchronization can be automated using NTP or PTP for vSphere and NTP for PowerStore.




Figure 2. Enabling NVMe over TCP on a VMkernel port in the vSphere Client




Figure 3. Adding an NVMe over TCP adapter in the vSphere Client




Figure 4. Using PowerStore Manager to configure a PowerStore host with an NVMe vVol initiator

Create Storage Container [X]

Info: To access this storage container you must have VASA provider registered and then create vVol datastore on vCenter server.

LAUNCH VCENTER

Name: MySC

Enable storage container capacity quota

Container Quota: 54.5 TB

Select the storage protocol:

- SCSI (Supports iSCSI or FC transport layer)
- NVMe (Supports NVMe TCP and FC transport layers)

CANCEL **CREATE**

Figure 5. Using PowerStore Manager to create an NVMe vVol Storage Container

Note: With NVMe-oF vVols, there is no physical Protocol Endpoint (PE): the PE is now a logical object representation of the ANA group where the vVols reside. Until a VM is powered on, the vPE does not exist. When a VM is powered on, the vPE is created so the host can access the vVols in the ANA group. For more information, see [What's New with vSphere 8 Core Storage](#).

NVMe/FC host configurations

When configuring an ESXi host for NVMe/FC, before you add it to a PowerStore appliance or cluster, you must change the NVMe Qualified Name (NQN), which is similar to an iSCSI Qualified Name (IQN), to the UUID format.

According to the *PowerStore Host Configuration Guide*, SSH to each ESXi host using root credentials, and issue the following command (reboot required):

```
esxcli system module parameters set -m vmknvme -p
vmknvme_hostnqn_format=0
```

To verify the host NQN was generated correctly after the reboot, use the following command:

NVMe over Fabric (NVMe-oF)

```
esxcli nvme info get
```

Also, you must enable NVMe support on the NVMe-capable HBAs.

Per the *PowerStore Host Configuration Guide*, depending on the NVMe HBA installed, issue the following commands with root privileges (reboot required):

For Marvell NVMe HBAs:

```
esxcli system module parameters set -p ql2xnvmesupport=1 -m qlnativefc
```

For Emulex NVMe HBAs:

```
esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3
```

Note: You must change the host NQN format parameter **before** adding the host in PowerStore Manager. Changing the `vmknvme_hostnqn_format` parameter after the host has already been added to the appliance changes its NQN, which causes the host to be disconnected from the array.

NVMe over TCP (NVMe/TCP)

NVMe over TCP support was introduced with vSphere 7.0 Update 3 and PowerStoreOS 2.1. When planning to implement this new protocol, confirm that the host's networking hardware is supported in the [VMware Compatibility Guide](#).

This section provides a high-level overview of configuration best practices, but for more information, see the PowerStore resources on the [Dell Technologies Info Hub](#).

NVMe/TCP host configurations

Like with iSCSI, ESXi 7.0 U3 adds a software adapter for NVMe over TCP (Figure 6). After the software adapter is added, it becomes associated with a physical network adapter (Figure 7).




Figure 6. Adding an NVMe over TCP software adapter




Figure 7. Associating the adapter with a physical NIC

The best practice for storage network redundancy is to add two NVMe over TCP adapters and associate them with their respective storage network’s physical NICs (see the following figure).




Figure 8. NVMe over TCP storage adapters

After you add the storage adapters, you can configure the cluster networking. The best practice is to use a vSphere Distributed Switch (VDS) with two distributed port groups, one for each of the redundant storage networks (see the following figure).




Figure 9. Distributed port groups used for storage networks

Note: It is a best practice to use two storage fabrics for redundancy.

Since each NVMe over TCP storage adapter is bound to a physical NIC, you must adjust the Teaming and Failover for each distributed port group. Set the physical uplink that is bound to the vmhba to Active, and set the other NICs to Unused (see the following figure).




