

# Dell PowerStore: Virtualization Integration

October 2022

H18152.9

## White Paper

### Abstract

This document discusses the virtualization features and integration points between the Dell PowerStore platform and VMware vSphere.

Dell Technologies

## 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 © 2020-2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Intel, the Intel logo, the Intel Inside logo and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Other trademarks may be trademarks of their respective owners.

Published in the USA [October 2022 H18152.9](#).

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

# Contents

<b>Executive summary .....</b>	<b>4</b>
<b>Introduction .....</b>	<b>5</b>
<b>vCenter connection.....</b>	<b>7</b>
<b>Internal and external ESXi hosts .....</b>	<b>10</b>
<b>vSphere Virtual Volumes .....</b>	<b>11</b>
<b>Virtual machines .....</b>	<b>18</b>
<b>VMware datastores .....</b>	<b>26</b>
<b>PowerStore X models .....</b>	<b>27</b>
<b>Migration .....</b>	<b>38</b>
<b>Metro Volume .....</b>	<b>39</b>
<b>vVol replication .....</b>	<b>39</b>
<b>VMware plug-ins.....</b>	<b>39</b>
<b>Conclusion.....</b>	<b>42</b>
<b>Appendix: Technical support and resources .....</b>	<b>43</b>

## Executive summary

### Overview

Virtualization offers many benefits such as consolidation, performance, availability, business continuity, load balancing, and ease of maintenance. Many applications are being virtualized today because of these advantages. It is important for data-center components not only to support, but also provide integration with, hypervisors and virtualized applications. This document details the many virtualization features and integration points that are available on Dell PowerStore.

### Audience

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

### Revisions

Date	Description
April 2020	Initial release: PowerStoreOS 1.0.0
August 2020	Minor updates
September 2020	Minor updates
December 2020	PowerStore 1.0.3 updates
April 2021	PowerStoreOS 2.0.0 updates
May 2021	Minor updates
January 2022	PowerStoreOS 2.1.0 updates; template update
April 2022	PowerStoreOS 2.1.1 updates
June 2022	PowerStoreOS 3.0.0 updates
October 2022	PowerStoreOS 3.2.0 updates

### 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:** Wei Chen

**Contributors:** Ethan Stokes, Stephen Granger

---

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

---

# Introduction

## Overview

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 offers investment protection through flexible payment solutions and data-in-place upgrades.

## PowerStore virtualization integration

PowerStore features multiple integration points with VMware vSphere virtualization technology that is used in data centers today. Many of these powerful integration points are embedded in the system and are designed with the end-user experience in mind. They can be easily managed directly from the HTML5-based PowerStore Manager user interface. In addition to the integration points that are built into the system, off-array software and plug-ins are available. These plug-ins enable PowerStore to be used with existing tools and fit the specific requirements of each organization. Storage and virtualization administrators can use these features to create simple, modern, flexible, and affordable solutions.

PowerStore is offered as a PowerStore T model or PowerStore X model appliance. Both models are designed to have deep integration with VMware vSphere. These integrations include VAAI and VASA support, event notifications, snapshot management, storage containers for VMware vSphere Virtual Volumes (vVols), and virtual machine discovery and monitoring in PowerStore Manager.

PowerStore X models provide flexibility and agility by providing AppsON functionality. This ability enables administrators to run applications directly on the storage system. Due to the embedded VMware ESXi hypervisor on the PowerStore X model nodes, other virtualization features and automation for the configuration process are available on this model. The vSphere hypervisor is embedded on each of the PowerStore X model nodes which allows applications to run directly on the PowerStore appliance. Simultaneously, it can be used as a standard external storage array, providing block-volume access to servers over Fibre Channel, iSCSI, or NVMe-oF.

## Terminology

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

**Table 1. Terminology**

Term	Definition
AppsON	A PowerStore X model appliance feature that enables running applications as virtual machines directly on PowerStore storage and compute. This integration brings applications closer to storage.
Controller VMs	Virtual machines that run a virtualized version of the PowerStoreOS on PowerStore X model appliances. Each PowerStore X model node has its own controller VM. Each controller VM reserves 50% of the available CPU and memory on the appliance, leaving the other 50% for user VMs.
Distributed Resource Scheduler (DRS)	A VMware feature that monitors resource utilization and spreads virtual machine workloads across ESXi hosts in a cluster.
Fibre Channel (FC) protocol	Protocol used to perform Internet Protocol (IP) and SCSI commands over a Fibre Channel network.
Internet SCSI (iSCSI)	A mechanism that provides access to block-level data storage over network connections.
PowerStore Manager	An HTML5 user interface used to manage PowerStore systems.
Storage container	A VMware term for a logical entity that consists of one or more capability profiles and their storage limits. This entity is known as a VMware vSphere Virtual Volumes (vVols) datastore when it is mounted in vSphere.
Storage Policy Based Management (SPBM)	Policies used to control storage-related capabilities for a VM and ensure compliance throughout its life cycle.
User VM	A virtual machine that is deployed by the administrator. A user VM can be using PowerStore storage with external compute hosts. A user VM can also be deployed using PowerStore storage and internal compute hosts by using AppsON.
Virtual machine (VM)	An operating system running on a hypervisor, which is used to emulate physical hardware.
vCenter	A VMware server that provides a centralized platform for managing VMware vSphere environments.
VMware vSphere Virtual Volumes (vVols)	A VMware storage framework that allows VM data to be stored on individual Virtual Volumes. This ability allows for data services to be applied at a VM level of granularity and according to SPBM. Virtual Volumes can also refer to the individual storage objects that are used to enable this functionality.
vSphere API for Array Integration (VAAI)	VMware APIs that improve ESXi host utilization by offloading storage-related tasks to the storage system.
vSphere APIs for Storage Awareness (VASA)	VMware vendor-neutral APIs that enable vSphere to determine the capabilities of a storage system. This feature requires a VASA provider on the storage system for communication.

Term	Definition
vSphere cluster	A group of ESXi hosts that are grouped to enable high availability, load balancing, and resource management.
vSphere datacenter	A container that consists of hosts, clusters, and other objects required to operate virtual machines.
vSphere Remote Office Branch Office (ROBO)	A VMware license with a limit of 25 virtual machines.

## vCenter connection

### Overview

To enable virtual machine (VM) discovery, monitoring, and snapshot management, the vCenter server must be registered in PowerStore Manager. This step enables PowerStore to monitor the VM attributes, capacity, storage and compute performance, and virtual volumes. It also enables PowerStore to subscribe to event notifications, alleviating the need for PowerStore to poll continuously for new information.

On PowerStore X model appliances, a vCenter server connection is required as part of the initial configuration process. This connection enables the VASA provider registration and vVol datastore creation to happen automatically. This ability allows users to begin using vVols immediately after deploying the system without any additional setup. PowerStore X models require the vCenter to be hosted on an external server.


On PowerStore T models, a vCenter server connection is optional. Starting with PowerStoreOS 2.0, the initial configuration includes a step that lets you configure the vCenter server connection. This connection is required on PowerStore X model appliances.

On PowerStore T model appliances, a vCenter server can also be connected after the initial configuration. To establish a vCenter server connection, open PowerStore Manager and browse **Compute > vCenter Server Connection**. You can connect a vCenter by entering in the **vCenter Server IP Address** (or FQDN), **User Name**, and **Password** for an existing vCenter server.

PowerStore T models can connect to any vCenter that is running vCenter version 6.0 Update 2 or later. For the supported vCenter versions on PowerStore X models, see the *PowerStore: Simple Support Matrix* at [Dell.com/powerstoredocs](http://Dell.com/powerstoredocs).


