

Using VMware vSphere Virtual Volumes and VASA with Dell PowerMax

A practical guide to implementing and working with vVols and the VASA Provider on Dell PowerMax

Abstract

This white paper discusses how VMware's vVols® are implemented on the PowerMax® using the Dell VASA Provider, SRDF replication, and VMware Live Site Recovery®.

October 2025

Revisions

Date	Description
October 2023	Initial release for VASA 4 from baseline VASA 3
April 2024	Troubleshooting for VASA registration
October 2024	VASA 5, PowerMaxOS 10.2.0
July 2025	vVol EOL
October 2025	PowerMaxOS 10.3.0

Acknowledgments

Author: Drew Tonnesen, Dell Engineering

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.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

Introduction

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © October 2025 Rev 1.0.4 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. Other trademarks may be trademarks of their respective owners. [9/23/2025] [Deployment and Configuration] [h19812]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	4
VMware Virtual Volumes end-of-life	9
1 Introduction.....	10
1.1 SRDF Overview	10
1.2 Minimum Support Level.....	11
1.3 Platforms.....	11
1.4 Protocols.....	12
1.4.1 IPV6	12
1.5 Audience.....	12
2 Virtual Volume Management.....	13
2.1 Storage Administrator role.....	13
2.2 VMware Administrator role	13
2.3 Upgrade/Migration	14
2.4 Replication	14
2.4.1 vVols 2.0 with PowerMax	14
2.4.2 SRDF	14
2.5 Data Reduction	15
2.5.1 PowerMax 2500/8500.....	18
2.6 Virtual Storage Integrator (VSI)	19
2.7 vSphere Cluster Services (vCLS).....	20
3 VAAI and vVols.....	22
3.1 ATS.....	22
3.2 Full Copy (XCOPY)	22
3.3 Block Zero (WRITE SAME)	22
3.4 Reclaim (UNMAP)	22
3.5 Thin Provisioning Out of Space	23
4 VASA Configuration	24
4.1 Installation.....	24
4.1.1 High Availability	24
4.1.2 Disaster Recovery	24
4.1.3 Backup.....	24
4.2 Management.....	24

4.2.1	PowerMax 2000/8000	24
4.2.2	PowerMax 2500/8500	28
5	Security Certificates	31
5.1	VASA 3 and 4	31
5.2	VASA 5	31
6	Registering the VASA Provider in vCenter	32
6.1	User authentication for registering the VASA Provider	32
6.2	VP registration wizard in VASA 3 and 4	35
6.3	VP registration wizard in VASA 5	38
6.3.1	Virtual host	41
6.3.2	Storage container provisioning and de-provisioning	42
7	Certificates in VASA 3 and 4	47
7.1	Signed certificates	47
7.2	Default certificate	48
7.3	Multiple vCenters and certificates	50
7.3.1	Unregister the VASA Providers	50
7.3.2	Retain VP Certificate	50
7.3.3	Self-signed certificate	50
7.3.4	CA certificate	53
8	Certificates in VASA 5	54
8.1	Signed certificates	54
8.1.1	VC trust store	56
8.2	Multiple vCenters	57
9	Configuring Virtual Volumes on the PowerMax	58
9.1	Using Unisphere with Virtual Volumes	58
9.1.1	VASA Provider Status	58
9.2	Creating the Storage Container in Unisphere	58
9.2.1	SC creation wizard	59
9.3	Creating the VASA Replication Group in Unisphere	63
9.3.1	VRG State	67
9.4	Provisioning the Protocol Endpoint in Unisphere	67
9.4.1	PE wizard	69
9.5	Using Solutions Enabler with Virtual Volumes	75
9.5.1	Creation of the Storage Container	75
9.5.2	Creation of the VASA Replication Group	77
9.5.3	Creation of the SCSI PE	78

9.5.4	Creation of the NVMe vPE	79
9.6	Host IO Limits/Storage IO Control (SIOC).....	79
10	Creating a vVol Datastore	80
10.1	vVol datastore wizard	80
10.2	Modifying the Storage Container in Unisphere.....	83
10.2.1	SC modification steps	83
10.2.2	Recognizing new SC size in vCenter.....	86
10.3	Out of space errors	88
11	Creating a VM Storage Policy for vVols	90
11.1	VM Storage Policy with SRDF Replication	90
11.1.1	VM Storage Policy wizard	91
11.2	Storage policy component	96
11.3	Creating a replicated VM with vVol storage	98
11.3.1	VM Creation	98
11.4	Changing VM Storage Policy for a VM	105
11.4.1	Change Storage Policy wizard	105
11.5	VMware CLI for vVols	107
11.5.1	Stats	109
11.5.2	Vmstats	112
11.5.3	vVol datastores in a cluster	112
11.5.4	VMware High Availability (HA)	114
11.5.5	Default profile/capability sets and default Storage Policy for vVol datastores.....	114
12	VMFS and vVol Cloning/Migrations	118
13	vVol identification and monitoring in Unisphere	119
13.1	Identifying vVol WWN in Unisphere.....	119
13.2	vVol Performance Monitoring in Unisphere.....	121
14	Scalability	123
14.1	ESXi	123
14.2	Storage Resources	123
14.3	VM snapshot sizing	126
14.3.1	Setup	126
14.3.2	Snapshot preserved space	127
14.4	Storage Demand Report.....	132
14.4.1	PowerMax 2000/8000	132
14.4.2	PowerMax 2500/8500	136
14.4.3	Solutions Enabler Storage Group Demand Report.....	137