Figure 10. Teaming and failover settings for the distributed port group

Next, add the VMkernel adapters to their respective distributed port groups, and enable the NVMe over TCP service (see the following figure). These VMkernel adapters supply the IP addresses for each of the storage adapters (for example vmhba66 or vmhba67 as shown in Figure 8).




Figure 11. VMkernel adapter with NVMe over TCP service enabled

After you configure the host and cluster networking pieces, the dual storage networks should look like the example cluster shown in the following figure.




Figure 12. vSphere Distributed Switch topology view

After you complete the prerequisite networking configuration, add the storage controllers to discover the PowerStore array ports and IP addresses. You can add the storage controllers manually, by using direct discovery, or automatically by using the [SmartFabric Storage Software](#) (SFSS) as a Centralized Discovery Controller (CDC). PowerStoreOS 3.0 added enhancements to automate PowerStore registration with the SFSS/CDC. For more information, see the [SmartFabric Storage Software \(SFSS\) for NVMe over TCP – Deployment Guide](#).

After controller discovery, add the respective PowerStore front-end ports to each storage adapter (see the following figure). For example, add storage network 1 ports to vmhba66, and add storage network 2 ports to vmhba67. This process can be streamlined when using zoning capabilities with SFSS.

Storage Adapters

ADD SOFTWARE ADAPTER ▾ REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

Adapter	Model	Type	Status	Identifier	Targets	Devices
vmhba32	USB Storage Controller	Block SCSI	Unknown	--	1	2
vmhba66	VMware NVMe over TCP Storage Adapter	NVME over T...	Unknown	--	2	4
vmhba67	VMware NVMe over TCP Storage Adapter	NVME over T...	Unknown	--	2	4

EXPORT ▾ 9 items

Properties Devices Paths Namespaces Controllers

ADD CONTROLLER REMOVE

Name	Subsystem NQN	Transport Type	FUSE Support	Model	Firmware Version
nqn.1988-11.com.dell.powerstore:00:e29c3...	nqn.1988-11.com.dell.powerstore:00:e29...	tcp	true	dellemc-powerstore	2.1.0.0 TCP
nqn.1988-11.com.dell.powerstore:00:e29c3...	nqn.1988-11.com.dell.powerstore:00:e29...	tcp	true	dellemc-powerstore	2.1.0.0 TCP

Figure 13. PowerStore ports added to each NVMe storage adapter

Finally, add the ESXi hosts to PowerStore Manager before provisioning volumes. If everything is configured correctly, the host NQN should be associated with both VMK IPs as listed in the Transport Address field as shown in the following figure.

Add Host

Host Initiators

Select the initiators you want to use for configuring host access to the storage in this cluster. You can select from the automatically discovered initiators. To view a list of all available initiators, ensure that the FC embedded modules are configured properly and connected to the cluster.

ADD INITIATOR MORE ACTIONS ▾ 1 Initiators, 1 selected

Identifier	Transport Address	Transport Type	Origin	Initiator Type
nqn.2014-08...:nvme:s1241	10.10.132.3, 10.20.132.3	TCP	Auto-Discovered	NVMe

Figure 14. PowerStore Manager—Adding NVMe/TCP host with both VMK Ips

Note: If an ESXi host has been previously configured with NVMe/FC, set the `vmknvme_hostnqn_format=1` variable back to the hostname option before configuring NVMe/TCP. For more information, see the *PowerStore Host Configuration Guide* at Dell.com/powerstoredocs.

Sizing and performance optimization

Introduction

There are several best practices for provisioning storage from a PowerStore appliance to an external vSphere cluster. The size of VMFS datastores and the number of virtual machines that are placed on each datastore can affect the overall performance of the volume and array.

Virtual machines running in PowerStore X model internal nodes use internal vVol storage by default. However, the PowerStore X model can simultaneously run VMs on internal nodes and serve storage to external ESXi hosts. Therefore, the VMFS datastore sizing and virtual machine placement strategies in this section only apply when using PowerStore as storage to external ESXi hosts.

Volume and VMFS datastore sizing

When a volume is created on PowerStore T models, the best practice is to create a volume no larger than needed and use a single VMFS partition on that volume.