Starting with PowerStoreOS 2.0, you can manage and monitor VASA registration from PowerStore Manager. This ability removes the need to log in to vSphere to view or update to the VASA registration. When registering a new vCenter, there is an option to provide the PowerStore credentials. If provided, these credentials are used to automatically register the VASA provider in vSphere. The credentials must be for an account with the VM Administrator, Storage Administrator, or Administrator role.

The following figure shows the vCenter Server registration page:



**Figure 1. Registering a vCenter server**

After a successful vCenter server connection is made, the IP address or hostname of the connected vCenter is displayed and the status changes to **Connected**. Starting with PowerStoreOS 2.0, the **VASA registration status** is displayed on this page. If the VASA provider was not connected during vCenter registration or becomes disconnected, this state is reflected in the status. Buttons to **Launch vSphere**, **Update Connection**, and **Disconnect** (PowerStore T model only) also become available, as shown in the following figure:




**Figure 2. vCenter connected**

Clicking **Launch vSphere** opens a new tab to the connected vCenter. This feature allows the administrator to browse the vCenter easily.

Use the **Update Configuration** button to update the connection with new information if the vCenter IP address, hostname, or credentials change. Each PowerStore cluster can only be registered to a single vCenter instance at a time. Do not use the update button to connect the PowerStore cluster to a separate vCenter instance. On PowerStore T models, the vCenter connection can be disconnected and then connected to the new vCenter instance. On PowerStore X models, you cannot change the vCenter connection to another vCenter instance. This limitation is due to the existence of vSphere objects such as the data center, cluster, PowerStore X model ESXi nodes, virtual distributed switches, and other configurations on the vCenter.

Starting with PowerStoreOS 2.0, you can use the **Update Configuration** button to manage the VASA registration status. For example, if the VASA provider is accidentally deleted in vSphere, the VASA registration status changes to **Not configured**. In this scenario, you can use the **Update Configuration** button to re-register the VASA provider directly from PowerStore Manager. If the VASA registration status is **Online**, the administrator is not prompted for the PowerStore credentials. The following figure shows the dialog box for updating the vCenter Server Configuration:



**Figure 3. Updating the configuration**

Use the **Disconnect** button to remove a vCenter connection. This feature is only available on PowerStore T models because the vCenter connection is mandatory on PowerStore X models. Starting with PowerStoreOS 2.0, the administrator has the option to remove the VASA provider registration when disconnecting a vCenter server. The following figure shows the confirmation dialog box that is displayed when you disconnect the vCenter server:




Figure 4. Disconnect vCenter Server confirmation dialog box

## Internal and external ESXi hosts

### Overview

External hosts can be registered in PowerStore Manager to enable access to storage resources. Internal hosts are the PowerStore X model ESXi nodes. Starting with PowerStoreOS 2.0, enhancements are added to display additional details and improve visibility for both internal and external ESXi hosts. These enhancements include:

- PowerStore X model ESXi node visibility in PowerStore Manager:
  - The internal ESXi nodes are shown alongside the external hosts that are registered on the cluster.
  - The **Host Type** column indicates whether the host is internal or external.
  - These internal ESXi nodes are visible on the **Hosts & Host Groups** and the **Storage Containers > ESXi Hosts** pages.
- vSphere hostname:
  - For internal and external ESXi hosts, the vSphere **Host Name** column displays the hostname that is displayed in the vSphere Web Client.
  - The **Host Name** column enables administrators to easily identify the host, even if it is registered with different names in PowerStore Manager and vSphere.
  - This column is on multiple pages within PowerStore Manager (**Hosts & Host Groups**, **Virtual Machines**, **Virtual Volumes**, and so on).
- ESXi version:
  - For internal and external ESXi hosts, the ESXi version is also displayed in PowerStore Manager.
  - PowerStore X model ESXi nodes in a PowerStore X cluster must all run the same ESXi version.

- The **ESXi Version** column is on the **Hosts & Host Groups** page.

The following figure shows the enhanced **Hosts & Host Groups** page.

Name	vSphere Host Name	Host/Host Group	Host Type	OS	Initiator Type	Initiators	Volume Mappings	ESXi Version
Appliance-WX-H6209-nod--	10.245	Host	Internal	ESXi	iSCSI	1	--	VMware ESXi 6.7.0.17167...
Appliance-WX-H6209-nod--	10.245	Host	Internal	ESXi	iSCSI	1	--	VMware ESXi 6.7.0.17167...
ESXi	10.245	Host	External	ESXi	iSCSI	1	--	VMware ESXi 6.7.0.14320...

**Figure 5. Internal ESXi host visibility, vSphere Host Name, and ESXi version visibility**

## vSphere Virtual Volumes

### Overview

PowerStore supports the VMware vSphere Virtual Volumes (vVols) framework through the VASA 3.0 protocol. This feature enables VM-granular data services and Storage Policy Based Management (SPBM). In traditional storage environments, volumes or file systems are formatted as VMFS or NFS datastores for VMs. Data services are applied at the volume or file-system level, which means all VMs that reside on that datastore are also affected.

With vVols, VM data is stored on dedicated storage objects that are called storage containers, which become vVol datastores in vSphere. A VM consists of multiple vVols depending on its configuration and status. PowerStore works with vSphere to track which vVols belong to which VM.

For data services, such as VM snapshots and clones, you can apply them at a VM-level of granularity because they are only applied to the relevant vVols. These data services are offloaded to PowerStore to maximize efficiency. Policies and profiles can be used to ensure that VMs are provisioned with the required storage capabilities.

### VASA provider

vSphere APIs for Storage Awareness (VASA) are VMware-defined and vendor-neutral APIs that enable vSphere to determine the capabilities of a storage system. The API requests basic storage information from PowerStore and uses it for monitoring and reporting storage details to the user in vSphere.

PowerStore includes a native VASA 3.0 provider, which enables the vVols storage framework. The VASA provider must be registered in vSphere in order to use vVols. On

PowerStore X models, the storage provider is registered in vSphere automatically as part of the initial configuration process.

On PowerStore T models, starting with PowerStoreOS 2.0, the storage provider can be optionally registered during the initial configuration process. After initial configuration is completed, this registration can be done as part of the vCenter server connection process in PowerStore Manager or manually registered in vSphere.

- To register the VASA provider directly from PowerStore Manager, browse to **Compute > vCenter Server Connection**.
- To register the VASA provider in vSphere, browse to **vCenter > Storage Providers > Configure**. Click **Add** and provide the information below, as shown in Figure 6.
  - Name: <name>
  - URL: https://<Cluster\_IP>:8443/version.xml
  - Username: User with administrator or VM administrator privileges
  - Password: <password>




Figure 6. New Storage Provider page

After a storage provider is successfully registered, additional details about the provider are displayed, as shown in the following figure:




Figure 7. Registered storage provider

## Storage containers

A storage container is used to present vVol storage from PowerStore to vSphere. vSphere mounts the storage container as a vVol datastore and makes it available for VM storage. When using AppsON, user VMs should **only** be provisioned on the vVol datastores. User VMs should **never** be provisioned on the PowerStore X model private datastores because those datastores are reserved for the controller VMs. PowerStore includes a default storage container that is named PowerStore <Cluster\_Name>, as shown in the following figure:





Figure 8. Default storage container

On PowerStore X models, the default storage container is mounted automatically on the internal ESXi nodes. PowerStore can also expose its storage containers to external ESXi hosts, enabling VM provisioning on external compute with PowerStore vVol storage. This functionality can be enabled as follows:


1. Register the PowerStore VASA provider (see [VASA provider](#)).
2. Establish iSCSI, Fibre Channel, or NVMe/FC connectivity between the external ESXi host and PowerStore.
3. Register the host as ESXi, and select its initiators in PowerStore Manager.
4. Initiate a rescan in vSphere.
5. Add the storage container as a vVol datastore in vSphere.