14.5	Queueing	139
14.5.1	Adjusting the PE queue.....	139
14.6	PowerMax code upgrade.....	140
14.7	vVol scaling.....	141
15	Using Virtual Volumes with VMware Live Site Recovery	142
15.1	VMware Live Site Recovery	142
15.1.1	VMware PowerCLI	142
15.2	Protocols.....	142
15.3	SRM restrictions	142
15.3.1	VMware	142
15.3.2	Dell	143
15.4	VM Replicated Objects	143
15.4.1	Compliance	143
15.5	Configuration	147
15.5.1	Placeholder datastores and Storage Policy Mappings	147
15.6	Create a Protection Group and Recovery Plan	150
15.7	TestFailover	153
15.8	Cleanup	154
15.9	Failover (Planned Migration or Failover)	154
15.10	Reprotect	154
15.10.1	Reprotect error.....	155
15.11	Configuration Monitoring	155
16	Conclusion.....	157
A	Appendix: Virtual Volume Troubleshooting	158
A.1	Directories.....	158
A.1.1	Creating a config vVol greater than the default.....	158
A.2	VASA Provider registration issues.....	164
A.2.1	Clock synchronization.....	164
A.2.2	Firewall	165
A.3	Cascaded Groups/Multi-host initiator groups	165
A.4	vVol datastore access issues	166
A.4.1	vVol datastore reports inaccessible.....	166
A.4.2	vVol shows as zero bytes	166
A.5	vCenter failure to recognize replication or service level.....	167
A.6	VASA Provider recovery	167
A.7	Orphaned virtual volumes.....	167

A.8	Changing Storage Policies with multi-writer flag	167
A.9	SRM TestFailover errors	167
A.10	SRM Failover	168
A.11	VASA Provider rescan errors	169
A.12	Timeouts	169
B	Appendix: Virtual Volume Operational Detail	170
B.1	Adding a TCP software adapter in vCenter	170
B.1.1	Add array controller	170
B.1.2	Add host.....	174
B.2	Obtaining the Host NQN and Host ID.....	177
B.3	vVol Config UNMAP	178
C	Appendix: Technical support and resources	181
C.1	Dell.....	181
C.2	Broadcom/VMware	181

VMware Virtual Volumes end-of-life

With the release of VMware vSphere Foundation (VVF) 9.0 and VMware Cloud Foundation (VCF) 9.0, Broadcom deprecated support for VMware Virtual Volumes (vVols). The feature will be completely removed from these products in version 9.1. Support will continue for vSphere 8.x and VCF 5.x until those products reach end-of-life support. Only critical bug fixes will be addressed. There will be no new features.

Existing PowerMax vVol implementations are not immediately impacted; however, it is likely that the VASA Provider will be deprecated in a future PowerMaxOS release to align with Broadcom. Customers are encouraged to review their use of vVols on the PowerMax as Broadcom will only offer limited support for vVols in VVF and VCF 9.0 on a case-by-case basis. If vVols are critical to the business, you are encouraged to contact your Broadcom representative about the possibility of continuing to use vVols in 9.0. Note that this would still be a temporary solution as there will be no support beyond 9.0 for any customer.

For customers wishing to migrate their virtual machines off vVol storage, use Storage vMotion. It is not possible to use SRDF or local TimeFinder replication technologies. You can move VMs from vVols to VMFS or NFS. This is an online task. Once all VMs and files are removed from a vVol datastore, it can be unmounted from the ESXi hosts. Customers may also choose to unregister the VASA Providers after unmounting all associated vVol datastores. The protocol endpoints (PE) and masking views can be deleted also. Removal of the VASA containers on the array, if desired, requires the assistance of Dell support.

1 Introduction

VMware vSphere® Virtual Volumes (vVols) using the embedded VASA Provider (EVASA), is available on PowerMax arrays running a minimum of PowerMaxOS 5978 Q3 2020 (5978.669.669). vVols are an integration and management framework (referred to as the vVol framework) that moves management from the datastore to the virtual machine. vVols virtualize storage, enabling a more efficient operational model that is optimized for virtualized environments and centered on the application instead of the infrastructure. vVols map virtual disks, configuration files, clones, and snapshots, directly to virtual volume objects on a storage system. This mapping allows VMware to offload more storage operations to the PowerMax, such as VM snapshots.

vVols virtualize SAN devices by abstracting physical hardware resources into logical pools of storage called Storage Containers which replace traditional storage volumes. These containers are created on the PowerMax and are assigned storage capabilities which can be all, or a subset, of the Service Level Objectives¹ the PowerMax array supports. These storage containers are presented to the vCenter when a VASA Provider is registered. The VASA Provider runs directly on the array as a GuestOS and integrates with vSphere Storage Monitoring Service (SMS) and ESXi vvolld service to manage all aspects of the storage in relation to vVols. All communication between VMware and the VASA Provider is out of band (IP).

From the presented storage containers, the VMware Admin creates vVol datastores through the wizard, just like a traditional datastore. A vVol datastore has a one-to-one relationship with a storage container and is its logical representation on the ESXi host. A vVol datastore holds VMs and can be browsed like any datastore. Each vVol datastore then will inherit the capabilities of the container, e.g., SLs, data reduction, replication, etc. The VMware Admin creates Storage Policies that are mapped to these capabilities so when a VM is deployed, the user selects the desired capabilities (if any) and is presented with compatible vVol datastores in which to deploy the VM.