While the maximum datastore size can be up to 64 TB, we recommended beginning with a small datastore capacity and increase it as needed. Right-sizing datastores prevents accidentally placing too many virtual machines on the datastore and decreases the probability of resource contention. Since datastore and VMDK sizes can be easily increased if a virtual machine needs extra capacity, it is not necessary to create datastores larger than required. For optimal performance, the best practice is to increase the number of datastores rather than increase their size.

If a standard for the environment has not already been established, the recommended starting size for a VMFS datastore volume is 1 TB as shown in the following figure.

The screenshot shows the 'Create Volumes' wizard interface. On the left, there are three tabs: 'Properties' (selected), 'Host Mappings', and 'Summary'. The main area is divided into 'Properties' and 'Additional Properties'. Under 'Properties', there is a 'General' section with fields for 'Name (Or Prefix)' (Datastore01), 'Description (Optional)', 'Category' (Virtualization), and 'Application' (Virtual Server). Below these are 'Quantity' (1) and 'Size' (1 TB), where the 'Size' field is highlighted with an orange border. The 'Additional Properties' section includes 'Associated Volume Group (Optional)' (None Selected), 'Volume Protection Policy (Optional)' (Production), and 'Volume Performance Policy' (Medium). At the bottom right, there are 'CANCEL', 'FINISH', and 'NEXT' buttons.

Figure 15. PowerStore volume creation wizard

Increasing the size of VMFS datastores

You can increase the size of VMFS datastores in PowerStore Manager by modifying a volume's properties and increasing the size. After rescanning the storage adapters on the ESXi hosts, increase the VMFS partition size. Open the wizard, right-click the datastore, and select **Increase Datastore Capacity**. The best practice is to extend datastores using contiguous space within a single volume, and to avoid spanning volumes due to recovery complexity.

Note: The VSI plug-in can automate the process of increasing the size of datastores with only a few clicks.




Figure 16. Increasing the VMFS datastore size

Performance optimizations

While ESXi storage performance tuning is a complex topic, this section describes a few simple methods to proactively optimize performance.

Note: The VSI plug-in allows administrators to quickly set host best practices for optimal operation and performance.

Virtual machines per VMFS datastore

While the recommended number of virtual machines per VMFS datastore is subjective, many factors determine the optimum number of VMs that can be placed on each datastore. Although most administrators only consider capacity, the number of concurrent I/Os being sent to the disk device is one of the most important factors in the overall performance. The ESXi host has many mechanisms to ensure fairness between virtual machines competing for datastore resources. However, the easiest way to control performance is by regulating how many virtual machines are placed on each datastore. The best way to determine if a datastore has too many virtual machines is by monitoring disk latency with either `esxtop` or PowerStore Manager. If the concurrent virtual machine I/O patterns are sending too much traffic to the datastore, the disk queues fill, and higher latency is generated.

Note: The vVol datastores in PowerStore X model arrays do not have the same technical architecture as VMFS datastores. The virtual machine placement strategies that are described in this section are not necessary.

Modifying VMFS queue depth

To regulate and ensure fairness of I/O sent from VMs to each datastore, ESXi has an internal mechanism to control how many I/Os each virtual machine can send to the datastore at a time. This mechanism is Disk.SchedNumReqOutstanding (DSNRO). Although you can tune DSNRO for each datastore using `esxcli`, the best practice is to not modify this setting unless operating in a test environment or directed by support personnel.

Multiple-VM and single-VM-per-volume strategies

Although there are small performance, data service, and troubleshooting benefits to placing a single virtual machine on a VMFS datastore, placing multiple virtual machines on each VMFS datastore is common practice. Typically, using vVols achieves the same or better performance benefits compared to placing a single VM on a datastore.

There are two disadvantages when placing a single VM on its own datastore: it reduces consolidation ratios and increases the management overhead of maintaining numerous items.