After Step 4, two protocol endpoints are automatically created on the ESXi host. These protocol endpoints are identified with LUN IDs 254 and 255 on the Storage Devices page, as shown in the following figure:



**Figure 9. Protocol endpoints with LUN IDs 254 and 255**

All registered ESXi hosts are automatically granted access to all the storage containers on PowerStore. These ESXi hosts can mount the datastore in vSphere after host connectivity is established and no further mapping is needed. The following figure shows the vVol datastore mounted in vSphere.




**Figure 10. PowerStore vVol datastore**

In addition to the default storage container, additional storage containers can also be created. On PowerStore X models, these additional storage containers are mounted automatically to the internal ESXi nodes. On PowerStore T models, these additional storage containers can be mounted as vVol datastores in vSphere.


By default, a storage container exposes all the free capacity available on the cluster. Storage containers can be configured with a quota to expose less or more storage to

vSphere. When configuring a quota on an existing storage container, a high-water mark can also be configured. When the utilization of the storage container exceeds the high-water mark, the system generates a notification. If the utilization falls below the high-water mark, the notification clears automatically. By default, the high-water mark is set to 85% and this setting is user configurable. The following figure shows setting a quota of 5 TB and a high-water mark of 85%:



**Figure 11. Storage container quota settings**

If a quota is set on an existing storage container, the size is not immediately updated in vSphere. To force a refresh, right-click the datastore and click **Refresh Capacity Information**. Alternatively, the capacity refreshes automatically every 15 minutes. The following figure shows the updated capacity on the vVol datastore after the quota is applied:



**Figure 12. vVol datastore capacity with quota**

With a multi-appliance cluster, the cluster creates a single storage container that exposes all storage from all appliances within the cluster. When a VM is provisioned on the storage container, resource balancer determines which appliance within the cluster its vVols are stored on. You can determine which appliance a vVol resides on by looking at the Virtual Volumes card within the VM or storage container properties page. vVols can also be migrated between appliances on-demand.

### Storage container protocol

With the introduction of PowerStoreOS 3.0, PowerStore supports the creation of either SCSI or NVMe storage containers. Before this release, all storage containers were SCSI by default. SCSI storage containers support host access through SCSI protocols, which include iSCSI or Fibre Channel. NVMe storage containers support host access through the NVMe/FC protocol.

When creating a storage container on a system running PowerStoreOS 3.0 or later, you can select either **SCSI (Supports iSCSI or FC transport layer)** or **NVMe (Supports NVMe FC transport layer)**. This selection specifies the protocol type for that storage container, and any hosts which mount the storage container as a vVol datastore must have appropriate connectivity and support.

Create Storage Container

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

LAUNCH VCENTER

Name  
NVMe-SC

Enable storage container capacity quota

Container Quota  
31.5 TB

Select the storage protocol

SCSI (Supports iSCSI or FC transport layer)

NVMe (Supports NVMe FC transport layer)

**Figure 13. Storage container protocol selection**

On the **Storage Containers** page in PowerStore Manager, a new column introduced in PowerStoreOS 3.0, **Storage Protocol**, is displayed by default. This column details the supported storage protocol for a given storage container. A storage container can either be SCSI or NVMe; there is no support for both protocols on the same storage container. This new feature has no impact on existing storage containers, which are all classified as SCSI.

<input type="checkbox"/> Name	Storage Protocol	Used Space	Quota	Total Space	Free Space
<input type="checkbox"/> NVMe-SC	NVMe	0 GB	0 GB	31.5 TB	31.5 TB
<input type="checkbox"/> SCSI-SC	SCSI	0 GB	0 GB	31.5 TB	31.5 TB

**Figure 14. Storage container storage protocol column**

Storage containers can convert their storage protocol between the two types, although this operation is disruptive. You must remove or unbound all vVols on the storage container. This process requires stopping all virtual machines on the associated vVol datastore, or use vSphere Storage vMotion to move all virtual machines and vVols to a different storage resource temporarily. Then, from the **Storage Containers** page in PowerStore Manager, select the storage container and click **MODIFY**. Complete the process by selecting the new protocol and clicking **APPLY**. Currently, virtual machines can be restarted or moved back onto the vVol datastore through vSphere Storage vMotion.

## Storage Policy Based Management

vVols use Storage Policy Based Management (SPBM) to ensure VMs have the appropriate storage capabilities through their entire life cycle. VM storage policies can be optionally created after the storage provider is registered. These policies are used to determine the required storage capabilities when a VM is being provisioned.

To create a storage policy, go to the **Policies and Profiles > VM Storage Policies** page in vSphere. Click **CREATE**, and then select **Enable rules for “Dell EMC PowerStore” storage**.

The QoS Priority rule determines the relative performance prioritization for the VM if the system experiences resource contention. You can select **HIGH**, **MEDIUM**, or **LOW** as the QoS Priority.

The Snapshot Schedule rule allows PowerStore to take snapshots of virtual machines at a given frequency. The Snapshot Schedule rule when creating a VM Storage Policy automatically displays all snapshot rules created on PowerStore. If you want to assign a Snapshot Schedule rule, you must create the snapshot rules on PowerStore before you create the VM Storage Policy in vSphere. The following figure shows the available PowerStore rules when you create a storage policy:




Figure 15. Create VM Storage Policy page

## Virtual machines

### Overview

VMs that are stored on PowerStore vVol datastores are automatically discovered and displayed in PowerStore Manager. All VMs stored on the vVol datastores are displayed. This listing includes VMs using internal compute on PowerStore X and external compute on an ESXi server. This page includes a list of VMs including the name, operating system, CPUs, memory, and more, as shown in the following figure:




Figure 16. Virtual Machines page

Click each VM to view more details such as capacity, compute and storage performance, alerts, protection, and virtual volumes for that VM. See the following figure:




Figure 17. VM storage performance

Starting in PowerStoreOS 3.0, you can determine the type of backing storage for virtual machines with the new **Datastore Type** column (see the figure below). This column shows if the virtual machine is deployed entirely onto NFS, VMFS, or vVol storage hosted on the PowerStore. If the virtual machine contains storage from two or more storage types, this column displays the Datastore Type **mixed**.

Name	Alerts	Power State	Datastore Type	vSphere Host Name	Datastore	Guest OS
vm windows-mixed	-	Powered On	Mixed	10.245.11.101	psd-ds-nfs, psd-ds-vmfs, psd-ds-v...	Microsoft Windows Server 2016 or later (64...
vm windows-nfs	-	Powered On	NFS	10.245.11.101	psd-ds-nfs	Microsoft Windows Server 2016 or later (64...
vm windows-vmfs	-	Powered On	VMFS	10.245.11.101	psd-ds-vmfs	Microsoft Windows Server 2016 or later (64...
vm windows-vvol	-	Powered On	vVol	10.245.11.101	psd-ds-vvol	Microsoft Windows Server 2016 or later (64...

Figure 18. VM datastore type


Virtual machines deployed onto storage classified as Datastore Type mixed will only contain the compute performance and virtual volumes tabs when viewing the details. The capacity, storage performance, alerts, and protection tabs are not available to these virtual machines. Virtual machines deployed onto storage classified as Datastore Type NFS or VMFS will only show the computer performance tab.

## Protection

The **Protection** card enables administrators to manage snapshots and configure protection policies for a VM. This page enables you to create a manual snapshot or modify and delete existing snapshots. Before PowerStoreOS 3.0, a protection policy could also be applied to the VM to take snapshots automatically, such as for volumes and file systems. With the release of PowerStoreOS 3.0, snapshot schedules are only applied to a

virtual machine through vSphere using VM Storage Policies. See the section [Storage Policy Based Management](#) for more details about VM Storage Policies.