Unlike traditional devices, ESXi has no direct device visibility to vVols on the PowerMax. Instead, ESXi hosts use a front-end access point called a Protocol Endpoint which may be physical (SCSI) or virtual (NVMe). A SCSI Protocol Endpoint (PE) is a special device created on the PowerMax and mapped and masked to the ESXi hosts. An NVMe Protocol Endpoint (vPE) is virtual yet provides the same type of functionality as a physical one. Each ESXi host requires a single, unique PE per array. The ESXi host uses the PE to communicate with the volumes and the disk files the vVols encapsulate. By utilizing PEs, ESXi establishes data paths from the virtual machines (VM) to the virtual volumes.

vVols simplify the delivery of storage capabilities to individual applications by providing finer control of hardware resources and allowing the use of native array-based data services such as SnapVX and SRDF® at the VM level. vVols present the VMware Admin a granularity of control over VMs on shared storage that cannot be achieved with traditional architecture due to the device limitations of the ESXi server. Note that as each vVol uses a host available address, each one counts toward the total number of devices on the array.

1.1 SRDF Overview

The Dell Symmetrix® Remote Data Facility (SRDF) family of products offers a range of PowerMax-based disaster recovery, parallel processing, and data migration solutions.

SRDF disaster recovery solutions are based on active remote mirroring and dependent-write consistent copies of data maintained at one or more remote locations. A dependent write is a write operation that cannot be issued by an application until a prior, related write I/O operation

¹ In more recent releases of the VMAX and the PowerMax, SL is renamed to SL or Service Level. Both terms will be used in this whitepaper interchangeably.

completes. Dependent-write consistency is required to ensure transactional consistency when the applications are restarted at the remote location.

SRDF configurations with vVols require at least two PowerMax systems. In two-site, non-SRDF/Metro configurations, these systems are usually also known as the primary and the secondary systems. Both sites can be located in the same room, in different buildings within the same campus, or hundreds to thousands of kilometers apart.

Dell supports a single mode of two-site replication with vVols, SRDF/Asynchronous.

SRDF/Asynchronous (SRDF/A) mode provides a dependent write consistent copy on the target (R2) device, which is only slightly behind the source (R1) device, though for vVols the RPO will be set at a maximum of 300 seconds. SRDF/A session data is transferred to the remote PowerMax array in predefined timed cycles or delta sets, which minimizes the redundancy of same track changes being transferred over the link. SRDF/A provides a long-distance replication solution with minimal impact on performance that preserves data consistency.

For more information on SRDF technology, including setup and configuration, refer to Dell Symmetrix Remote Data Facility (SRDF) documentation available at <http://support.dell.com>.

This paper will start with an explanation of how virtual volumes are managed including how the VMware storage APIs relate to vVols. There will also be a brief introduction to VASA before embarking on the technical details of the implementation, namely, how to configure and use the embedded VASA Provider with virtual volumes in VMware vSphere 8 and 9 environments with PowerMax storage arrays. This will include how to utilize the capabilities of the array like quality of service through service levels, and replication through SRDF. An understanding of the principles that are exposed here will allow the reader to deploy and utilize virtual volumes in the most effective manner.

Note: Dell does not support any mode of replication with vVols except asynchronous. Therefore, SRDF/S and SRDF/Metro are not supported. VMware enables vVol replication support for both asynchronous and active/active (if supported by the vendor).

1.2 Minimum Support Level

Support for VMware vVols with the embedded VASA Provider (EVASA or eVASA) requires a minimum of PowerMaxOS 5978 Q3 2020 release, Management Software Version 9.2, and vSphere 8.x.² This paper covers VASA 3, 4, and 5 up to PowerMaxOS 10.3.0.0 and vSphere 9.0. VASA 5 requires a minimum of vSphere 8.0 U1.

Note: Dell PowerMax VASA Provider does not support SCSI-3 persistent reservations (PR). These reservations are required when using Windows Server Failover Cluster (WSFC), and therefore that solution cannot be run on vVols with the PowerMax. Note that Oracle RAC does not require SCSI-3 PR and thus is supported.

1.3 Platforms

vVols with the VASA Provider is supported on two different PowerMax platforms. For simplicity in this paper, the models 2000 and 8000 will be referred to as the V3; and the models 2500 and 8500 will be known as the V4. Other than management of the VASA Provider and some screen

² vSphere 7 is now out of support and therefore vSphere 8 is the minimum required release with vVols.

differences, the implementation of Virtual Volumes is the same on the V3 and V4 platforms. Differences between the two will be called out in the paper if required.

1.4 Protocols

The PowerMax arrays supports FC, iSCSI and NVMe/TCP with vVols. NVMe/TCP requires vSphere 8.0 U1 and higher and PowerMaxOS 10.1.0.0 and higher. PowerMax supports using both SCSI and NVMe on the same ESXi host. Such configurations would be common for protocol migrations.

Note: NVMe/FC is not supported with vVols on the PowerMax.

1.4.1 IPV6

PowerMaxOS 10.3.0 adds support for IPv6 with the VASA Provider on PowerMax 2500 and 8500 systems; however, note that with the deprecation of vVols this is not a Broadcom certified configuration and thus not recommended.

1.5 Audience

This technical white paper is intended for VMware administrators and PowerMax administrators responsible for deploying VMware vSphere 8.x and 9.x, with VMware vVols, on the PowerMax.

Note: Due to the nature of development, the exact minor revisions of products in this paper may not match those available to customers at the time of general availability (GA). Note, however, that every effort was made to ensure the functionality between the versions in this paper and those at GA are the same. In addition, screenshots come from various supported releases, however all current functionality is included.

