

# Dell PowerStore: Virtualization Integration

May 2024

H18152.12

## White Paper

### Abstract

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

## 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 © 2020-2024 Dell Inc. or its subsidiaries. All Rights Reserved. Published in the USA May 2024 H18152.12.

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>6</b>
<b>vCenter connection</b> .....	<b>8</b>
<b>ESXi hosts</b> .....	<b>12</b>
<b>vSphere Virtual Volumes</b> .....	<b>13</b>
<b>Virtual machines</b> .....	<b>22</b>
<b>VMware datastores</b> .....	<b>29</b>
<b>VAAI</b> .....	<b>30</b>
<b>Migration</b> .....	<b>31</b>
<b>Metro Volume</b> .....	<b>31</b>
<b>vVol replication</b> .....	<b>31</b>
<b>VMware plug-ins</b> .....	<b>31</b>
<b>Conclusion</b> .....	<b>34</b>
<b>Technical support and resources</b> .....	<b>35</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	Part number/ revision	Description
April 2020	H18152	Initial release: PowerStoreOS 1.0.0
August 2020	H18152.1	Minor updates
September 2020	H18152.2	Minor updates
December 2020	H18152.3	PowerStore 1.0.3 updates
April 2021	H18152.4	PowerStoreOS 2.0.0 updates
May 2021	H18152.5	Minor updates
January 2022	H18152.6	PowerStoreOS 2.1.0 updates Template update
April 2022	H18152.7	PowerStoreOS 2.1.1 updates
June 2022	H18152.8	PowerStoreOS 3.0.0 updates
October 2022	H18152.9	PowerStoreOS 3.2.0 updates
May 2023	H18152.10	PowerStoreOS 3.5.0 update <ul style="list-style-type: none"> <li>Support for VASA 4.0</li> </ul>
September 2023	H18152.11	PowerStoreOS 3.6.0 updates
May 2024	H18152.12	PowerStoreOS 4.0.0 updates <ul style="list-style-type: none"> <li>vCenter certificate verification</li> <li>Dynamic node affinity for vVols</li> <li>PowerStore Q model appliances</li> </ul> Dynamic AppsON updates Removed references to PowerStore X

### 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, Andrew Sirpis

---

**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 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 can also be deployed in a Dynamic AppsON configuration, which is an exclusive integration between compute-only Dell VxRail dynamic nodes and a Dell PowerStore storage system. This enables independent scaling of compute and storage, providing flexibility of choice by increasing the extensibility of both platforms. It provides VxRail environments access to PowerStore enterprise efficiency, data protection, and resiliency features. It also helps PowerStore environments quickly expand compute for CPU-intensive workloads in a traditional three-tier architecture. For more information about Dynamic AppsON, see the document [Dell VxRail and Dell PowerStore: Better Together Through Dynamic AppsON](#).

### Terminology

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

**Table 1. Terminology**

Term	Definition
Distributed Resource Scheduler (DRS)	A VMware feature that monitors resource utilization and spreads virtual machine workloads across ESXi hosts in a cluster.

Term	Definition
Dynamic AppsON	A configuration that provides an exclusive integration between compute-only Dell VxRail dynamic nodes and a Dell PowerStore storage system.
Fibre Channel (FC) protocol	A 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.
NVMe over Fibre Channel (NVMe/FC)	A protocol used to perform Non-Volatile Memory Express (NVMe) commands over a Fibre Channel network.
NVMe over TCP (NVMe/TCP)	A protocol used to perform Non-Volatile Memory Express (NVMe) commands over an Ethernet network.
PowerStore Manager	An HTML5 user interface used to manage PowerStore systems.
PowerStore Q model	A container-based storage system that is running on purpose-built hardware. This storage system supports unified (block and file) workloads, or block-optimized workloads. The PowerStore Q model supports Quad-Level Cell (QLC) NVMe SSDs for data storage.
PowerStore T model	A container-based storage system that is running on purpose-built hardware. This storage system supports unified (block and file) workloads, or block-optimized workloads. The PowerStore T model supports Triple-Level Cell (TLC) NVMe SSDs for data storage.
Storage container	A logical entity that consists of one or more supported storage capabilities along with its available capacity. 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.
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.

## 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.

Starting with PowerStoreOS 2.0, the initial configuration includes a step that lets you optionally configure the vCenter server connection. A vCenter server can also be connected after the initial configuration. To establish a vCenter server connection, open PowerStore Manager and go to **Compute > vCenter Server Connection**. You can connect a vCenter by entering the **vCenter Server IP Address** (or FQDN), **User Name**, and **Password** for an existing vCenter server. For the supported vCenter versions on PowerStore, see the *PowerStore: Simple Support Matrix* at [Dell.com/powerstoredocs](http://Dell.com/powerstoredocs).

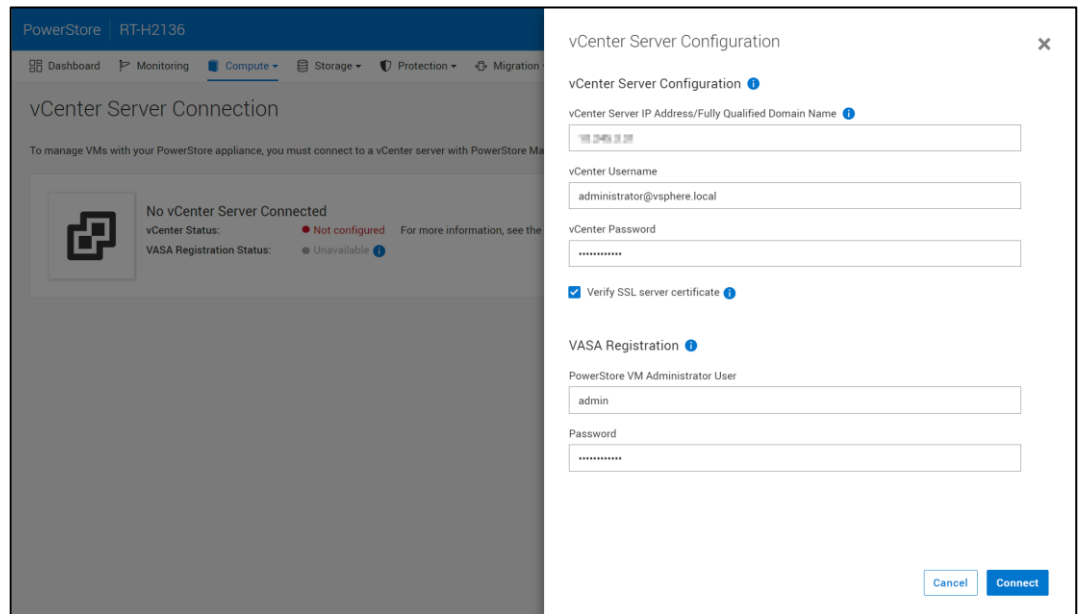
Starting with PowerStoreOS 2.0, you can manage and monitor VASA registration from PowerStore Manager. This ability removes the need to log in to vSphere to view or update the VASA registration. When registering a new vCenter, there is an option to provide the PowerStore credentials. If provided, these 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.

### Certificate verification

Starting with PowerStoreOS 4.0, administrators have the option to verify the vCenter SSL certificate. This enables the administrator the ability to confirm that the certificate matches the one that is displayed in vCenter before allowing communication. It improves security by ensuring that PowerStore is communicating with the intended vCenter. This option is available during initial configuration and on already configured systems through the GUI, PSTCLI, and REST API.

The following figure shows the vCenter Server registration page with the **Verify SSL server certificate** checkbox.