The following figure shows the VM protection page where you can configure snapshots and protection policies:



**Figure 19. VM protection**

VM snapshots are visible in both PowerStore Manager and vCenter, regardless of where they are created. You can view information about VM snapshots in the **Manage Snapshots** page in vCenter. You can also initiate a revert operation from here to restore the VM using the snapshot. You can revert to any snapshot in the snapshot tree.

Snapshots that are taken from PowerStore **do not** include the guest VM memory. This behavior means that the VM memory contents and power state are not preserved, but the snapshot is crash consistent. After the snapshot restore operation completes, the VM reverts to a powered-off state and can be powered back on. The following figure shows a VM with manual and scheduled snapshots that are created from PowerStore:



**Figure 20. VM snapshots**

vSphere enforces a limit of **31** snapshots for each VM. If this limit is reached, the oldest snapshot is automatically deleted chronologically starting with the oldest when the next snapshot is created by the policy. Although manually created snapshots count towards this limit, they are never automatically deleted because they do not have an expiration date.

In large environments, it is possible to initiate many snapshot requests to vCenter at once. To prevent overloading vCenter, PowerStore sends a maximum of **five** simultaneous create snapshot operations to vCenter. The remaining operations are queued and started as each create snapshot operation completes. PowerStore also sends a maximum of **five** simultaneous delete snapshot operations to vCenter. Although the create snapshot operations are sent individually, delete snapshot operations can be sent in batches, up to a limit of five. Because these two limits are different, it is possible to have a total of five create and five delete snapshot operations simultaneously on different VMs.

For more information about snapshots and protection policies, see the document [PowerStore: Snapshots and Thin Clones](#).

## Virtual Volumes

The type of vVol provisioned depends on the type of data that is being stored:

- **Data:** Stores data such as VMDKs, snapshots, full clones, and fast clones. At least one data vVol is required per VM to store its hard disk.
- **Config:** Stores standard VM configuration data such as .vmx files, logs, and NVRAM. At least one config vVol is required per VM to store its .vmx configuration file.
- **Swap:** Stores a copy of the VM memory pages when the VM is powered on. Swap vVols are automatically created and deleted when VMs are powered on and off. The swap vVol size matches the VM memory size.
- **Memory:** Stores a complete copy of VM memory on disk when suspended, or for a with-memory snapshot.

At a minimum, three vVols are required for each powered-on VM: **data** for the hard disk, **config** for the configuration, and **swap** for the memory pages.

The **Virtual Volumes** card provides details about the vVols used for the VM. PowerStore uses the VASA protocol to communicate with vSphere to create, bind, unbind, and delete vVols automatically, as needed. Manual management of these vVols is not required. This page also provides options to migrate vVols, manage the Watchlist, and collect support materials.

Information such as the vVol name, type, capacity, storage container, appliance, and I/O priority are displayed, as shown in the following figure:




Figure 21. Virtual Volumes

### Virtual Volume migration


In PowerStoreOS 1.0, vVols can be migrated between appliances within the cluster. However, it is limited to vVols that are not in use so the virtual machine must be powered off before any of its vVols can be migrated. Starting with PowerStoreOS 2.0, online vVol migration is supported. This functionality allows vVols that are used for powered-on virtual machines to be migrated between appliances within the cluster.

In order to support online vVol migration, the ESXi host must be running VMware ESXi 6.7 P02 or higher. Previous versions of VMware ESXi do not support online vVol migration as this functionality requires ESXi vVol rebind orchestration. In this scenario, the vVol must be unbound manually by powering off the virtual machine or the ESXi host must be upgraded to the appropriate version.

The online migration operation is transparent to the virtual machine and no rescans are required. Like volume migrations, both manual and assisted migrations are available for vVols. The migration traffic flows over the first two ports of the four-port card using the Intra-Cluster Management (ICM) and Intra-Cluster Data (ICD) IPv6 networks.

It is possible to have multiple vVols for a single virtual machine spread across multiple appliances. The best practice recommendation is to have all vVols for a virtual machine on the same appliance. Online vVol migration can be used as a nondisruptive method to consolidate a virtual machine's vVols onto a single appliance.

vVol migrations can be initiated from the **VM Details > Virtual Volumes** or **Storage Container Details > Virtual Volumes** pages. The following figure shows the migrate operation:



**Figure 22. vVol migration**

Here is the workflow for an online vVol migration:

1. Administrator creates a migration session. The system creates a connection between the source and destination appliances
2. Initial sync: The source vVols data, fast clones, and snapshots are migrated to the destination.
3. Delta sync and nondisruptive cutover.
  - a. A final delta copy is completed.
  - b. PowerStore and ESXi coordinate to perform rebind events, enabling an automated and nondisruptive cutover to the new appliance.

For more information about manual and assisted migrations, see the document [Dell PowerStore: Clustering and High Availability](#).

## vVol storage and virtual machine compute

To enable optimal virtualization performance, accounting for a virtual machine's compute and storage placement is important. This section provides recommendations for when you are using PowerStore storage with external compute and internal compute (AppsON).

### vSphere Virtual Volumes storage with external compute

For optimal performance, keep all vVols for a VM together on a single appliance. When provisioning a new VM, PowerStore groups all its vVols onto the same appliance. In a multi-appliance cluster, the appliance that has the highest amount of free capacity is selected. This selection is maintained even if the provisioning results in a capacity imbalance between appliances afterwards. If all vVols for a VM cannot fit on a single appliance due to space, system limits, or health issues, the remaining vVols are provisioned onto the appliance with the next-highest amount of free capacity.

When provisioning a VM from a template or cloning an existing VM, PowerStore places the new vVols onto the same appliance as the source template or VM. This action

enables the new VM to take advantage of data reduction to increase storage efficiency. For VM templates that are frequently deployed, it is recommended to create one template per appliance and evenly distribute VMs between appliances by selecting the appropriate template.

When taking a snapshot of an existing VM, new vVols are created to store the snapshot data. These new vVols are stored on the same appliance as the source vVols. In situations where the source vVols are spread across multiple appliances, the vVols created by the snapshot operation also become spread. vVol migrations can be used to consolidate a VM's vVols on to the same appliance.

In this configuration, PowerStore provides storage and an external hypervisor provides compute. The external hypervisor connects to the storage system through an IP or FC network. Since the external hypervisor always travels through the SAN to communicate with the storage system, no further considerations are needed for compute placement.

### **vVol storage with internal compute (AppsON)**

On PowerStore X model appliances, AppsON enables customers to run their applications using the internal ESXi nodes. When using AppsON, using the same appliance for a virtual machine's storage and compute minimizes latency and network traffic. In a single appliance cluster, compute and storage for AppsON VMs are always collocated, and no further considerations are needed for compute placement.

Starting with PowerStoreOS 2.0, PowerStore X model appliances can be configured in a PowerStore cluster. Clustering improves ease of management by providing a single point of management and enables easy migration of volumes and vVols between appliances within the cluster. As of PowerStoreOS 3.2, multi-appliance clustering on PowerStore X is no longer supported.

When a multi PowerStore X cluster is configured, this action also creates an ESXi cluster in vSphere with all the PowerStore X model ESXi nodes. From a vSphere perspective, each PowerStore X model ESXi node is weighted equally so it is possible for a VM's storage and compute to be separated. This configuration is not ideal because it increases latency and network traffic. For example, if a virtual machine's compute is running on Node A on appliance-1 but its storage resides on appliance-2. Then, I/O must traverse through the top-of-rack (TOR) switches for the compute node to communicate with the storage appliance.

For optimal performance, it is recommended to keep all vVols for a VM together on a single appliance. When provisioning a new VM, PowerStore groups all its vVols onto the same appliance. This grouping is maintained even if the provisioning results in a capacity imbalance between appliances afterwards. If all vVols for a VM cannot fit on a single appliance due to space, system limits, or health issues, the remaining vVols are provisioned on to the appliance with the next-highest amount of free capacity.