Virtual machine affinity with Model X appliances

When using PowerStore Model X appliances in a multi-appliance cluster, remember that while the cluster may have multiple ESXi hosts, each virtual machine's virtual disks are bound to an individual appliance. This means that virtual machine storage performance may experience higher latency if the virtual machine is vMotioned to other appliances in the cluster. The best practice is to always keep the virtual machine running in the two hosts of the appliance that own the volumes. If the virtual machine must be relocated to another appliance in the cluster, you can nondisruptively move the virtual disks using the migration wizard in the PowerStore user interface. To prevent the VMware Distributed Resource Scheduler (DRS) from automatically migrating a virtual machine to non-optimal nodes in the cluster, assign the integrated VM/Host Rules in vCenter to any virtual machine. See the following figure.

The screenshot shows the vCenter interface for Cluster-PS-9, specifically the VM/Host Rules configuration page. The page is divided into a left sidebar with navigation options and a main content area. The main content area is titled 'VM/Host Rules' and contains a table of rules and a 'VM/Host Rule Details' section.

Name	Type	Enabled	Conflicts	Defined By
PS-9-appliance-1	Run VMs on Hosts	Yes	0	System
PS-9-appliance-2	Run VMs on Hosts	Yes	0	System
PS-9-appliance-3	Run VMs on Hosts	Yes	0	System

The 'VM/Host Rule Details' section shows the rule 'PS-9-appliance-1 VMs Group Members' with a list of hosts: W2019SQLServer, 100.88.130.94, and 100.88.130.95.

Figure 17. Virtual machine affinity rules

Partition alignment

Due to the PowerStore architecture, manual partition alignment is not necessary.

Guest vSCSI adapter selection

When creating a new virtual machine, vSphere automatically suggests the disk controller option based on the operating system selected (see the following figure). The *PowerStore Host Configuration Guide* recommends using the VMware Paravirtual SCSI controller for optimal performance. You can find more information about the Paravirtual adapter, including its benefits and limitations, in [VMware documentation](#).




Figure 18. Selecting the virtual SCSI controller

Array offload technologies

VMware can offload storage operations to the array to increase efficiency and performance. This action is performed by vStorage APIs for Array Integration (VAAI), a feature that contains primitives for both block and file storage types:

Block

- Write Same (Zero): Also known as Block zeroing. This primitive is primarily used for the ESXi host to instruct the storage to zero out eagerzeroedthick VMDKs.
- XCOPY (Extended Copy): Also known as Full copy. Instead of the ESXi host performing the work of reading and writing blocks of data, this primitive allows the

host to instruct the array to copy data which saves SAN bandwidth. This operation is typically used when cloning VMs.

- **Atomic Test & Set (ATS):** Also known as Hardware accelerated locking. This primitive replaces SCSI-2 reservations to increase VMFS scalability with changing metadata on VMFS datastores. With SCSI-2 reservations, the entire volume had to be locked, and all other hosts in the cluster had to wait while that ESXi host changed metadata. The hardware accelerated locking primitive allows a host to lock only the metadata on disk it needs, not hampering I/O from other hosts while the operation is performed.
- **UNMAP:** Also known as dead space reclamation. This primitive uses the SCSI UNMAP command to release blocks that are no longer in use back to the array. For example, after deleting a VM, the ESXi host issues a series of commands to the PowerStore array to indicate that it is no longer using certain blocks within a volume. This capacity is returned to the pool so that it can be reused.

File

- **Full File Clone:** Enables the offloading of powered-off virtual disk cloning to the array. Similar to XCOPY for block.
- **Fast File Clone/Native Snapshot Support:** Enables the creation of virtual machine snapshots to be offloaded to the array.
- **Extended Statistics:** Enables visibility into actual space usage on NAS datastores and is especially useful for thin-provisioned datastores.
- **Reserve Space:** Enables provisioning virtual disks using the Lazy Zeroed or Eager Zeroed options on NFS storage.

Note: For more information about VAAI, see [VMware vSphere APIs: Array Integration \(VAAI\)](#).