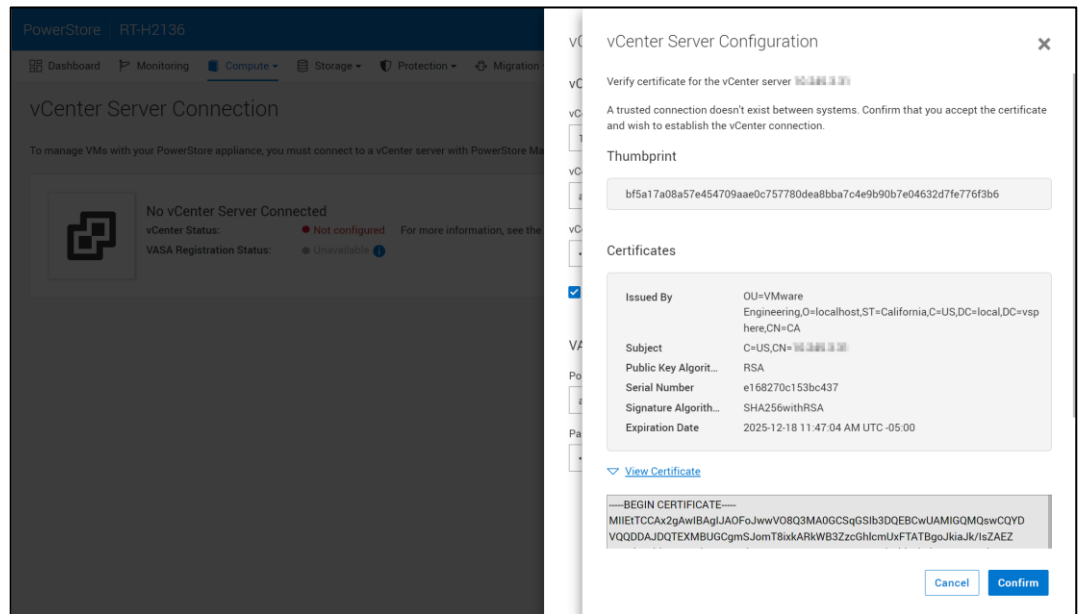




**Figure 1. Registering a vCenter server**

Certificate verification is enabled by default and is highly recommended. When it is enabled, the vCenter certificate is retrieved and displayed, as shown in [Figure 2](#). The administrator can review the certificate and confirm that it matches the certificate in vCenter. If the administrator confirms that the certificate matches, PowerStore stores a copy of the certificate, marks it as trusted, and allows communication. If the administrator notes that the certificate does not match, the system leaves the certificate untrusted, disallows communication, and the registration process is canceled.

The system also reviews the certificate's details to ensure that they are valid. It checks to ensure that the Fully Qualified Domain Name (FQDN) or IP address matches, that the certificate's start date has passed, that the expiration date has at least 30 days left, and so on. If any of these system checks fail, the connection is blocked due to an invalid certificate.



**Figure 2. vCenter certificate verification**

If certificate verification is disabled, the registration process proceeds without validating the certificate's details and does not display the certificate to the administrator.

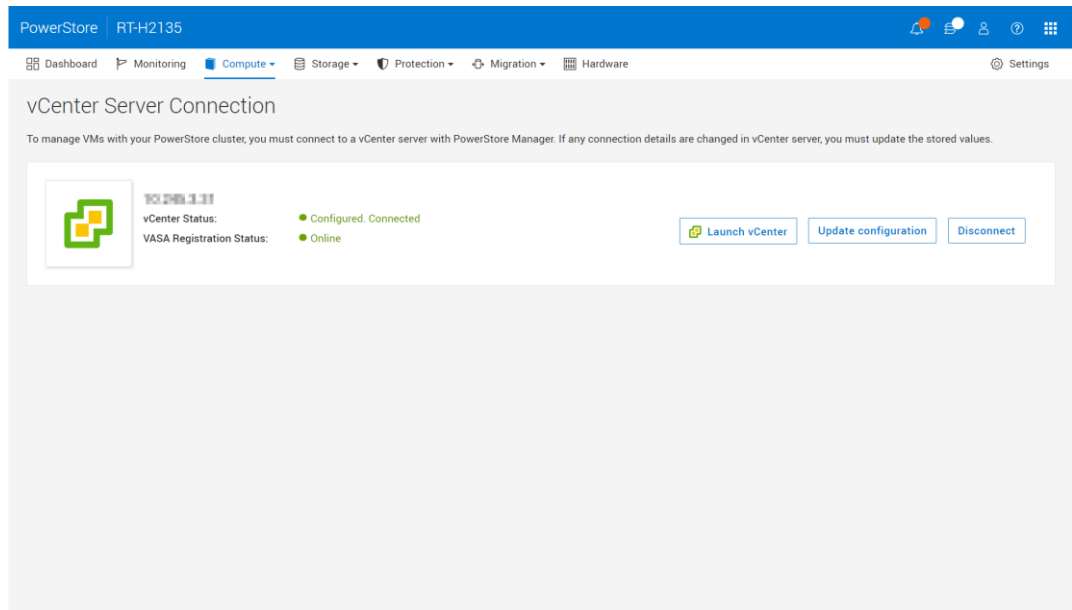
With certificate verification enabled, the vCenter certificate that is stored on PowerStore is used to compare to the one reported by the vCenter. However, normal operations such as certificate renewals can cause the vCenter certificate to change. If this happens, an alert is generated to inform the administrator that there is a certificate mismatch. If this is expected, the administrator can update the configuration to verify the new certificate.

If a PowerStore system with an already connected vCenter is upgraded to PowerStoreOS 4.0, certificate verification remains disabled. An alert is generated to inform the administrator that the connection is not secure. The administrator can update the configuration and enable certificate verification to clear the alert. If certificate verification is not desired (not recommended), you can enable/disable certificate verification or unregister/re-register the vCenter without certificate verification to clear the alert.

A vCenter registered with certificate verification enabled can be disconnected from PowerStore. When the vCenter is disconnected, the stored certificate is removed and is no longer trusted. Communication with this vCenter is no longer allowed. If desired, an administrator can connect PowerStore to the same vCenter again later using the same workflow as a new vCenter connection.

## vCenter connected

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** also become available, as shown in the following figure:



**Figure 3. 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. The vCenter connection can be disconnected and then connected to the new vCenter instance.

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:

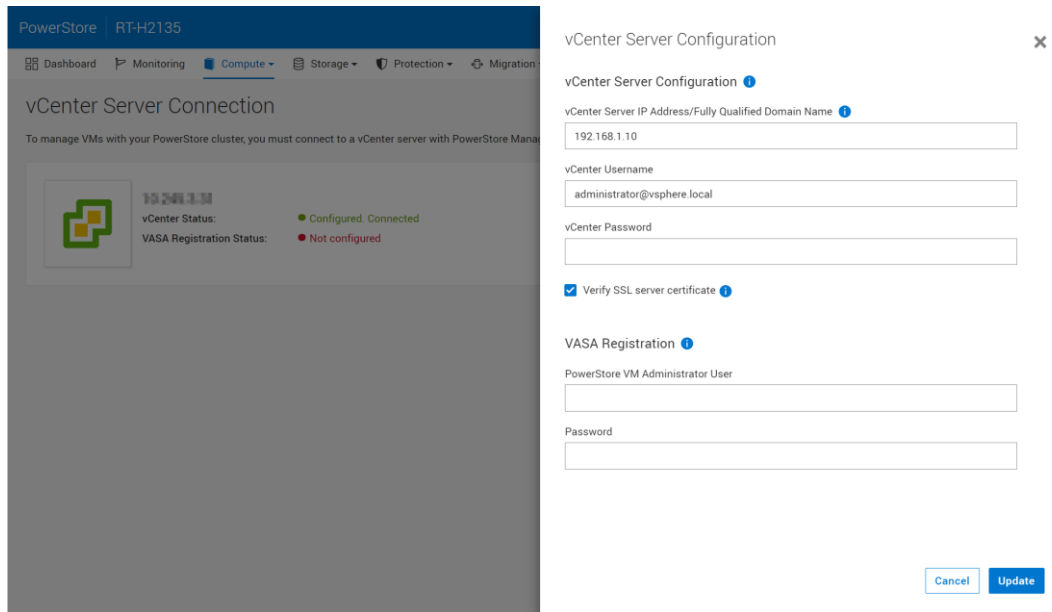


Figure 4. Updating the configuration

Use the **Disconnect** button to remove a vCenter connection. 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:

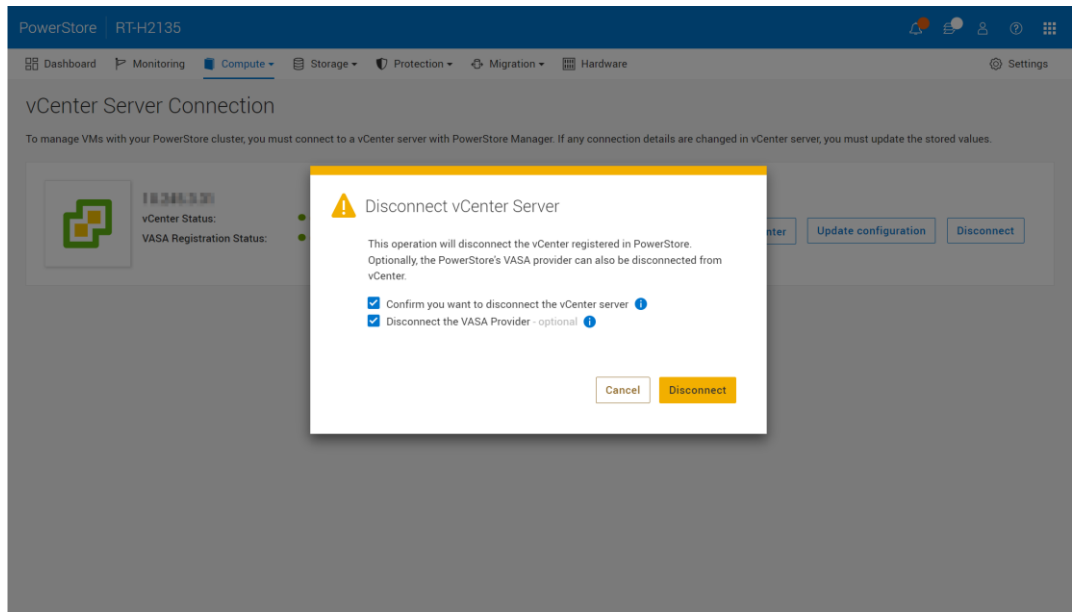


Figure 5. Disconnect vCenter Server confirmation dialog box

## ESXi hosts

### Overview

ESXi hosts can be registered in PowerStore Manager to enable access to storage resources. Starting with PowerStoreOS 2.0, enhancements are added to display additional details and improve visibility for ESXi hosts. These enhancements include:

- vSphere hostname:
  - The **vSphere Host Name** column displays the hostname that is displayed in the vSphere Web Client.
  - The **vSphere 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:
  - The **ESXi Version** column is also displayed in PowerStore Manager on the **Hosts & Host Groups** page.

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

Name	Alerts	vSphere Host Name	Host/Host Group	OS	Initiator Type	Initiators	Volume Mappings	ESXi Version
ESXi	-	10.245.1.100	Host	ESXi	iSCSI	1	-	VMware ESXi 7.0.3...

Figure 6. 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 and VASA 4.0 protocols. 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 and VASA 4.0 provider, which enables the vVols storage framework. The VASA provider must be registered in vSphere in order to use vVols.

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

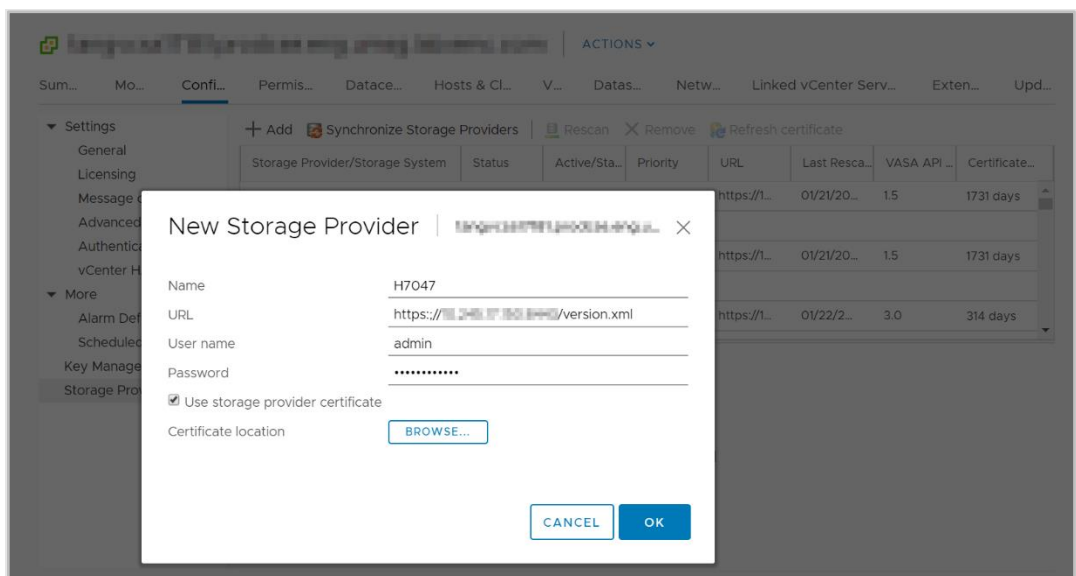
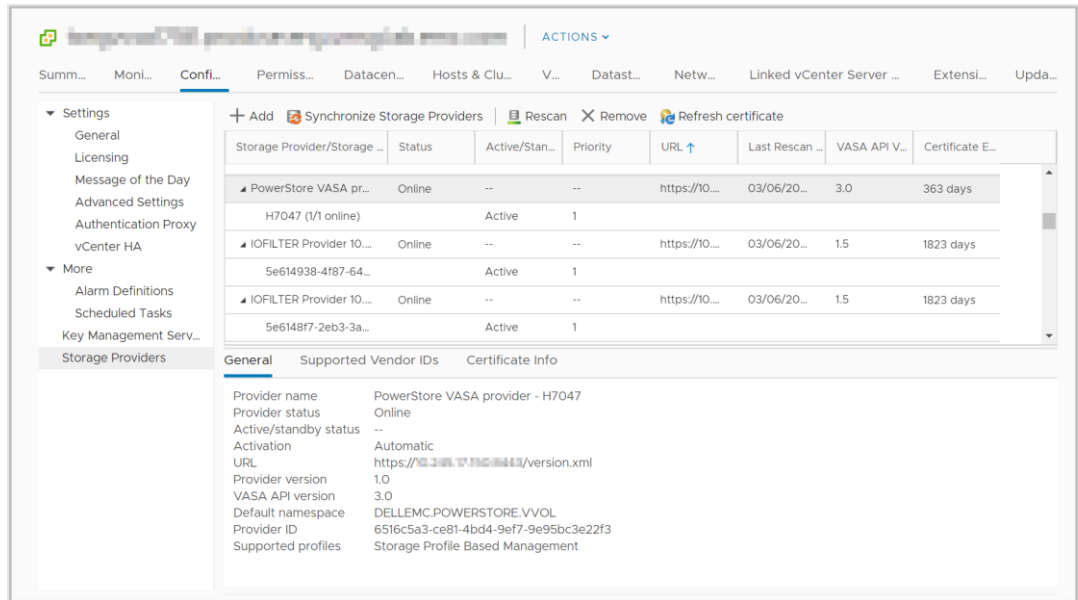