When provisioning a new AppsON VM, the administrator can control the vVol storage placement. When deploying a VM to the vSphere cluster, the VM's vVols are placed on the appliance with the highest amount of free capacity. When deploying a VM to a specific host within the vSphere cluster, its vVols are stored on the appliance to which the node belongs.

When deploying a new AppsON VM using a template or cloning an existing VM, PowerStore places the new vSphere Virtual Volumes onto the same appliance as the source template or VM. This action enables the new VM to take advantage of data reduction to increase storage efficiency. For VM templates that are frequently deployed, it is recommended to create one template per appliance and evenly distribute VMs between appliances by selecting the appropriate template.

Regardless of how the VM is deployed, the compute node is always determined by VMware DRS when the VM is initially powered on. If DRS chooses a compute node that is not local to the vVol's storage appliance, compute and storage are not collocated. It is also possible for DRS to move virtual machines afterwards so that its compute and storage become separated later.

When taking a snapshot of an existing AppsON VM, new vVols are created to store the snapshot data. These new vVols are stored on the same appliance as the source vVols. In situations where the source vVols are spread across multiple appliances, the vVols created by the snapshot operation also become spread. vVol migrations can be used to consolidate a VM's vVols onto the same application.

To confirm compute and storage collocation for an AppsON virtual machine, browse to **Compute > Virtual Machines > Virtual Machine details > Virtual Volumes**. The **vSphere Host Name** column displays the vSphere name of the compute node for that vVol. The **Appliance** column displays the name of the storage appliance where that vVol is being stored. The following figure displays an optimal configuration:



**Figure 23. Virtual Volumes page for a virtual machine**

For an optimal configuration, store all vVols for a specific virtual machine on a single appliance. Also, the compute node for these vVols should be one of the two nodes of the appliance that is being used for storage. If there are any discrepancies, vSphere vMotion and PowerStore vVol migration can be used to move compute or storage to create an aligned configuration.

Starting with PowerStoreOS 2.0, PowerStore automatically creates a host group, VM group, and VM/host affinity rule that ties them together in VMware vSphere. The host

group contains the two internal ESXi hosts, and one host group is created per appliance. The VM group is initially empty, and one VM group is created per appliance.

Administrators should manually add the relevant virtual machines into the VM group based on where their storage resides. The affinity rule states that the VMs in the group should run on the specified appliance. This rule ensures that VMs run on a compute node that has direct local access to its storage. These groups and rules are automatically added and removed as appliances get added and removed from the cluster.

To manage the affinity rules, browse to **Cluster > Configure > VM/Host Rules** in the vSphere Web Client. When a host group is selected, the two internal ESXi nodes for that appliance are displayed in the members list below. Any VMs that have storage residing on this appliance can be added in the VM group, as shown in the following figure. If VM storage is migrated to another appliance within the cluster, update these rules to reflect the new configuration.




Figure 24. Host/VM Rules

## VMware datastores

### Overview

PowerStore displays tight integration with VMware, supporting vVol, VMFS, and NFS datastores backed by storage containers, volumes, and the file system, respectively. PowerStore natively supports visibility into vVol datastores, pulling all virtual machines hosted on PowerStore vVol datastores into PowerStore Manager for direct monitoring. With the introduction of PowerStoreOS 3.0, this VMware visibility is expanded to include NFS and VMFS datastores backed by PowerStore storage.

### vVol datastores

vVol datastores are fully supported on PowerStore and are backed by storage container objects. See the section [vSphere Virtual Volumes](#) for a detailed explanation of vVols and their support on PowerStore.

### NFS datastores

NFS datastores use the PowerStore File System, a 64-bit file system architecture, which includes several advantages and a maximum size of 256 TB. Other features include file

system shrink, extend, replication, snapshots, and more. For more information about PowerStore file systems, see the document [Dell PowerStore: File Capabilities](#).

Before using NFS datastores, create an NFS-enabled NAS server. You must create a file system associated with this NAS server and an NFS export. VMware ESXi hosts require read/write and root access to the NFS export. In vSphere, administrators must create an NFS datastore that uses the PowerStore file system.

With the introduction of PowerStoreOS 3.0, a new **VMware** type file system is supported on PowerStore. This file system is designed for VMware NFS datastore use cases and contains several enhancements for VMware environments. For more information about VMware Filesystems on PowerStore, see the document *Dell PowerStore: File Capabilities*.

**VMFS datastores** VMFS datastores are accessed through block protocols, and SCSI (Fibre Channel or iSCSI) or NVMe over Fabrics (NVMe/TCP or NVMe/FC) connectivity is required. After the communication path is established, ensure the VMware ESXi hosts for these datastores are registered by creating host objects on PowerStore. Then, you can create block volumes and map them to the VMware ESXi hosts. In vSphere, administrators must create a VMFS datastore that uses the PowerStore volume that is mapped to the VMware ESXi host.

With the introduction of PowerStoreOS 3.0, PowerStore Manager provides visibility into VMFS datastores created on PowerStore volumes. If the vCenter is registered with PowerStore, you can use PowerStore Manager to view the virtual machines on the datastore, along with their compute and storage metrics. The **Volumes** page has a new **Datastore** column (hidden by default) which shows the mapping from the volume to the VMFS datastore.

## PowerStore X models


### Licensing

Each PowerStore X model node has VMware ESXi installed on it. Each node requires a VMware vSphere Enterprise Plus license, which can be applied after the appliance is installed. You can provide your own license or purchase one along with the PowerStore X model appliance.

Starting with PowerStoreOS 1.0.3, vSphere Remote Office Branch Office (ROBO) licenses can be installed on PowerStore X model nodes. PowerStore X models support both the vSphere ROBO Advanced and ROBO Enterprise licenses. ROBO edition licenses are limited to 25 virtual machines, including the PowerStore X controller VMs. During initial configuration of the PowerStore X model appliance, the appliance automatically enables Distributed Resource Scheduler (DRS) in partially automated mode. vSphere ROBO Advanced licenses do not support DRS, and vSphere ROBO Enterprise licenses only support DRS for entering maintenance mode. Before installing a ROBO license on a PowerStore X model node, you must disable DRS on the ESXi cluster. When using a vSphere ROBO license, the user must manually initiate VM load balancing.

For more information about the ROBO license support, see the VMware document [VMware vSphere Compute Virtualization: Licensing, pricing and packaging](#).

Starting with PowerStoreOS 3.2.0, PowerStore X systems issue alerts about expiring internal ESXi licensing. The system provides a warning-level alert that the underlying internal ESXi host’s license will expire in x number of days. Once a permanent license is applied to the internal ESXi host, the warning alert is cleared automatically, and the system issues an informational alert of ESXi internal host is permanently licensed. The alerts are shown in PowerStore Manager under **Monitoring > Alerts**, as shown in Figure 25.




**Figure 25. ESXi licensing alert**

## Performance best practices

When configuring a new PowerStore X model appliance, it is highly recommended to apply these best practices to enable maximum performance. These settings should be changed before provisioning any resources on the appliance to avoid interruptions.

Starting with PowerStoreOS 1.0.3, these best practices can be applied during the Initial Configuration Wizard (ICW). In the ICW, an optional optimization step is displayed after the cluster is configured. At this step, the administrator can customize the MTU size and provide additional IPs to be used as iSCSI targets. Starting with PowerStoreOS 2.0, clustering is supported on PowerStore X model appliances, and the number of additional IP addresses requested by the system depends on the appliance count and model. The system does not request any additional IPs for any PowerStore 1000X model systems in the cluster because those models do not require additional iSCSI targets. Then, the system automatically configures the cluster with the best practices described in this section, and no further action is needed. The following figure shows the **Optimization** page of the ICW for a PowerStore X model cluster:



**Figure 26. ICW Optimization step**

We recommend completing best practices optimizations before adding a PowerStore X model appliance to an existing cluster. Starting with PowerStoreOS 2.0, the add appliance wizard has an **Optimize Performance** checkbox. If this box is selected, the add appliance wizard requests additional IP addresses for the new appliance. The system then automatically optimizes the newly added appliance to be consistent with the other appliances in the cluster. If the cluster is not optimized and there are no plans to change this state, you can add the new appliance without selecting **Optimize Performance**. Mixing optimized and non-optimized appliances within a cluster is not supported.