2 Virtual Volume Management

vVols involve both a storage role and a VMware role. In some companies these two roles are consolidated, but in most large enterprises these are distinct positions, and as such there is a desire for bifurcation of tasks when it comes to virtualization and storage. vVols offer that separation with the storage administrator (SA) maintaining control over the physical storage requirements as well as where that storage is made available, and how much. The VMware administrator, meanwhile, maintains the ability to create VMs and select from available service levels (SL) that the SA has presented to the storage container(s).

2.1 Storage Administrator role

The SA is responsible for three main tasks:

1. Provision and present PE(s) to the ESXi host(s).
2. Create storage container(s) and assign storage resources in the desired storage amounts and SLs.
3. Create VASA Replication Groups for the storage containers to enable replication. (optional)

These tasks can be accomplished through Unisphere for PowerMax, or Solutions Enabler. Unisphere is recommended as the wizards provide an easy interface to deploy the objects. There is an additional capability available to the SA, and that is general monitoring of the VASA Provider (VP). Unisphere will monitor the viability of the VASA Provider by making direct calls to the array automatically, without any configuration from the SA.

All tasks are covered in the sections Using Unisphere with Virtual Volumes and Using Solutions Enabler with Virtual Volumes.

One thing to note with vVols is that the SA has no control over local or remote replication of the vmdks or virtual machines. Snapshots that utilize TimeFinder technology can only be accessed through the vSphere interface. It is not possible, for instance, to use Unisphere or Solutions Enabler to snapshot a vVol device. Similarly, devices that are replicated with SRDF as part of the VM are controlled through the VASA Provider and cannot be manipulated outside of that.

2.2 VMware Administrator role

Once the PE is presented to the ESXi hosts and the Storage Administrator completes the other tasks, the VMware administrator is then able to proceed with his/her tasks:

1. Register the embedded VASA Providers in VMware vSphere vCenter.
2. Create vVol datastores from presented storage container(s).
3. Create Storage Policies for each service level supported and advertised by the system and add replication if configured and desired.
4. Create virtual machines.
5. Create snapshots. (optional)
6. Configure VMware Site Recovery Manager. (optional)

The VMware admin manages the lifecycle of virtual machines, including snapshots, clones, fast clones, replication, etc., through the vSphere Client attached to the vCenter. For example, when the

VMware admin takes a VM snapshot, TimeFinder SnapVX³ technology is utilized so the snapshot is far more efficient than the traditional VMware implementation. VMware no longer has to keep track of multiple delta files which can grow beyond the original size of the VM and impact performance. Each snapshot will create no more than a single vVol (only if memory is included, otherwise the snapshot is targetless) and is maintained by Dell on the PowerMax.

Dell limits the number of snapshots to twelve (12), not including the source. VMware's limit is 31, not including the source. If the user attempts to take more than twelve snapshots, the error in [Figure 1](#) will be generated. While it is not possible to adjust this value directly, if the twelve-snapshot restriction is deleterious to the business, a customer may open a support ticket with Dell Support and a request made to increase it as there are implications to doing so.

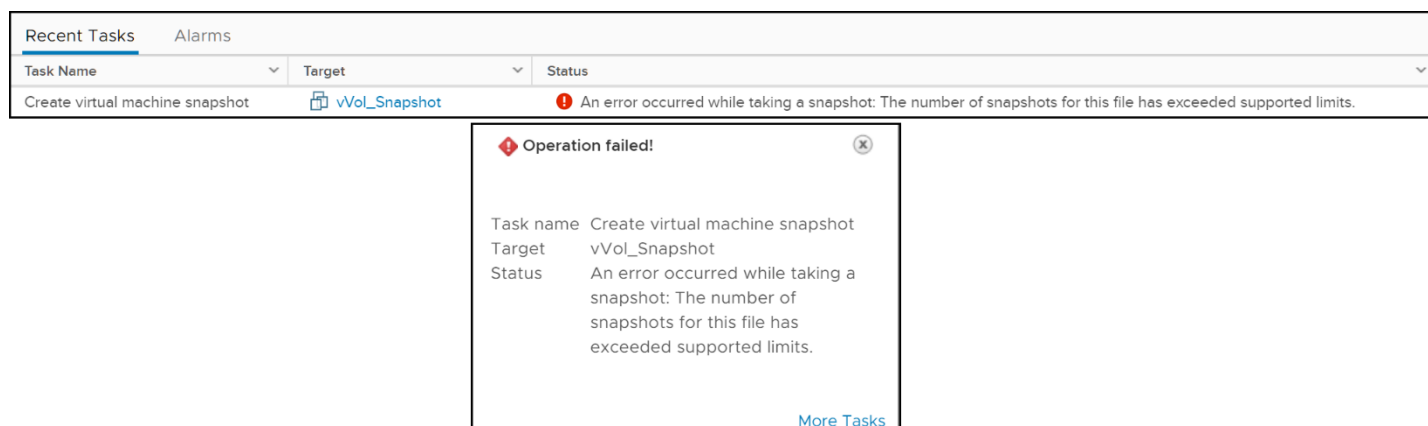


Figure 1. Error upon taking too many snapshots

2.3 Upgrade/Migration

There is no migration path from an external VASA Provider to an embedded VASA Provider, whether internal to an array or between arrays. Any VMs using vVol storage with an external VASA Provider should be migrated with Storage vMotion to VMFS storage first, then migrated to new storage containers created with the embedded VASA Provider.

2.4 Replication

2.4.1 vVols 2.0 with PowerMax

vVols 2.0 and higher is required for remote replication. Dell supports SRDF with vSphere 8.x and 9.x. Both SCSI and NVMe are supported.