Figure 7. New Storage Provider page

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

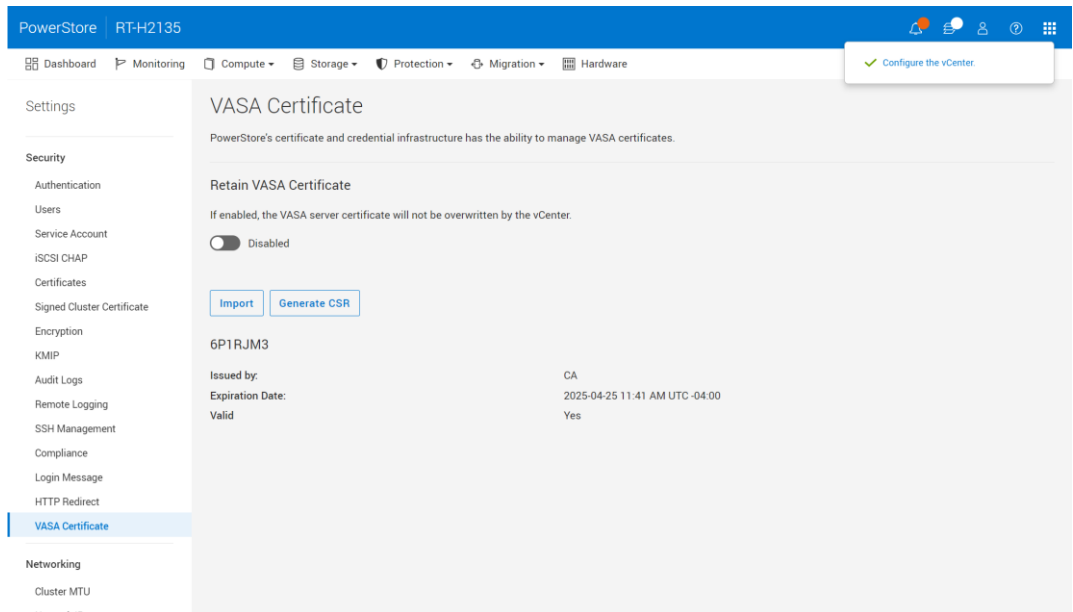


**Figure 8. Registered storage provider**

## VASA certificates

Starting with PowerStoreOS 3.5, you can implement a third-party certificate for the PowerStore VASA provider. You can replace the system's self-signed certificate with the third-party certificate, or you can retain the system's self-signed certificate and also add a third-party certificate. Implementing a third-party certificate for the PowerStore VASA provider allows for support of multiple standalone vCenter instances to leverage a single PowerStore VASA provider.

- To view or implement a third-party certificate for VASA directly from PowerStore Manager, go to **Settings** and select **VASA Certificate** under the **Security** section, as shown in the following figure.



**Figure 9. Import third-party VASA certificate**

Under **VASA Certificate**, you have the option to retain the current VASA certificate, whether it is a self-signed PowerStore certificate, a third-party signed certificate, or a vCenter certificate. If enabled, the retain option prevents the VASA server certificate on the PowerStore system from being overwritten by the vCenter server.

To implement a third-party certificate for VASA, follow these steps:

1. In PowerStore Manager, go to **Settings > VASA Certificate**.
2. Click **Generate CSR** and provide the following information to generate a Certificate Signing Request (CSR):
  - Common Name
  - IP Address
  - Organization
  - Organizational Unit
  - Locality
  - State
  - Country / Region
3. Have the CSR request shown in Figure 10 signed by a third-party Certificate Authority.
4. Click **Import** to import the signed certificate into PowerStore Manager.

---

**Note:** Ensure that the vCenter server trusts the CA of the certificate being imported; otherwise, VASA functionality, including the management of a vVol VM, will be unavailable.

---



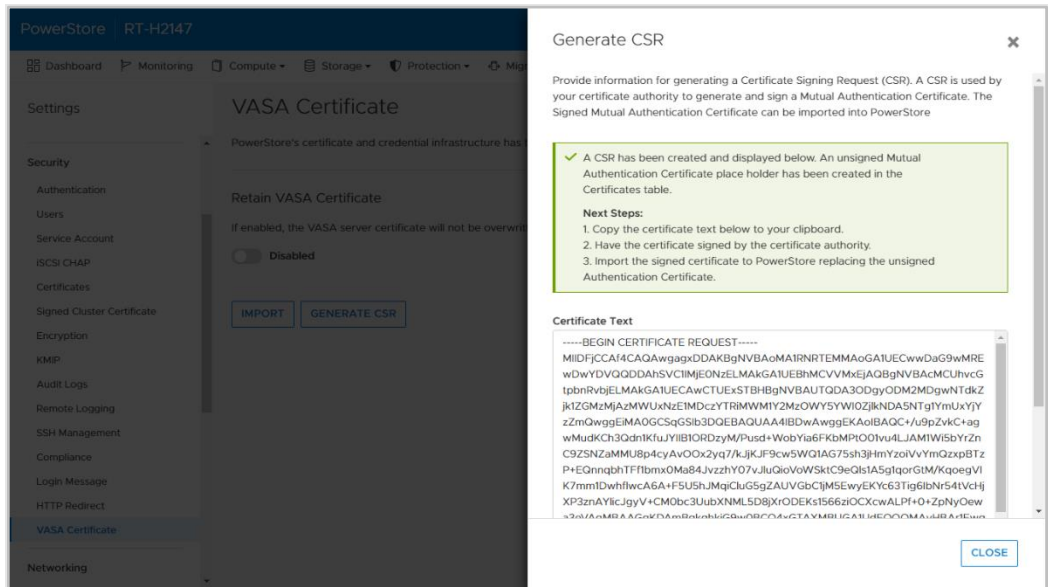


Figure 10. Generate CSR request for VASA

## 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. PowerStore includes a default storage container that is named PowerStore <Cluster\_Name>, as shown in the following figure:

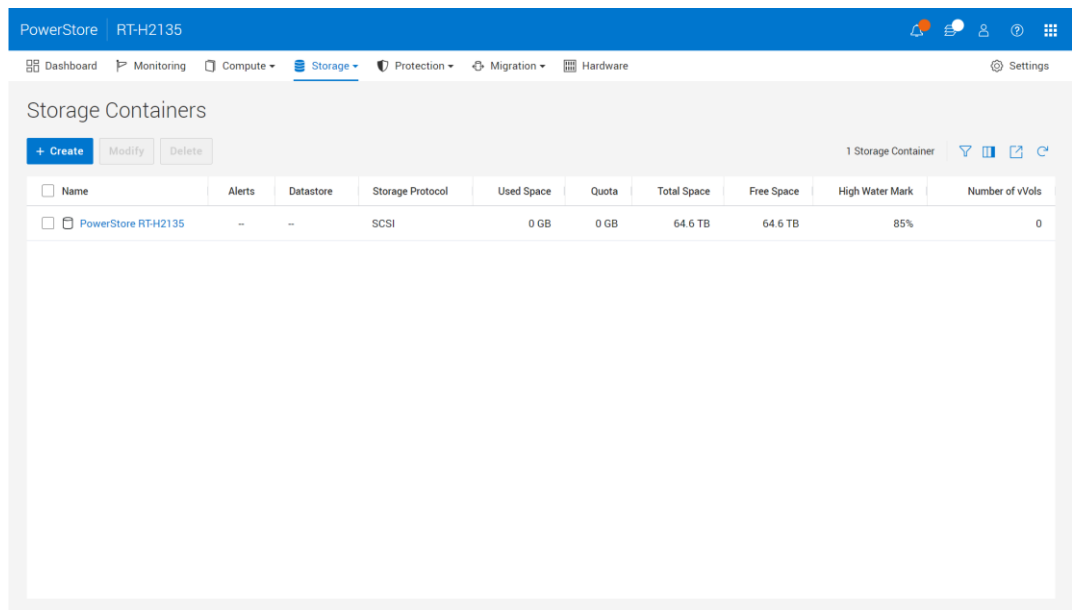


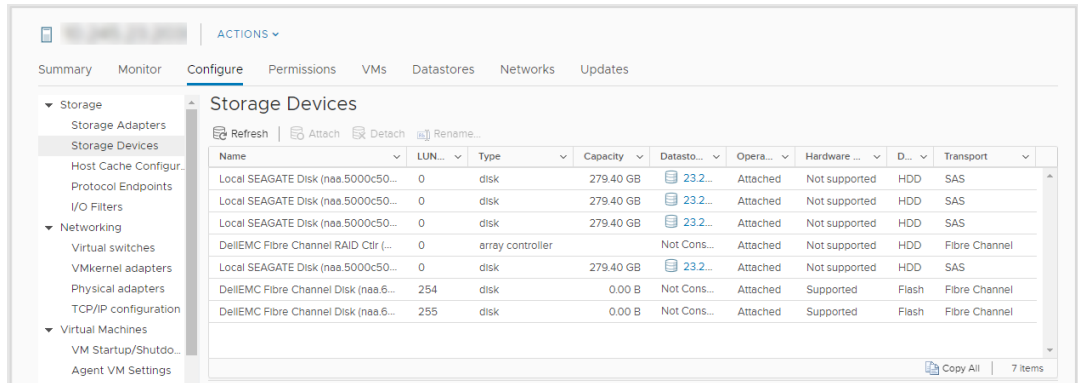
Figure 11. Default storage container

PowerStore can expose its storage containers to 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, NVMe/FC, or NVMe/TCP (PowerStoreOS 3.6 and later) connectivity between the 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 12. Protocol endpoints with LUN IDs 254 and 255**