Management and monitoring

Introduction

This section describes PowerStore features used to manage and monitor storage.

Mapping or unmapping practices

After a volume is created, mapping specifies the hosts that the PowerStore array presents storage to.

Cluster mappings

For ESXi hosts in a cluster, we recommend using host groups to uniformly present storage to all initiators for reduced management complexity (see the following two figures). This practice allows a volume or set of volumes to be mapped to multiple hosts simultaneously and maintain the same logical unit number (LUN) across all hosts.

Note: It is required to use consistent LUN numbers for standard volume mappings and Metro Volume mappings across hosts within the same vSphere cluster, hosts within other vSphere clusters, or hosts not in a cluster. For additional information, please see the following references:

Dell KB Article 000191503 PowerStore: Inconsistent Logical Unit Numbers between hosts could result in data access or data consistency issues




Figure 19. Example of vSphere cluster containing multiple ESXi hosts




Figure 20. Host group details for the vSphere cluster showing three ESXi hosts

Properly unmapping volumes from ESXi hosts

If a VMFS datastore is no longer required, the best practice is to unmount the datastore from the vCenter (see the following figure). Then, detach the disk devices from **each** host in the cluster (see Figure 22) before unmapping and deleting the volume from the PowerStore Manager. This gracefully removes a datastore and can prevent an all paths down (APD) state from occurring.




Figure 21. Unmount the datastore




Figure 22. Detach the disk device from a host

Thin clones

PowerStore thin clones make block-based copies of a volume or volume group and can also be created from a snapshot. Because the thin clone volume shares data blocks with the parent, the capacity usage of the child volume mainly consists of the delta changes from after it was created. Thin clones are advantageous in a vSphere environment because a VMFS datastore full of virtual machines can be duplicated for testing purposes, all while consuming less storage. For example, if a vSphere administrator has to clone a multi-terabyte database server for a developer to run tests, the VM can be isolated and tested. Also, the VM only consumes blocks that changed.

Within the PowerStore architecture, thin clones have several advantages for storage administrators:

- The thin clone can have a different data protection policy from the parent volume.
- The parent volume can be deleted, and the thin clones become their own resource.
- VMs can be cloned for testing monthly patches or development.

Data encryption

Data at rest encryption (D@RE) is enabled by default on the PowerStore array. No configuration steps are necessary to protect the drives.

Space reclamation

The VAAI dead space reclamation primitive is integrated into the array through the SCSI protocol. Depending on the version of ESXi the host is running, the primitive can automatically reclaim space.

VMFS-6 and ESXi versions that support automatic UNMAP: The best practice is to keep or reduce the reclamation rate to the **Low** or **100 MB/s**, which is the default setting (see the following figure).




Figure 23. Setting the space reclamation rate

Note: vSphere 8 introduces the ability to configure the space reclamation rate to as low as 10 MB/s. This can be useful for environments where space reclamation at a higher rate can be disruptive to the storage fabric or its consumers.

VMFS-5 and ESXi versions that do not support automatic UNMAP: The best practice is to set the reclamation rate to 200.

```
esxcli storage vmfs unmap --volume-label=volume_label --reclaim-unit=200
```

Note: In certain older versions of ESXi, you must manually invoke the dead space reclamation primitive. See the [VMware Knowledge Base](#) for more information about which versions require additional steps.

VASA

VMware vSphere APIs for Storage Awareness (VASA) is a feature that allows vSphere hosts to gain insight into the storage types backing the datastores and enables vSphere to manage storage. For example, the VASA provider that is embedded into PowerStore allows it to manage vVols.