2.4.2 SRDF

The PowerMax enables replication with virtual volumes through SRDF. VMware only supports asynchronous replication with vVols. The maximum Recovery Point Objective (RPO) is 300 seconds, or five minutes. It cannot be adjusted. The RPO is achieved by PowerMax taking a

³ TimeFinder must be licensed on the array or vVol snapshot creation will fail.

SnapVX targetless snapshot of each VASA Replication Group on the target array every 300 seconds. It will cycle through a total of five snapshots. When a testfailover or failover is executed, PowerMax will use the latest available snapshot, which will always have an RPO less than, or equal to, 300 seconds. Replication is controlled at the VM-level and cannot be manipulated outside of the vCenter interface. While replicated pairs are viewable through Solutions Enabler, no actions are permitted.

As with SRDF with non-vVols, bi-directional replication is supported, though separate VASA Replications Groups are required since the mode is asynchronous.

2.5 Data Reduction

Virtual Volumes supports data reduction (compression/deduplication) on the PowerMax arrays (PowerMax 2000/2500/8000/8500). Data reduction is enabled at the storage resource level as in Figure 2.

The screenshot shows the 'Add Storage Resources' window. A tab labeled '1 Storage Resources' is selected. Below it is a table with columns: Name, SRP, Service Level, Limit (GB), and Data Reduction. Three resources are listed: Demo_resource_1 (Gold, 1000 GB), Demo_resource_2 (Silver, 10.06 GB), and Demo_resource_3 (Platinum, 100 GB). All three have a checkmark in the Data Reduction column. A tooltip at the bottom right states: 'Compression will be enabled and Deduplication will be enabled where supported'. At the bottom, it shows 'Total Resource Subscribed Limit 1110.06 GB' and 'Total Resources 3'. There are 'CANCEL' and 'ADD TO JOB LIST' buttons at the bottom right.

Name	SRP	Service Level	Limit (GB)	Data Reduction
Demo_resource_1	SRP_1	Gold	1000	<input checked="" type="checkbox"/>
Demo_resource_2	SRP_1	Silver	10.06	<input checked="" type="checkbox"/>
Demo_resource_3	SRP_1	Platinum	100	<input checked="" type="checkbox"/>

Total Resource Subscribed Limit **1110.06 GB** Total Resources **3**

CANCEL **ADD TO JOB LIST**

Figure 2. Enabling Data Reduction on storage resource

Data reduction is a combination of compression and deduplication on the PowerMax. While compression is only applicable to data within a storage resource where it is active, deduplication applies to all data in an SRP when that data is in a storage resource or storage group with data reduction active. The following paragraph provides a visual example of how this feature works.

A single virtual machine, VM_1, is deployed from a template in storage resource A with data reduction active. Its OS vmdk (vVol) is 100 GB and identified as E3. The array recognizes that many tracks can be compressed and separates them into smaller extent pools. These are named DG1_F_x in Figure 3.

```

root@dsib2017:~
[root@dsib2017 ~]# symcfg -sid 355 list -tdev -dev E3 -detail

Symmetrix ID: 000197600355

Enabled Capacity (Tracks) : 425136600
Bound Capacity (Tracks) : 819210

-----
S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym      Bound      Flags      Total      Pool      Pool      Exclusive      Comp
Pool     Name      ESPT      Tracks    Subs      Allocated    Allocated    Ratio
                                         Tracks (%)   Tracks
000E3 -      FS.B      819210    -          0    0      60732    1.4:1
      DG1_F_1    -.--      -          -      3055    0        -        -
      DG1_F_5    -.--      -          -      26155   3        -        -
      DG1_F_7    -.--      -          -      14994   2        -        -
      DG1_F_8    -.--      -          -      9121    1        -        -
      DG1_F_9    -.--      -          -      8571    1        -        -
      DG1_F_F    -.--      -          -      80296  10        -        -
Total
Tracks      819210    -      142192  17      60732

Legend:
Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
      (S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
      (P)ersistent Allocs : A = All, S = Some, . = None
      S(T)atus      : B = Bound, I = Binding, U = Unbinding, A = Allocating,
                     D = Deallocating, R = Reclaiming, C = Compressing,
                     N = Uncompressing, F = FreeingAll, . = Unbound

[root@dsib2017 ~]#

```

Figure 3. Compressed/Deduplicated tracks in a vVol

In addition, the array is also able to remove tracks through deduplication. The total number of tracks allocated in a pool is 142192, however in the green circle in Figure 3, the number of tracks that are exclusive to this device are only 60732. This means that 81460 tracks are shared with other devices that have deduplication active in the SRP. Note that it does not matter whether those devices are vVols or regular TDEVs or even in use by VMware.

Now, suppose a second VM, VM_2, is deployed from template into storage resource B with data reduction active and that new vVol device in VM_2 is identified as E7. If the tracks are viewed for this new VM, note how in Figure 4 there is only a single exclusive track.

```

root@dsib2017:~
[root@dsib2017 ~]# symcfg -sid 355 list -tdev -dev E7 -detail

Symmetrix ID: 000197600355

Enabled Capacity (Tracks) : 425136600
Bound Capacity (Tracks) : 819210

S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym      Bound      Flags      Total      Pool      Pool      Exclusive      Comp
Pool      ESPT      Tracks      Sub      Allocated      Allocated      Tracks      Ratio
Name                                     (%)      Tracks (%)      Tracks
000E7 -      FS.B      819210      -      0      0      1      1.4:1
      DG1_F_1      -.--      -      -      3089      0      -      -
      DG1_F_5      -.--      -      -      26155      3      -      -
      DG1_F_7      -.--      -      -      14993      2      -      -
      DG1_F_8      -.--      -      -      9121      1      -      -
      DG1_F_9      -.--      -      -      8571      1      -      -
      DG1_F_F      -.--      -      -      80263      10      -      -
Total
Tracks      819210      -      142192      17      1

Legend:
Flags: (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
      (S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
      (P)ersistent Allocs : A = All, S = Some, . = None
      S(T)atus : B = Bound, I = Binding, U = Unbinding, A = Allocating,
      D = Deallocating, R = Reclaiming, C = Compressing,
      N = Uncompressing, F = FreeingAll, . = Unbound

[root@dsib2017 ~]#