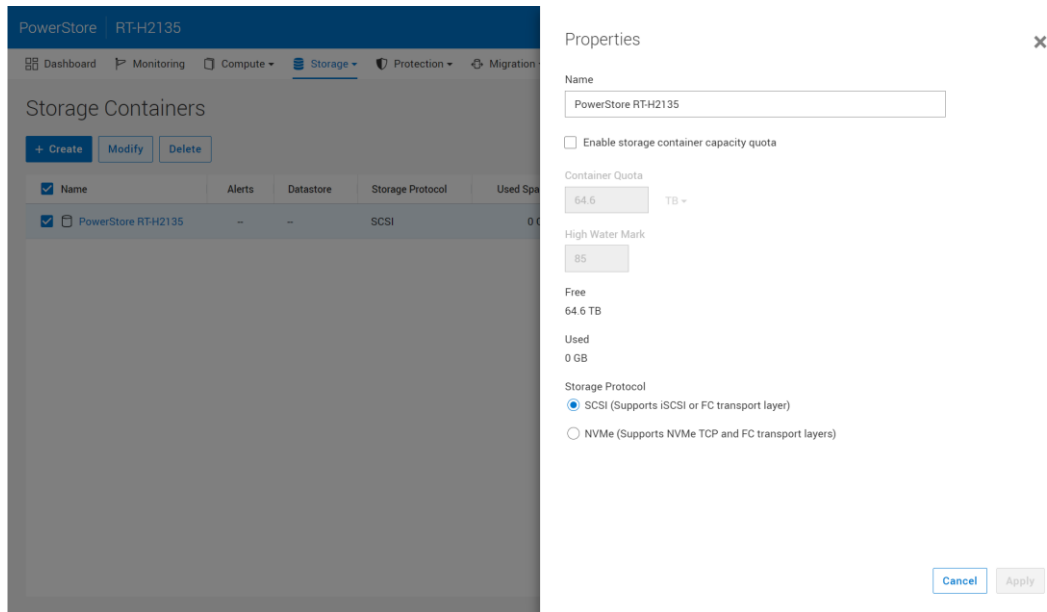
All registered ESXi hosts are automatically granted access to all the storage containers on PowerStore. These ESXi hosts can mount the datastore in vSphere after host connectivity is established and no further mapping is needed. The following figure shows the vVol datastore mounted in vSphere.



**Figure 13. PowerStore vVol datastore**

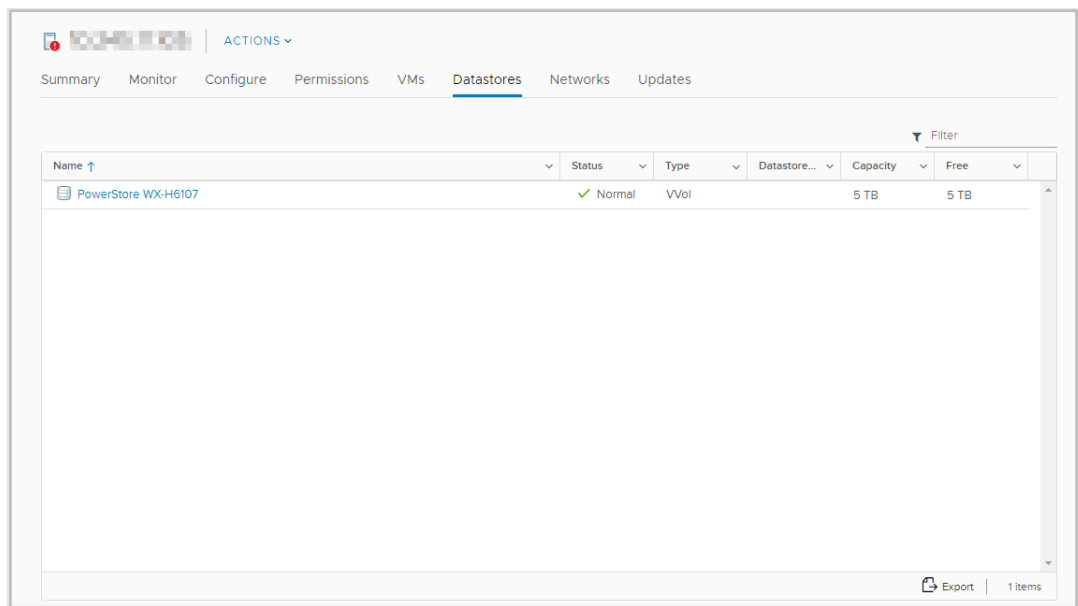
In addition to the default storage container, additional storage containers can also be created. These additional storage containers can be mounted as new vVol datastores in vSphere.

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 14. 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 15. 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. SCSI storage containers support host access through SCSI protocols, which include iSCSI or Fibre Channel. NVMe storage containers support host access through the NVMe TCP and FC protocols. PowerStoreOS 3.6 and later supports NVMe storage container host access through the NVMe/TCP protocol for vVols. NVMe/TCP vVols requirements include VMware 8.0 Update 1 and PowerStoreOS 3.6 or later. Users utilizing NVMe vVols with their ESXi host must choose either NVMe/FC or NVMe/TCP. Using both at the same time is not supported.

Note that ESXi hosts use unique and separate identification (host NQN and host ID pairs) to access vVol and non-vVol namespaces over NVMe. Due to this, it is required to create separate host entries of type "NVMe-vVol" to access vVols and "NVMe" to access non-vVol (VMFS) datastores. An ESXi host in PowerStore cannot use both NVMe vVol and VMFS datastores with a single host entry.

With support for NVMe storage containers, PowerStore systems running PowerStoreOS 3.0 and later leverage VASA version 4.0 for all NVMe storage containers. Any storage container using the SCSI protocol continues to use VASA version 3.0.

When creating a storage container on a system running PowerStoreOS 3.6 or later, you can select either **SCSI (Supports iSCSI or FC transport layer)** or **NVMe (Supports NVMe TCP or 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.

The screenshot shows a 'Create Storage Container' dialog box. At the top, there is a warning message: 'To access this storage container you must have VASA provider registered and then create vVol datastore on vCenter server.' Below the warning is a 'LAUNCH VCENTER' button. The 'Name' field contains 'NVMe-SC'. There is an unchecked checkbox for 'Enable storage container capacity quota'. The 'Container Quota' is set to '54.4' with a unit dropdown menu currently showing 'TB'. Under 'Select the storage protocol', the 'NVMe (Supports NVMe TCP and FC transport layers)' option is selected with a radio button, while the 'SCSI (Supports iSCSI or FC transport layer)' option is unselected.