Note the following points about the PowerStore VASA provider:

- The provider is automatically registered on PowerStore X models.
- The provider must be registered when using vVols.
- The provider can be manually or automatically registered on PowerStore T models.
 - Registering the VASA provider is required for vCenter integration with PowerStore.
- Before you create the vCenter server connection, you must enter the VASA provider details on the vCenter Storage Providers screen in vCenter (see the following figure).
 - a. In PowerStore Manager, create an account named **vmadmin**, and assign the VMAdmin role.
 - b. In vCenter, add the storage provider, and specify the name, vmadmin credentials, and the URL in the format: **https://<management IP address>:8443/version.xml**
 - c. Do **not** select **Use storage provider certificate**.
 - d. For more information, follow the instructions in PowerStore Manager help by searching for **vCenter Server Connection**.

The screenshot shows the vSphere Client interface with the 'Config' tab selected for the 'tsvc05' storage provider. The interface displays a table of storage providers and their details.

Storage Provider/Storage Sys...	Status	Active/Sta...	Priority	URL	Last Rescan TL...	VASA API Versi...	Certificate Expi...
VMware vSAN	Online	--	--	https://tsvc0...	12/06/2019, 4...	3.0	--
PowerStore	Online	--	--	https://100.8...	12/09/2019, 6...	3.0	364 days
PSB (1/1 online)		Active	1				
IOFILTER Provider s1242	Online	--	--	https://s1242...	12/06/2019, 4...	1.5	1784 days
IOFILTER Provider s1241	Online	--	--	https://s1241...	12/06/2019, 4...	1.5	1784 days
IOFILTER Provider s1240	Online	--	--	https://s1240...	12/06/2019, 4...	1.5	1784 days

Below the table, the 'General' tab is selected, showing the following details:

- Provider name: PowerStore
- Provider status: Online
- Active/standby status: --
- Activation: Automatic
- URL: https://[redacted]:8443/version.xml
- Provider version: 1.0
- VASA API version: 3.0
- Default namespace: DELLEMCPowerStore VVOL
- Provider ID: 84fclaa8-345a-4c4d-s355-77d4a31fa450
- Supported profiles: Storage Profile Based Management

Figure 24. PowerStore VASA provider registration

Caution: The VASA certificate is set with a one-year expiration by default, and you should periodically renew it through the vCenter Storage Providers by clicking **Refresh Certificate**. If the certificate is not refreshed before expiration, see [KB 190731: How to renew the PowerStore VASA storage provider certificate after expiration](#).

Note: To use the PowerStore VASA provider across multiple vCenter Servers, there are two options: Use Enhanced Linked Mode or Share the vCenter root certificates across vCenter Servers. For more information, see [PowerStore: Using vVols across multiple vCenters: How to register the PowerStore VASA provider across multiple vCenters](#).

Virtual Volumes

VMware vSphere Virtual Volumes (vVols) are used as the primary storage mechanism for virtual machines running on PowerStore X model appliances internal nodes. They can be used by external ESXi hosts from either PowerStore X or T models. vVols is a new storage methodology that runs on top of existing storage protocols such as Fibre Channel and iSCSI. It enables administrators to have more granular control over virtual machines regarding performance, snapshots, and monitoring.

One of the key features of vVols is that it allows administrators to use storage policy-based management (SPBM) for their environment. This enables you to align application needs with the appropriate storage resources in an automated manner.

Because vCenter is required for binding and unbinding vVols from the protocol endpoints during power-on, power-off, and other operations such as vMotion, you should regard vCenter as a tier 1 application.

Caution: Never migrate the VMware vCenter virtual appliance to a vVol datastore or storage container. Because vCenter is required for bindings to power on vVol-based virtual machines, this action might prevent powering on vCenter after the VM is shut down or has experienced an unplanned outage.

Note: PowerStoreOS 3.0 introduced NVMe-vVol host connectivity supporting NVMe/FC vVols. PowerStoreOS 3.6 introduced NVMe-vVol host connectivity supporting NVMe/TCP vVols.

Scripting and automation

The PowerStore platform has a REST API and PowerShell cmdlets to automate management tasks. Find more information at [Dell Support](#).