If you are planning to apply these best practices to an already configured system, it is recommended to first upgrade the system to PowerStoreOS 1.0.3 or later. Starting with PowerStoreOS 1.0.3, because some of the steps in the procedure are automated, fewer manual steps are required to achieve the same settings and results.

For more information about how to apply the PowerStore X performance best practice tuning, reference article HOW17288 on Dell Support.

When implementing the best practices in this document, we also recommend reviewing and applying the VMware vSphere settings that are described in the *PowerStore Host Configuration Guide* at [Dell.com/powerstoredocs](https://www.dell.com/powerstoredocs) and the *Dell PowerStore: Best Practices Guide* on the [PowerStore Info Hub](https://www.dell.com/powerstoreinfohub). The Dell Technologies *Virtual Storage Integrator* can also be used to automatically apply these best practices to your host.

## Initial configuration

The ICW prompts for the vCenter server details on PowerStore X model appliances. You must provide the details for an existing vCenter server that is hosted on an external server. This page is not displayed when configuring a PowerStore T model appliance.

The vCenter information enables automation during the initial configuration process. These steps can include establishing the vCenter connection, creating the vSphere cluster, configuring objects such as virtual distributed switches, and registering the VASA storage provider. If an existing datacenter name is specified, the cluster is created

underneath that datacenter. Otherwise, a new datacenter with the specified name is automatically created for this cluster.

On PowerStore X models, you cannot change the vCenter connection to another vCenter instance. This limitation is due to the existence of the objects such as the datacenter, cluster, PowerStore X model ESXi nodes, virtual distributed switches, and other configurations. The following figure shows the PowerStore X model ICW **Hypervisor** page:




Figure 27. PowerStore X model Initial Configuration > Hypervisor page


## AppsON

Integration of the PowerStore container-based architecture with the onboard VMware ESXi results in a new level of consolidation for enterprise storage. This ability combines the benefits of a local on-array application environment with unmatched integration with the vSphere management environment and server resources. This integration allows users to bring applications closer to storage, by running applications as virtual machines directly on PowerStore.

Benefits of the AppsON capability include a new level of agility for application deployments. This feature enables seamless movement between the PowerStore appliances and VMware ESXi servers. It also shrinks the stack by eliminating server and networking footprint for space-efficient edge and remote deployments. AppsON is ideal for data-intensive applications that require low latency or heavy storage use compared to compute.

## vCenter

Due to the embedded VMware ESXi hypervisor on PowerStore X model appliances, these nodes can be managed and monitored in vCenter, along with other ESXi hosts. For PowerStore X models, the vCenter must be hosted on an external server. Standard vSphere concepts such as datacenter, cluster, hosts, and virtual distributed switches are applied to the PowerStore X model objects. The following figure shows these objects along with the controller VMs in vSphere:



**Figure 28. PowerStore X model objects in vSphere**

## ESXi cluster

During initialization of the PowerStore X model appliance, the system creates a ESXi cluster that contains the PowerStore X model ESXi nodes. This ESXi cluster is created under a new or existing datacenter in the vCenter.

### Adding an external ESXi host into a PowerStore X model ESXi cluster

External ESXi hosts can also be added to the PowerStore X model ESXi cluster with an approved Request for Product Qualification (RPQ). The RPQ is needed due to considerations for networking and CPU model compatibility with vSphere Enhanced vMotion Compatibility (EVC).

Adding an external ESXi host into the cluster enables the ability to use external compute for VM load balancing and high availability. It is not necessary to have the external ESXi host in the same cluster to enable vMotion and Storage vMotion between PowerStore X model internal and external ESXi hosts.

### Non-Disruptive Upgrade (NDU) and external ESXi hosts

When you upgrade to PowerStoreOS 2.1.1 from previous versions of PowerStore, the PowerStore X internal ESXi cluster's vSphere Distributed Virtual Switch (DVS) is automatically upgraded to DVS 7. This upgrade of the DVS fails if there are any external ESXi hosts in the PowerStore X internal ESXi cluster not running vSphere 7 or later.

If the DVS upgrade fails, an alert notifies the user to upgrade the external ESXi hosts to vSphere version 7 or later and manually complete the DVS upgrade. The PowerStoreOS NDU can complete successfully even if the DVS upgrade fails, and a failed DVS upgrade is nondisruptive to the cluster. For more information about NDU, see the *Dell PowerStore Software Upgrade Guide*.

## Controller VMs

Each PowerStore X model appliance includes two controller VMs, one for each node. These VMs run a virtualized version of the PowerStore operating system. Each controller VM reserves 50% of the available CPU and memory on the appliance, leaving the other 50% for user VMs. Resources are guaranteed for the controller VMs, so that there is no

resource contention between user VMs and controller VMs. It is normal for high CPU and memory alerts to be generated for these controller VMs in vCenter due to the nature of guaranteeing resources for the Controller VMs.

Each controller VM resides on a private datastore, which is provisioned on the internal M.2 device on each physical node. These private datastores are reserved for the controller VMs and should *never* be used for user VMs. The controller VM must always reside on its associated node and can *never* be migrated. Since these VMs are fully dedicated and vital to the PowerStore X model storage operations, it is crucial that you *never* make changes to the controller VMs. *Do not* replicate or take snapshots of the controller VMs.

The controller VMs are named **PSTX-<DST>-<A/B>**, where **DST** is the Dell Service Tag for the appliance. They are stored on a private local VMFS6 datastore that is named **PRIVATE-<DST>-<A/B>.INTERNAL**. These private datastores are reserved only for the controller VMs and should not be used to store any user VMs. All user VMs should be stored on the vVol datastore instead.

## Networking

PowerStore X model appliances have a vSphere Distributed Virtual Switch (DVS), multiple port groups, and NIC teaming that is configured automatically as part of the initial configuration process. The DVS has a naming convention of **DVS-<Cluster\_Name>**. Then, the DVS name is prepended to each port group name, along with a dash.

The DVS has the following port groups created, by default:

- PG\_MGMT: PowerStore management
- PG\_MGMT\_ESXi: ESXi management
- PG\_Storage\_INIT1 - 2: VMkernel adapters for iSCSI connectivity from ESXi to the controller VM
- PG\_Storage\_TGT1 - 4: iSCSI targets on the controller VM for internal and external connectivity
- PG\_vMotion1: vMotion network used for VM mobility

The vSphere DVS groups the physical adapters from both nodes together into uplinks. The uplinks are used on each of the port groups to indicate which ports are active, standby, or unused. The following table shows the mapping between the vSphere uplink, vSphere physical adapter, and PowerStore Manager port names:

**Table 2. Uplink to physical port mappings**

vSphere Uplink	vSphere Physical Adapter	PowerStore Manager Port
Uplink1	vmnic8	4PortCard-hFEPort1
Uplink2	vmnic9	4PortCard-hFEPort0
Uplink3	vmnic6	4PortCard-hFEPort3
Uplink4	vmnic7	4PortCard-hFEPort2

The following table shows the management port groups that are used for the controller VM and ESXi management. Both management networks are configured with Uplink1 and

Uplink2 as active for high availability. Uplink3 and Uplink4 are configured as standby in case the primary uplinks become unavailable.