**Figure 16. 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 17. 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:

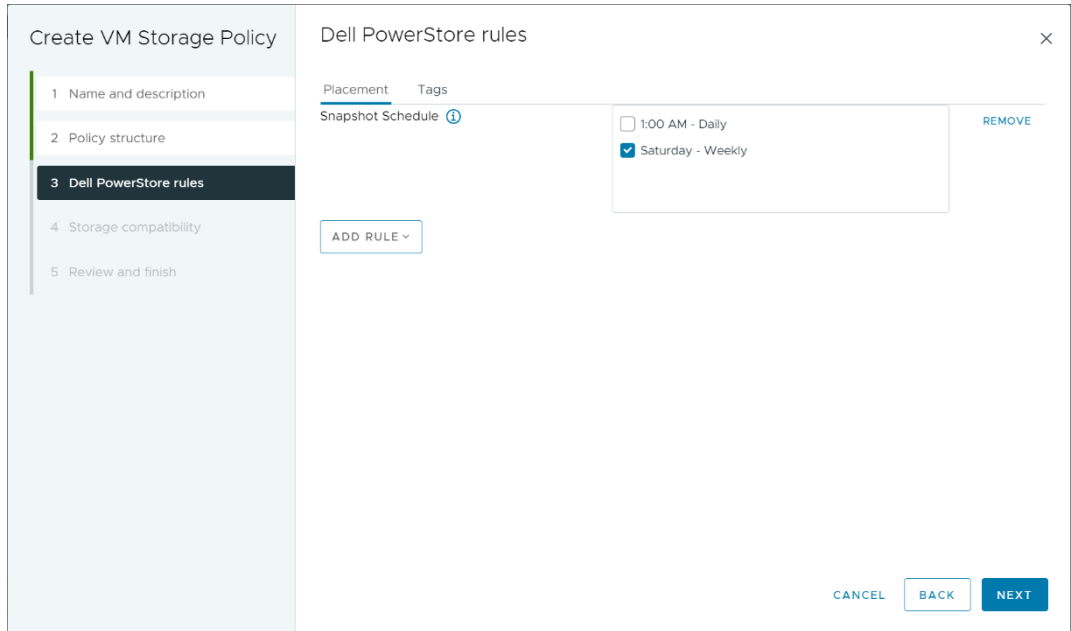


Figure 18. Create VM Storage Policy page

## 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 page includes a list of VMs including the name, operating system, CPUs, memory, and more, as shown in the following figure:

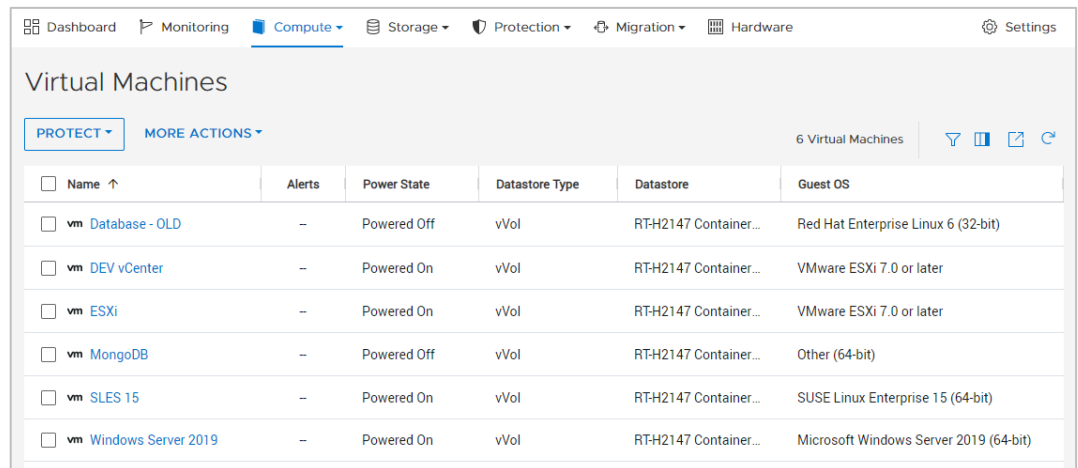


Figure 19. Virtual Machines page

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:

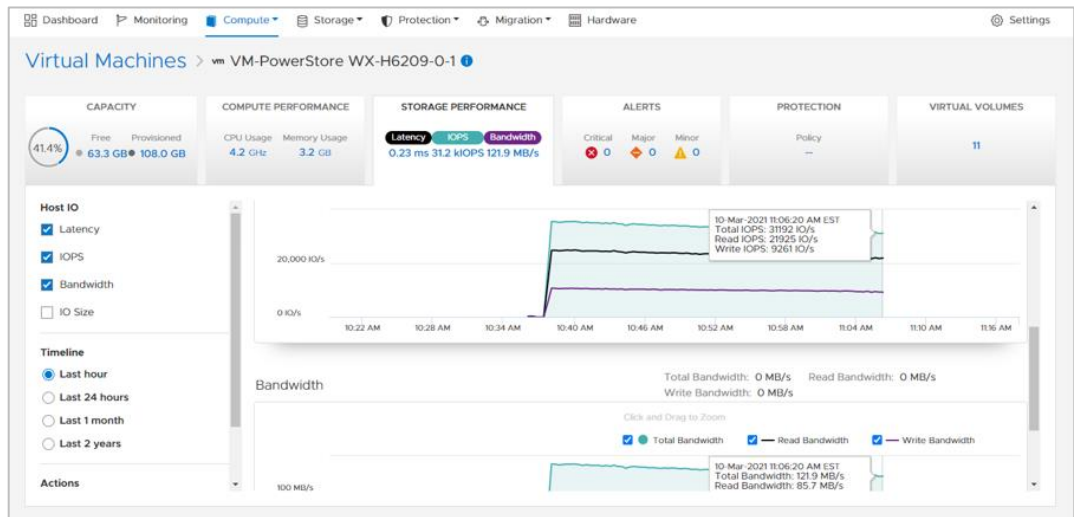


Figure 20. VM storage performance

Starting in PowerStoreOS 3.0, you can determine the type of backing storage for virtual machines with the new **Datastore Type** column (see Figure 21). 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**.

Virtual Machines						
PROTECT ▾		MORE ACTIONS ▾		6 Virtual Machines		
<input type="checkbox"/> Name ↑	Alerts	Power State	Datastore Type	Datastore	Guest OS	
<input type="checkbox"/> vm Database - OLD	--	Powered Off	VMFS	RTH2147 VMFS(RT...	Red Hat Enterprise Linux 6 (32-bit)	
<input type="checkbox"/> vm DEV vCenter	--	Powered On	Mixed	RTH2147 Container...	VMware ESXi 7.0 or later	
<input type="checkbox"/> vm ESXi	--	Powered On	vVol	RTH2147 Container...	VMware ESXi 7.0 or later	
<input type="checkbox"/> vm MongoDB	--	Powered On	VMFS	RTH2147 VMFS(RT...	Other (64-bit)	

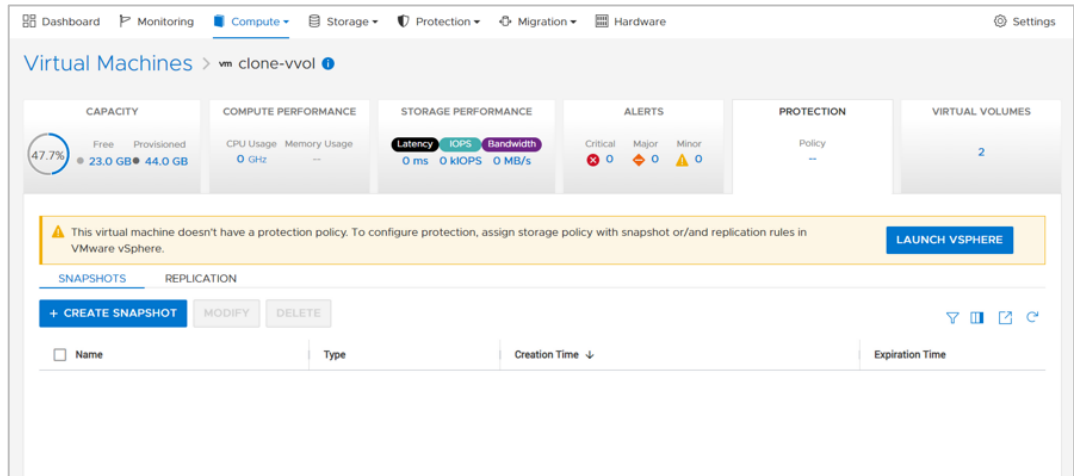
Figure 21. VM datastore type

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.