Virtual Storage Integrator

Another tool for storage management is the Virtual Storage Integrator (VSI) which is the vSphere web client plug-in for PowerStore and other Dell Technologies storage products. You can download the VSI appliance as an OVA and install it into the virtual infrastructure. It is given an IP address and added to vCenter as part of the installation process. This plug-in allows administrators to create datastores, expand datastores, apply ESXi host best-practice settings, perform capacity monitoring, and more.




Figure 25. Virtual Storage Integrator vSphere web client plug-in

vRealize Orchestrator

PowerStore also supports VMware vRealize Orchestrator (vRO) as a plug-in. vRO enables administrators to automate many common workflows with PowerStore appliances. Download the vRO plug-in for PowerStore at [Dell Support](#).




Figure 26. PowerStore vRealize Orchestrator plug-in example

Data protection and disaster recovery

Introduction

PowerStore has integrated snapshot and replication capabilities to protect data, and it is policy driven for ease of administration.

Snapshots and recoveries

To automate and simplify protecting data, PowerStore uses protection policies. These policies are a set of snapshot and replication rules that are applied to a volume or group of volumes. Snapshot policies can also be applied to file systems and starting with

PowerStoreOS 3.0, replication policies can be assigned to NAS servers. Protection policies help protect data, set retention policies, and help guarantee recovery point objectives (RPOs) for an organization.

Also, protection policies can be applied to individual volumes or to volume groups. When a protection policy is applied to a volume group, it allows multiple volumes to have snapshots taken, to be replicated, or to be recovered, simultaneously. This ability allows protecting complex applications that are interdependent and span across multiple volumes.

You can take vVol snapshots from either the PowerStore Manager or the vCenter client, but they are inherently managed by vCenter. When you create virtual machine snapshots from the vCenter client, the best practice is to disable the option for virtual machine memory which may increase snapshot time significantly.

Secure snapshots

Starting with PowerStoreOS 3.5, the secure snapshot setting can be enabled for snapshots on volumes and volume groups. Snapshot rules can also be configured to create secure snapshots automatically. With secure snapshots enabled, the snapshots and parent resource are protected from accidental or malicious deletion and serve as a cost-effective line of defense against ransomware attacks. If an unauthorized user gains access to a system, the attacker cannot delete secure snapshots and cause data loss.

The figure displays two side-by-side screenshots of the 'Create Snapshot' workflow in PowerStore Manager. The left window is titled 'Create Snapshot of Volume' and shows settings for 'SQLDatabase'. The right window is titled 'Create Snapshot of Volume Group' and shows settings for 'DatabaseApp-VG'. Both windows have a 'Name' field with a timestamp, a 'Description' field, and 'Local Retention Policy' options: 'No Automatic Deletion' and 'Retain until' (selected). Below the retention policy is a 'Secure Snapshot' checkbox, which is highlighted with a red box in both screenshots. The text below the checkbox explains that secure snapshots prevent accidental or intentional deletion and cannot be manually deleted during the retention period. A 'Learn more' link is provided. At the bottom of each window are 'CANCEL' and 'CREATE SNAPSHOT' buttons.

Figure 27. Workflow to create a snapshot of a volume or volume group

Snapshots and options for application backup and restore

Using array-based snapshots is an effective way to protect virtual machine data and establish an RPO. In the PowerStore architecture, you can create the snapshot schedule using protection policies. Each protection policy can define snapshot rules to establish a schedule and retention, and replication rules to specify a destination array and RPO.




Figure 28. Protection policies screen in PowerStore Manager

PowerStore has data-recovery mechanisms that behave differently depending on the usage scenario (see the following figure).

- **Snapshot:** These represent a specific point in time for data stored on a volume that can be used for recoveries to refresh, restore, or create thin clones.
- **Thin clone:** This takes an existing snapshot from a parent volume and creates a child volume from that point in time.
- **Refresh:** This allows snapshot data to replace existing data in the volume. The existing data is removed, and snapshot data from the new source is copied to it in-place. A parent volume can refresh a child, and a child can refresh a parent.
- **Restore:** The restore operation replaces the contents of a parent storage resource with data from an associated snapshot. Restoring resets the data in the parent storage resource to the point in time the snapshot was taken.