```

Figure 4. Compressed/Deduplicated tracks in a cloned vVol

All but one track is now shared with E3. In fact, if E3 is viewed after this clone, all its exclusive tracks are gone because it has the reciprocal relationship with E7 (Figure 5).

```

root@dsib2017:~
[root@dsib2017 ~]# symcfg -sid 355 list -tdev -dev E3 -detail

Symmetrix ID: 000197600355

Enabled Capacity (Tracks) : 425136600
Bound Capacity (Tracks) : 819210

S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym      Bound      Flags      Total      Pool      Pool      Exclusive      Comp
Pool Name ESPT    Tracks  Subs      Allocated Allocated  Allocated  Ratio
                                (%)      Tracks  (%)      Tracks
000E3 -      FS.B      819210    -          0      0          0      1.4:1
      DG1_F_1    -.--      -          -          3055    0          -      -
      DG1_F_5    -.--      -          -          26155   3          -      -
      DG1_F_7    -.--      -          -          14994   2          -      -
      DG1_F_8    -.--      -          -          9121    1          -      -
      DG1_F_9    -.--      -          -          8571    1          -      -
      DG1_F_F    -.--      -          -          80296   10         -      -
Total
Tracks      819210    -          142192   17          0

Legend:
Flags:  (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
        (S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
        (P)ersistent Allocs : A = All, S = Some, . = None
        S(T)atus      : B = Bound, I = Binding, U = Unbinding, A = Allocating,
                        D = Deallocating, R = Reclaiming, C = Compressing,
                        N = Uncompressing, F = FreeingAll, . = Unbound

[root@dsib2017 ~]#

```

Figure 5. Removal of exclusive tracks after the second VM is deployed

2.5.1 PowerMax 2500/8500

On the PowerMax V4, the data reduction algorithm has been updated along with the hardware and software. As part of its efficiency, it no longer separates out tracks into different pools as shown in Figure 5. Instead, there is only a single pool, an example of which is provided in Figure 6.

```

root@dsib2005:~
[root@dsib2005 ~]# symcfg -sid 302 list -tdev -dev 10018 -detail

Symmetrix ID: 000120200302

      S Y M M E T R I X   T H I N   D E V I C E S
-----
Sym  Pool Name  Flags    Provisioned    Pool    Exclusive    Total    Data
      ESPT      Total    Effective    Effective    Unreducible    Reduction
      Tracks    Tracks    Tracks    Tracks    Tracks    Ratio
-----
10018 -      F...    2252805    0    0    -      0    1.0:1
      DG_01    -.--      -      500752    22    -      -      -
Total
Tracks    2252805    500752    0    -      0

Legend:
Flags:  (E)mulation : A = AS400, F = FBA, 8 = CKD3380, 9 = CKD3390
        (S)hared Tracks : S = Shared Tracks Present, . = No Shared Tracks
        (P)ersistent Allocs : A = All, . = None
        S(T)atus      : A = Allocating, D = Deallocating, R = Reclaiming,
                        F = FreeingAll, . = N/A
[root@dsib2005 ~]#

```

Figure 6. Data reduction pool in PowerMax 2500/8500

2.6 Virtual Storage Integrator (VSI)

There is no version of VSI that supports vVols. If VSI is installed in the vSphere Client, for instance, VSI will be unable to resolve any information about the datastore and will produce the standard message in [Figure 7](#): *This datastore is not on a registered Dell EMC storage system.*

Currently, from within vCenter, it is not possible to map a vVol vmdk back to the underlying storage device, however this can be done in Unisphere for PowerMax and is covered later in the paper.

Note that while VSI cannot interpret vVol datastores, beginning in VSI 10.6 it is possible to register the VASA Provider in the vCenter using the plug-in, though there are restrictions. Refer to the VSI Product Guide for more detail.