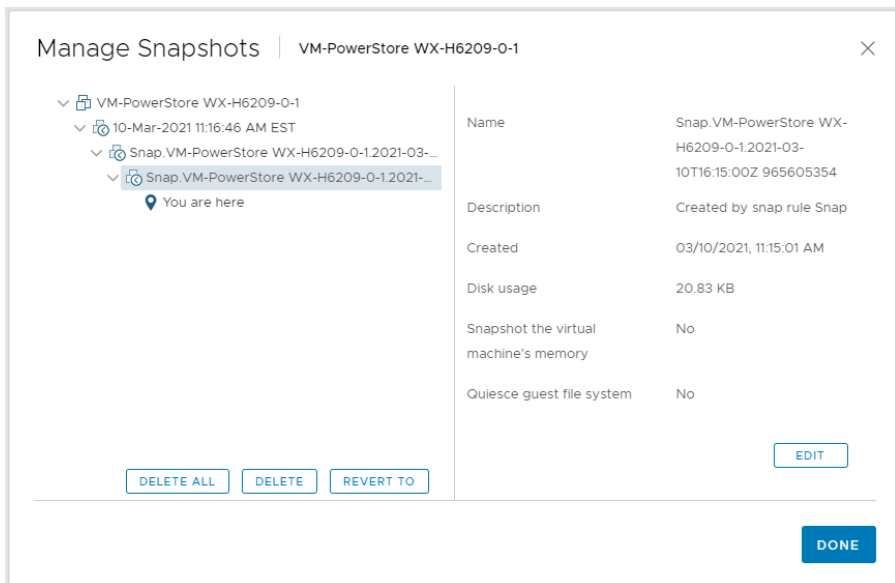
The following figure shows the VM protection page where you can configure snapshots and protection policies:



**Figure 22. 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 23. 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:

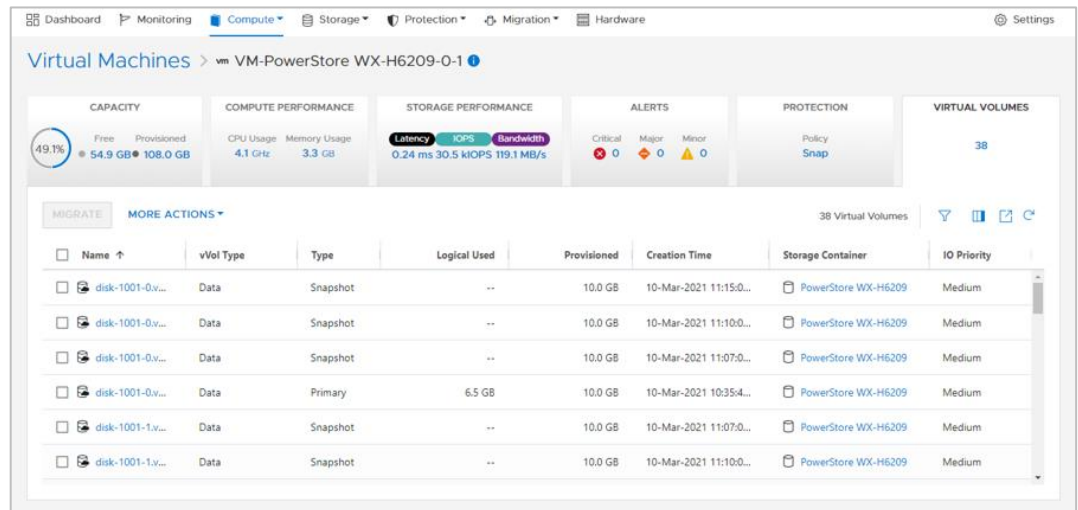


Figure 24. Virtual Volumes

### Virtual Volume migration

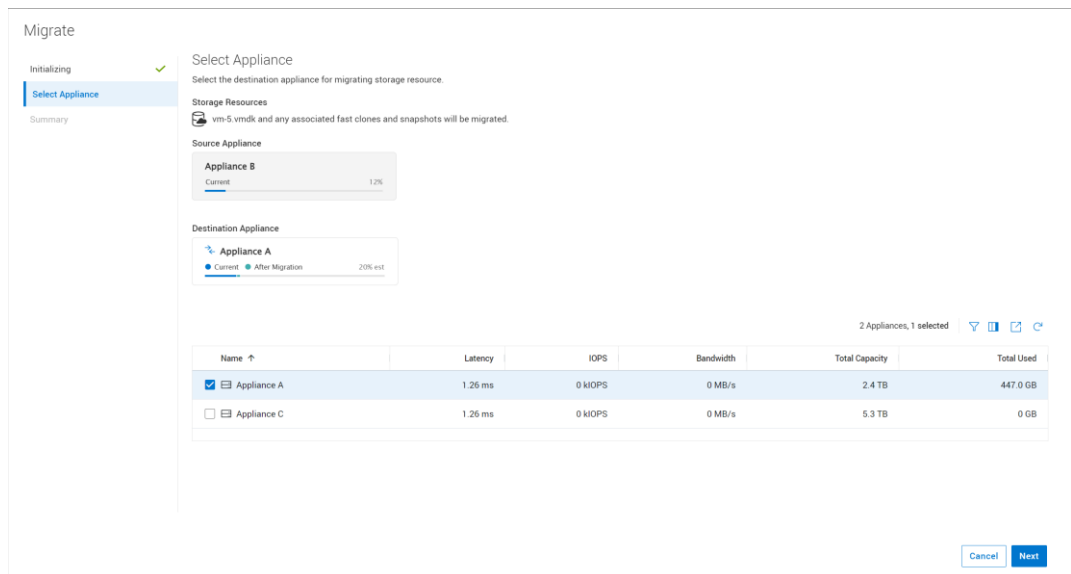
In PowerStoreOS 1.0, vVols can be migrated between appliances within the cluster. However, it is limited to vVols that are not in use so the virtual machine must be powered off before any of its vVols can be migrated. Starting with PowerStoreOS 2.0, online vVol migration is supported. This functionality 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 25. 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](#).

### Virtual Volume storage placement

To enable optimal virtualization performance, accounting for a virtual machine's storage placement is important. This section provides placement recommendations when using vVol storage.

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 afterward. 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 onto the same appliance.

### Virtual Volume node affinity

Each vVol has an attribute that is designated by PowerStore's Resource Balancer called node affinity. The node affinity specifies the preferred path to the vVol, either through Node A or Node B. Paths on one node are designated as active/optimized while the paths on the other node are designated as active/non-optimized. Because both nodes on PowerStore are active and have full access to the vVol, the non-optimized paths can continue servicing IO if all optimized paths become unavailable. The architecture enables both load balancing and high availability.

In PowerStoreOS 4.0, dynamic node affinity is available for vVols. This provides Resource Balancer the ability to dynamically change a vVol's node affinity based on performance metrics. This allows the appliance to maintain relatively consistent utilization, latency, and performance across both nodes. This has several benefits including balancing hardware utilization, optimizing performance, and providing the ability to adapt as workloads change.

This feature is designed to be completely transparent to the administrator and clients. The system makes changes automatically and intelligently as needed. It is enabled automatically after upgrading to PowerStoreOS 4.0 or later. It does not require any user intervention and is non-disruptive to clients.

Performance metrics are collected every five minutes and the rebalance operation happens every 30 minutes. For a rebalance to be initiated, the following criteria must be met:

- One node must have at least 50% or higher CPU utilization
- One node must have a sustained 20% or more deviation from the ideal balance of 50/50
- An internal latency check must pass to confirm that the rebalance is beneficial

If these checks pass, PowerStore attempts to first migrate volumes to achieve balance. If no further balancing can be done on volumes and there is still an imbalance, vVol node affinity can be changed in the next cycle. One vVol can be changed per cycle. Resource Balancer aims to achieve an even balance between the two nodes based on the throughput and bandwidth workload, regardless of the resource type.