Caution: Using the refresh and restore operations on active virtual machine volumes may cause unexpected results and behaviors. All host access to the volume must cease before attempting these operations.




Figure 29. PowerStore snapshot and recovery

If a virtual machine residing on a VMFS datastore requires recovery, the best practice is to create a thin clone from a snapshot. The high-level steps are as follows:

1. In PowerStore Manager, create a thin clone from a snapshot, and present it to the vSphere cluster.
2. In the vSphere client, rescan the storage, add a datastore, select the newly created volume, and assign a new signature (see the following figure).
3. Register the VM from the `snap-xxxxxxx-originaldatastorename` datastore.
4. Use Storage vMotion to migrate the virtual machine back to the original datastore, if applicable.




Figure 30. New datastore wizard > Assign a new signature

Crash consistent and application consistent snapshots

When taking array-based snapshots of virtual machines, remember that snapshots taken without application coordination are considered **crash consistent**. Crash consistency is the storage term for data that has a snapshot taken in-progress without application awareness. While most modern applications can recover from crash consistent data, their

recovery can yield varying levels of success. For example, when recovering a Microsoft Windows virtual machine, as the operating system boots, it responds as if it has encountered an unexpected power-loss event and can potentially check the disk (chkdsk) on startup.

Application consistent snapshots are supported by products such as Dell AppSync. This enables coordination between the array and the application to help assure that the data is quiesced, the caches are flushed, and the data is preserved in a known good state. Application consistent snapshots such as these offer a higher probability of recovery success.

Note: When taking managed snapshots such as with vVols in PowerStore Manager, virtual machine memory is not included. When performing snapshots from vCenter, we recommend not including virtual machine memory.

Replication and remote recovery

PowerStore offers asynchronous replication of block storage as well as asynchronous replication of NAS Servers and its underlying File Systems and NFS Exports as of PowerStoreOS v3.0. In addition, synchronous replication with Metro Volume support was also added as of PowerStoreOS v3.0. Replication is used to transfer data to one or more remote PowerStore clusters. When the remote cluster is in a different location than the local cluster, this feature can help to protect virtual machine data from localized geographical disasters. Replication RPOs and options are set within protection policies (see the following figure).

Figure 31. Creating a replication rule

Note: The document [Dell PowerStore: Replication Technologies](#) provides comprehensive information about replication between PowerStore clusters. See this document for more details about replication.

**PowerStore
backup to
PowerProtect DD
series
appliances**

Starting with PowerStoreOS 3.5, PowerStore can back up virtual machine data to a physical PowerProtect appliance or to a PowerProtect DD Virtual Edition (DDVE) either on premises or in the cloud. This enables a simplified native backup solution for volumes and volume groups and the virtual machine data contained within them. No backup application is required. By configuring remote backup sessions from PowerStore Manager, users can quickly back up resources, retrieve remote snapshots, and initiate instant access sessions. Instant access allows the host to view the content of a remote snapshot without retrieving the data back to the PowerStore system. Users can instantly access deleted, corrupted, or modified data within the snapshot and copy it back to the host for quick recovery. Users can also discover and retrieve snapshot backups created on a different PowerStore cluster.

References

Dell Technologies documentation

The following Dell Technologies documentation provides other information related to this document. Access to these documents depends on your login credentials. If you do not have access to a document, contact your Dell Technologies representative.

See the following documents on the [PowerStore Info Hub](#):

- Dell PowerStore Host Configuration Guide
- Dell PowerStore Protecting Your Data
- Dell PowerStore Replication Technologies
- Dell PowerStore Snapshots and Thin Clones
- Dell PowerStore: Synchronous Replication and Metro Volume

VMware documentation

See the following links for related VMware resources:

- [VMware Documentation](#)
- [VMware Hardware Compatibility List \(HCL\)](#)
- [VMware vSphere APIs: Array Integration \(VAAI\)](#)
- [What's New with vSphere 8 Core Storage](#)