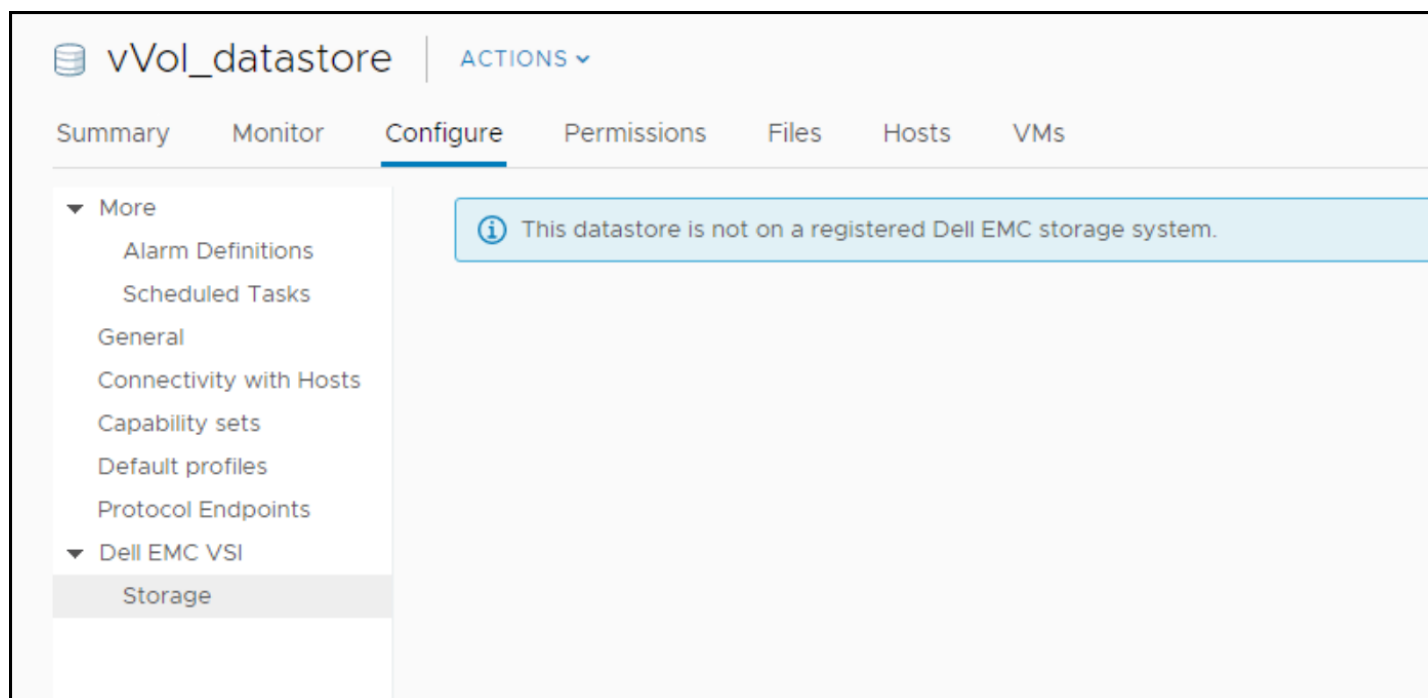


Figure 7. VSI with vVols

2.7 vSphere Cluster Services (vCLS)

vSphere Cluster Services or vCLS, is a feature that ensures the cluster services vSphere DRS and vSphere HA are all available regardless of the vCenter status. Currently, vCLS is a foundational release. Future releases will enable all functionality. vCLS is enabled automatically, meaning vSphere HA and DRS cannot operate without it. VMware creates one to three VMs, eponymously named, to manage each vSphere cluster. VMware places these VMs on datastores of its choosing, including vVol datastores.

On a VMFS datastore, the vCLS VMs are small in size - 2 GB vmdk, 128 MB swap – but when they are assigned to a vVol datastore it means the creation of three new devices on the array of the following size and type: 255 GB config, 2 GB data, and 128 MB swap. VMware creates and destroys these VMs as needed, which can impact array resources when the VMs are on vVol datastores. Furthermore, when VMware is unable to create vCLS VMs, it continues to try and may overwhelm the VASA Provider and lead to Orphaned virtual volumes. To avoid this, Dell recommends assigning the vCLS VMs to non-vVol datastores, either shared or local. This can be accomplished using the **Datastores** option under **vSphere Cluster Services** in the **Configure** tab of the vSphere Client as shown in [Figure 8](#).

Boston_Cluster

ACTIONS

Summary

Monitor

Configure

Permissions

Hosts

VMs

Datastores

Networks

Updates

Quickstart

General

Key Provider

VMware EVC

VM/Host Groups

VM/Host Rules

VM Overrides

I/O Filters

Host Options

Host Profile

Licensing

vSAN Cluster

Supervisor Cluster

Trust Authority

Alarm Definitions

Scheduled Tasks

vSphere Cluster Services

Datastores

Datastores

vSphere Cluster Services (vCLS) VM disks are placed by a default datastore selection logic. To override the default vCLS VM disk placement for this cluster, you can specify a set of Allowed datastores by clicking Add. Some datastores cannot be selected for vCLS 'Allowed' as they are blocked by solutions as SRM or vSAN maintenance mode where vCLS cannot be configured. Users cannot add/remove Solution blocked datastores for vCLS VMs. [Learn more](#)

VCLS ALLOWED

SOLUTION BLOCKED

ADD

REMOVE

<input type="checkbox"/>	Name	Type	Status	Capacity	Free
<input type="checkbox"/>	ALL_HOSTS_18F_302	VMFS 6	✓ Normal	8 TB	5.28 TB
<input type="checkbox"/>	ALL_HOSTS_DEV_10A_SID_302	VMFS 6	✓ Normal	4 TB	1.44 TB
<input type="checkbox"/>	TCP	VMFS 6	✓ Normal	511.75 GB	410.32 GB

Figure 8. vCLS datastore assignment

3 VAAI and vVols

With vVols, there is less reliance on VAAI primitives, either because they are not needed or the VASA APIs are more efficient and used instead; however, VAAI and thin-provisioning primitives do co-exist with vVols.

For vVols, ESXi will generally try to use the VASA or VAAI API primitives as the default behavior. If these are not supported, it will fall back to software.

This paper will not cover the VAAI primitives in detail, rather it will focus on when they might be used in a vVol environment. For an in-depth look at the VAAI primitives see the following whitepaper: [Using VMware vSphere Storage APIs for Array Integration with Dell PowerMax](#).