When a vVol rebalance is initiated, PowerStore coordinates with all of the ESXi hosts to update the peer protocol endpoint to be the active/optimized path. When complete, this results in ESXi directing IO to the protocol endpoint on the peer node. ESXi automatically quiesces IOs until the pathing update is complete.

Data and config vVols are eligible for balancing. Data vVols store data, such as VM hard disks, snapshots, full clones, and fast clones and generally have the most IO. Config vVols store standard VM configuration data, such as .vmx files, logs, and NVRAM.

Some vVols are excluded from balancing. vVols that are unbound are not in use, such as a powered-off VM, so these are excluded from balancing. Also, vVols that have a migration in progress or were recently moved are excluded to avoid interference with ongoing migrations or moving resources back too quickly.

## 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.

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.

PowerStore VMFS datastores can be used to provision clustered virtual disks for Windows Server Failover Clustering (WSFC) deployments. VMware supports accessing clustered

VMDKs through FC and NVMe/FC. For more information about clustered virtual disks, see the VMware vSphere documentation.

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.

## VAAI

### Overview

vSphere Storage APIs - 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

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 from one ESXi host to another. 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:

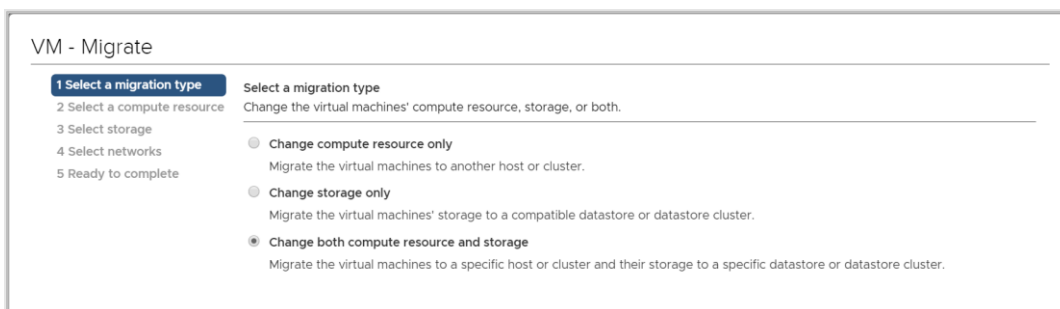


Figure 26. Compute and Storage vMotion

## Metro Volume

Metro Volume is a high availability and data mobility feature for PowerStore storage that works with 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

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:

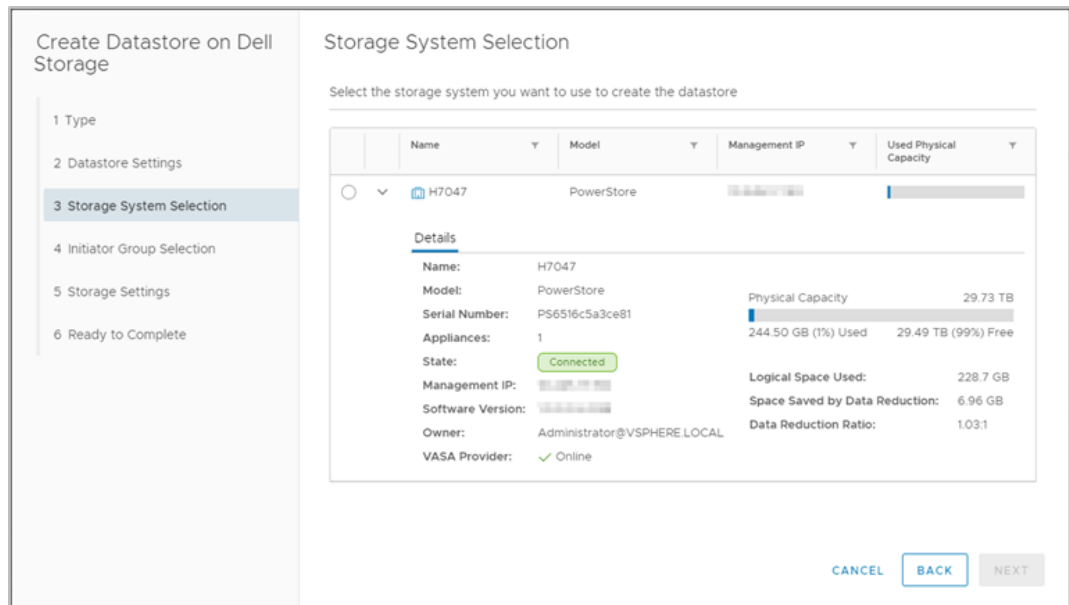


Figure 27. Creating a datastore using VSI

## 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:



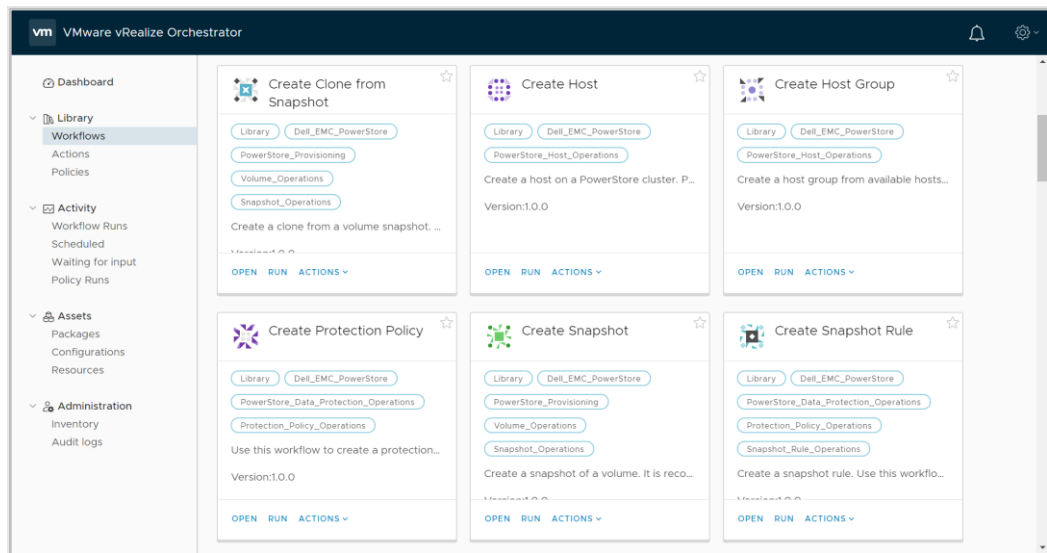


Figure 28. vRealize Orchestrator

## 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:

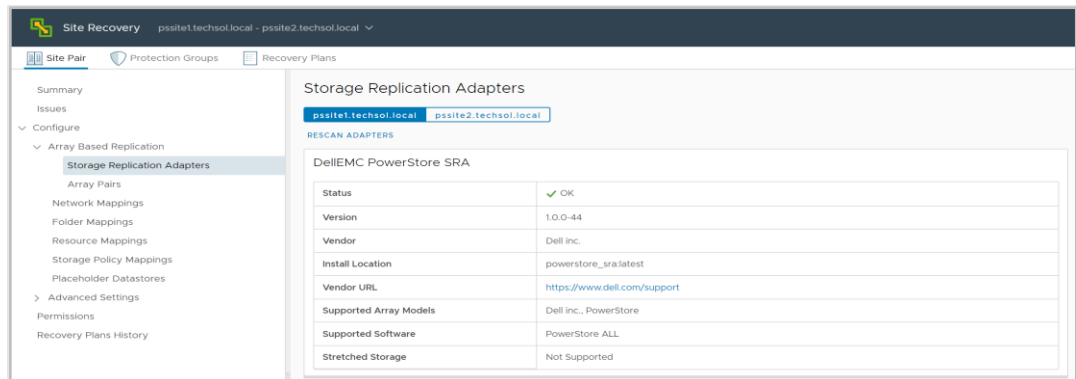


Figure 29. PowerStore SRA

## 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

## Conclusion

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

## Conclusion

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. 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.

## Technical support and resources

### Resources

The [Dell Technologies Storage Info Hub](#) 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.