The PowerStore X model ESXi node management interface is configured on a VMkernel adapter named **vmk0**. Because the PowerStore management interface resides on the controller VM, it does not require a VMkernel adapter.

**Table 3. Management port group uplinks**

vSphere VMkernel Adapter	vSphere Port Group	vSphere Active Uplinks	vSphere Standby Uplinks
N/A	PG_MGMT	Uplink2 Uplink1	Uplink3 Uplink4
vmk0	PG_MGMT_ESXi	Uplink2 Uplink1	Uplink3 Uplink4

The following table shows the VMkernel adapters that are created for storage connectivity. The PowerStore X model ESXi nodes use these VMkernel adapters to connect to the iSCSI targets on the controller VMs. There are two VMkernel adapters on each node for multipathing purposes. The VMkernel adapters are active on one uplink and there are no standby uplinks.

The communication between the node and controller VM is used to establish iSCSI sessions, create protocol endpoints, and run I/O to the vVol datastore. Because the controller VM runs on the node itself, the traffic on these networks remains local to the node.

**Table 4. VMkernel adapters**

vSphere VMkernel Adapter	vSphere Port Group	vSphere Active Uplink
vmk1	PG_Storage_INIT1	Uplink1
vmk2	PG_Storage_INIT2	Uplink2

The following table shows the controller VM iSCSI targets that are created. These targets enable both the PowerStore X model ESXi node and external hosts to establish iSCSI connectivity. A minimum of one per node is required, which is configured automatically as part of the initial configuration process. By default, this target is active on Uplink1 on each node. The remaining uplinks are configured in standby mode.

**Table 5. Controller VM iSCSI targets**

vSphere Port Group	vSphere Active Uplink	vSphere Standby Uplinks
PG_Storage_TGT1	Uplink1	Uplink2 Uplink3 Uplink4
PG_Storage_TGT2	Uplink2	Uplink1 Uplink3 Uplink4
PG_Storage_TGT3	Uplink3	Uplink1 Uplink2 Uplink4

vSphere Port Group	vSphere Active Uplink	vSphere Standby Uplinks
PG_Storage_TGT4	Uplink4	Uplink1 Uplink2 Uplink3

The storage network can be scaled out to enable connectivity on the remaining ports on the four-port card. When this action is done, the additional uplinks become active, as shown in the table above. Depending on the model appliance, this action may be a best practice for maximum performance.

The following table shows information about the virtual ports that are available on the appliance. The virtual ports page can be used to map additional ports for the storage network or tag additional ports for the replication network. By default, vFEPort1 is tagged for both storage and replication.

**Table 6. PowerStore virtual ports**


PowerStore Manager Virtual Port	vSphere Network Adapter	vSphere Port Group	Purpose
vFEPort0	Network adapter 1	PG_MGMT	PowerStore management
vFEPort1	Network adapter 2	PG_Storage_TGT1	Storage and replication network
vFEPort2	Network adapter 3	PG_Storage_TGT2	Storage and replication network scaling
vFEPort3	Network adapter 4	PG_Storage_TGT3	Storage and replication network scaling
vFEPort6	Network adapter 5	PG_Storage_TGT4	Storage and replication network scaling
vFEPort7	Network adapter 6	PG_Internal	Internal system use

The following table shows the VMkernel adapters that are created for vMotion operations. This network is used when moving VMs between the two PowerStore X model ESXi nodes and from external hosts.

**Table 7. vMotion port group uplink**

vSphere VMkernel Adapter	vSphere Port Group	vSphere Active Uplinks	vSphere Standby Uplinks
vmk3	PG_vMotion1	Uplink3	Uplink1 Uplink2 Uplink4

The following figure shows these networks as they appear in vCenter:



**Figure 29. vSphere networks**

Before deploying a user VM on the internal ESXi nodes, create a new port group for the external network. This process is completed by right-clicking the **DVS > Distributed Port Group > New Distributed Port Group**. Provide the information for the new port group and configure a VLAN, if necessary. After the new port group is configured, user VMs can be deployed and can use this port group for network connectivity.

If the PowerStore X model ESXi node, iSCSI, or vMotion interfaces must change, they must be updated in PowerStore Manager. This action updates the configuration and propagates the necessary changes to vSphere automatically. Changing these interfaces directly in vSphere is **not supported**.

## Volumes

PowerStore X model appliances can provision volumes and volume groups to external hosts. For example, you can provision volumes to external ESXi nodes for Virtual Machine File System (VMFS) datastores or Raw Disk Mappings (RDMs).

By default, AppsON virtual machines use the PowerStore efficient vVol implementation due to its simple nature, design optimizations, and integration within PowerStore Manager. Due to these benefits, it is recommended to use vVols for all AppsON virtual machines. Starting with PowerStoreOS 2.0, PowerStore X model appliances also support VMFS datastores for the storage of virtual machines within AppsON. This process is done by allowing the mapping of block volumes to the PowerStore internal ESXi hosts using the PowerStore REST API or CLI, or both. For more information about configuring VMFS on PowerStore X model appliance internal nodes, reference article KB182913 on Dell Support.

## Distributed Resource Scheduler Monitoring Service

PowerStore X model ESXi nodes are designed to work with VMware Distributed Resource Scheduler (DRS). During initial configuration of the PowerStore X model appliance, the appliance automatically enables DRS in partially automated mode. Partially automated mode automatically applies DRS for initial VM placement and makes suggestions for load balancing, which the administrator can initiate.

Because the appliance is optimized for and expects this configuration, changing the DRS automation level is **not supported**. The DRS Monitoring Service polls vSphere every 15

seconds and confirms that the DRS automation level is set to partially automated. If a change is detected, it automatically heals by reverting the automation level back to partially automated.

Starting with PowerStoreOS 1.0.3, vSphere ROBO licenses can be installed on PowerStore X model nodes. vSphere ROBO Advanced licenses do not support DRS; vSphere ROBO Enterprise licenses only support DRS for entering maintenance mode. Before you install a ROBO license on a PowerStore X model node, you must disable DRS on the ESXi cluster. When using a vSphere ROBO license, the user must manually initiate VM load balancing.

### Serviceability

On a PowerStore T model appliance, a node can be rebooted or powered off in PowerStore Manager. On a PowerStore X model appliance, these operations are not available in PowerStore Manager. Instead, after you put the PowerStore X model into maintenance mode, initiate a reboot or power off the system. This helps prevent accidental reboots of PowerStore X model ESXi nodes that have VMs running on them. The following figure shows that these operations are unavailable on a PowerStore X model appliance:





Figure 30. PowerStore X model operations

### Maintenance Mode Service

Because DRS cannot migrate controller VMs, PowerStore X model appliances include the Maintenance Mode Service (MMS), which manages the controller VMs during maintenance mode operations. Instead of being moved, the controller VM is gracefully powered off.

Before shutting down or rebooting a PowerStore X model ESXi node, put the node into maintenance mode first. Entering maintenance mode ensures that there are no VMs running on this node before it is shut down or rebooted. When maintenance mode is entered, DRS migrates any running VMs to the peer node in the vSphere cluster. The following figure shows the maintenance mode operations that are available in vCenter:



**Figure 31. Maintenance Mode**

If maintenance mode is initiated on a PowerStore X model ESXi node, MMS automatically initiates a shutdown of the node controller VM. The shutdown of the controller VM is started after the migration of all the user VMs has completed. After the controller VM successfully powers off, the ESXi node enters maintenance mode. When maintenance mode is entered, the ESXi node can be rebooted or shut down from vCenter without impact.

When one of the controller VMs is powered off or rebooted, the services fail over to the other controller VM. To avoid interruption, place only a single node of an appliance in maintenance mode at a time. To restore high availability, maintenance mode must be exited on the node. When the administrator initiates an exit maintenance mode operation, MMS automatically powers up the controller VM. When the controller VM fully powers up, controller VM redundancy is restored.

After entering or exiting maintenance mode on a PowerStore X model ESXi node, wait a few minutes before issuing another maintenance mode operation. This period provides the controller VMs enough time to completely fail over all resources and services before starting the next operation.

## Upgrades

You can use the PowerStore X model upgrade image to upgrade the PowerStore X model cluster to the latest software version. However, the PowerStore X model ESXi nodes can only use ESXi versions that are validated by Dell Technologies and available on [Dell Support](#). See Table 12 in the PowerStore Simple Support Matrix for the supported ESXi versions for each PowerStore release. **Do not use ESXi update images that are obtained from VMware or any other source.** When a new version is available for update, a notification is posted. For more information, see the upgrade procedure in the *Dell PowerStore Virtualization Guide* on the [PowerStore Info Hub](#).

## VMware VAAIOverview

vSphere API for Array Integration (VAAI) improves ESXi host utilization by offloading storage-related tasks to PowerStore. Because the array processes these tasks, the ESXi host CPU, memory, and network utilization are reduced. For example, an operation such as provisioning full clones from a template VM can be offloaded to PowerStore. PowerStore processes these requests internally, performs the write operations, and returns an update to the ESXi host when the requests are complete.

The following primitives are supported with PowerStore:


- **Block:**
  - **Atomic Test and Set (ATS):** Enables arrays to perform locking at the block level of a volume, instead of the whole volume, which enables multiple ESXi hosts to access a volume simultaneously. This is also known as Hardware-Assisted Locking.
  - **Block Zero:** Enables arrays to zero out many blocks, which speeds VM provisioning by accelerating the disk zeroing operation. This is also known as Hardware-Assisted Zeroing or Write Same.
  - **Full Copy:** Enables arrays to make full copies of data within the array without the need for the ESXi host to read and write the data. This is useful when cloning VMs, and it is also known as Hardware-Assisted Move or XCOPY. (XCOPY is not standardized in NVMe specifications, and NVMe/TCP and NVMe/FC do not support Full Copy offload.)
  - **Thin Provisioning – Unmap:** Enables arrays to reclaim unused blocks on a thin LUN. Unmap is also known as Dead Space Reclamation.
- **File:** These primitives are introduced in PowerStoreOS 3.0 and require installation of the VAAI plug-in on the ESXi hosts.
  - **Fast File Clone:** Enables the creation of virtual machine snapshots to be offloaded to the array.
  - **Full File Clone:** Enables the offloading of virtual-disk cloning to the array.
  - **Reserve Space:** Enables provisioning virtual disks using the Thick Lazy and Eager Zeroed options over NFS.
  - **Extended Statistics:** Enables visibility into space usage on NAS datastores and is especially useful for thin-provisioned datastores.

## Migration

### Overview

PowerStore is designed to integrate easily and seamlessly into an existing VMware vSphere environment. Native vSphere features and tools can be used between PowerStore and external ESXi hosts.

This ability enables performing quick and simple migrations by using tools such as vMotion and Storage vMotion. vMotion can be used to move VM compute off the current ESXi host and onto a PowerStore X model node. Storage vMotion can be used to move VM storage off the current datastore and onto the PowerStore vVol datastore. You also have the option of performing both a vMotion and Storage vMotion simultaneously, as shown in the following figure:



**Figure 32. Compute and Storage vMotion**

## Metro Volume

### Overview

Metro Volume is a high availability and data mobility feature for PowerStore storage and VMware vSphere. It provides symmetric active/active data access to Metro Volumes for proactive use cases between PowerStore clusters. The architecture also lays a foundation for VMware vSphere Metro storage cluster designs. For a detailed review of Metro Volume, see the document [Dell PowerStore: Metro Volume](#).

## vVol replication

### Overview

PowerStoreOS versions 3.0 and later support VASA 3.0 native storage-based asynchronous replication for vVol-based VMs. This feature uses VMware Storage Policies and requires VMware Site Recovery Manager instances in both sites. Asynchronous replication of vVol-based VMs is included at no extra cost for supported PowerStore clusters. See the document [Dell PowerStore: VMware Site Recovery Manager Best Practices](#) or the VMware Site Recovery Manager product documentation for more information.


## VMware plug-ins

### Introduction

To further enhance the VMware integration that is built into the system, plug-ins for off-array software are available. These plug-ins provide flexibility and enable PowerStore to easily integrate into your environment using existing tools.

### Virtual Storage Integrator

Virtual Storage Integrator (VSI) brings storage provisioning, management, and monitoring capabilities to the standard VMware vSphere Client interface. Viewing and performing common storage tasks can be accomplished directly from vSphere, without needing to launch PowerStore Manager. The VSI plug-in also provides visibility into the storage system, enabling administrators to see the underlying storage on which their VMs are running. When connecting external ESXi hosts to PowerStore, use VSI to scan the host and apply best practices for performance and availability. The following figure shows the datastore creation wizard in VSI:



**Figure 33. Creating a datastore using VSI**

## vRealize Orchestrator

VMware vRealize Orchestrator (vRO) enables creating automation workflows to streamline VMware and PowerStore tasks. The PowerStore plug-in includes many workflows such as provisioning storage, managing hosts, configuring protection, and viewing the details of the resources.

The vRO framework allows individual workflows to be put together to build a custom workflow. For example, you can create a custom vRO workflow that connects an ESXi host to the iSCSI target on the PowerStore appliance and then registers the host on the appliance. The vRO workflow engine can be used with vRealize Automation to create a policy-based self-service environment.

The following figure shows some of the workflows available in vRO with the PowerStore plug-in:




Figure 34. vRealize Orchestrator

## Storage Replication Adapter

A PowerStore Storage Replication Adapter (SRA) is available for customers that are using array-based replication and VMware Site Recovery Manager (SRM) for disaster recovery. For SRM to manage PowerStore replication properly, the SRA must be installed on the SRM Server hosts at the recovery and protected sites. The following figure shows the PowerStore SRA in SRM:




Figure 35. PowerStore SRA

## Best practices

Site Recovery Manager ships with a default configuration that is tuned for a large cross-section of environments. However, each environment is unique in terms of architecture, infrastructure, size, and recovery time objectives. Larger and more complex SRM environments may require tuning adjustments for SRM to work properly. See the document [Dell PowerStore: Site Recovery Manager Best Practices](#) for more information.

## RecoverPoint for Virtual Machines

PowerStore also supports VM-granular replication services by using RecoverPoint for Virtual Machines. RecoverPoint for VMs is a software-only replication solution that provides any-point-in-time asynchronous and synchronous protection on a per-VM basis. It is storage agnostic and works in the hypervisor layer with all storage types supported by

VMware, including vVols. For more information about RecoverPoint for VMs, see the document *RecoverPoint for Virtual Machines Administrator's Guide* on [Dell Support](#).

# Conclusion

## Overview

PowerStore was designed to include a comprehensive set of integration points with VMware virtualization technology. Because many of these powerful integration points are embedded in the system, they can be managed through the HTML5-based PowerStore Manager and vCenter. PowerStore X model appliances feature deeper integration by allowing applications to run directly on the appliance and seamlessly integrate into the virtualized environment. Off-array software and plug-ins are also available to enable PowerStore to be used with your existing tools. Both storage and virtualization administrators can use PowerStore to create a solution that meets requirements and supports today's business needs for maximum infrastructure flexibility.

## Appendix: Technical support and resources

### Resources

The [Dell Technologies Info Hub > Storage](#) site provides expertise that helps to ensure customer success with Dell Technologies storage platforms.

[Dell.com/powerstoredocs](https://www.dell.com/powerstoredocs) provides detailed documentation about how to install, configure, and manage PowerStore systems.