3.1 ATS

All config vVols (represented as directories in a vVol datastore) are formatted with VMFS-6 and therefore support ATS commands. This support is detected based on ATS support for a PE LUN to which vVols are bound.

As ATS is a SCSI command, the equivalent used for an NVMe vVol implementation is Compare and Write.

3.2 Full Copy (XCOPY)

The PowerMax does not support XCOPY commands with vVols. In most use cases the VASA APIs and host-based copy are utilized. A listing of functions and the commands used is available in the section VMFS and vVol Cloning/Migrations.

NVMe does not have an equivalent for XCOPY.

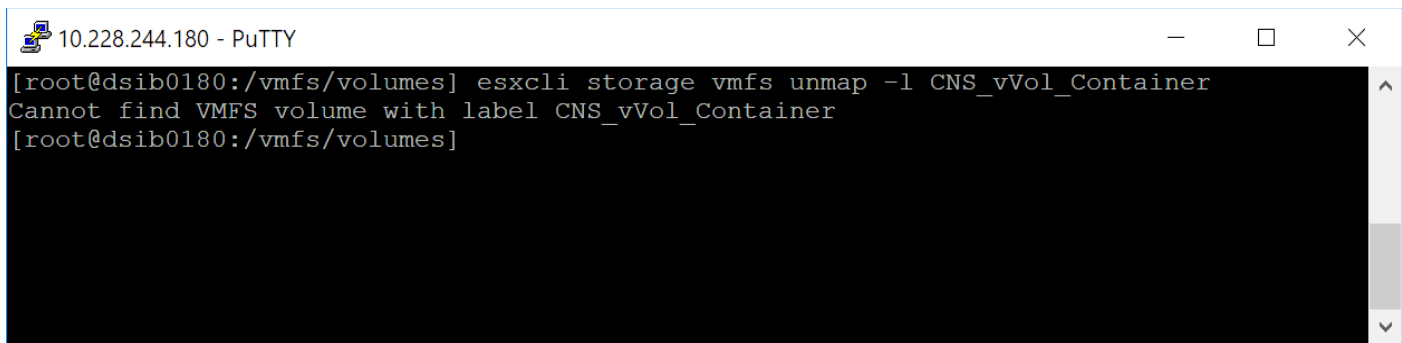
3.3 Block Zero (WRITE SAME)

Block Zero is most commonly used to zero out space on VMFS – in particular for eagerzeroedthick disks. As vSphere does not use eagerzeroedthick for vVols, the use of WRITE SAME is limited to zeroing extents when writing to or clearing storage.

As WRITE SAME is a SCSI command, the equivalent used for an NVMe vVol implementation is Write Zeroes.

3.4 Reclaim (UNMAP)

In traditional VMFS, the *unmap* command can be issued at the datastore level (manual or auto) on the ESXi host to reclaim space vacated by disk (VM) deletion or Storage vMotion. Since vVols are basically individual vmdk files, it is not possible to issue UNMAP against the vVol datastore. If an attempt is made to run reclaim on a vVol datastore, the error in [Figure 9](#) will be returned, indicating the datastore is not of the VMFS type.



```

10.228.244.180 - PuTTY
[root@dsib0180:/vmfs/volumes] esxcli storage vmfs unmap -l CNS_vVol_Container
Cannot find VMFS volume with label CNS_vVol_Container
[root@dsib0180:/vmfs/volumes]

```

Figure 9. UNMAP attempt on vVol datastore

The fact that UNMAP does not work on vVol datastores is of no great concern. This is because a vVol is both a vmdk and a device on the PowerMax array. When you delete a vVol, therefore, the PowerMax device (TDEV) is also deleted and the space it once occupied is immediately freed from the Storage Resource Pool (SRP).

Because vVols are vmdks and by default of thin type, they do support UNMAP at the Guest OS level. There are some restrictions/prerequisites that must be met on the VM and the Guest OS. These can be found in VMware article 301290. This functionality ensures that no space on the array is wasted.

Beginning with vSphere 8.0.2, UNMAP can be issued on config vVols. See vVol Config UNMAP in the appendix.

As UNMAP is a SCSI command, the equivalent used for an NVMe vVol implementation is Deallocate.

3.5 Thin Provisioning Out of Space

Storage container “out of space” warnings will be advertised to vSphere.

4 VASA Configuration

The VASA 3, 4 and 5 Providers are only available as an embedded GuestOS on the PowerMax array. VASA 3 requires PowerMaxOS 5978 Q3 2020 or PowerMaxOS 10.0.0.0 and 10.0.1.0. VASA 4 requires PowerMaxOS 10.1.0.0. VASA 5 requires PowerMaxOS 10.2.0.0 and higher. As with all GuestOS installations on the array, there are two instances of the VASA Provider for high availability. Both providers are registered in the vCenter with one recognized as primary and one as standby.

4.1 Installation

The embedded VASA Providers come pre-configured on the PowerMax array. Customer IPs are supplied prior to installation and will be used during the registration process in vCenter and to perform any required maintenance activity on the providers (e.g., user creation).

Note: There is no external VASA Provider for PowerMax.

4.1.1 High Availability

The VASA Provider is comprised of two instances on the array, and thus already in an HA configuration. The PowerMax handles the assignment of roles (active/standby) and will seamlessly switch between them as needed. Likewise, when both providers are registered in the vCenter, VMware is able to determine the active provider and send the command(s) appropriately.

4.1.2 Disaster Recovery

There is no disaster recovery solution for the VASA Provider. If both instances on the array fail or are otherwise unreachable, Dell support should be contacted immediately.

4.1.3 Backup

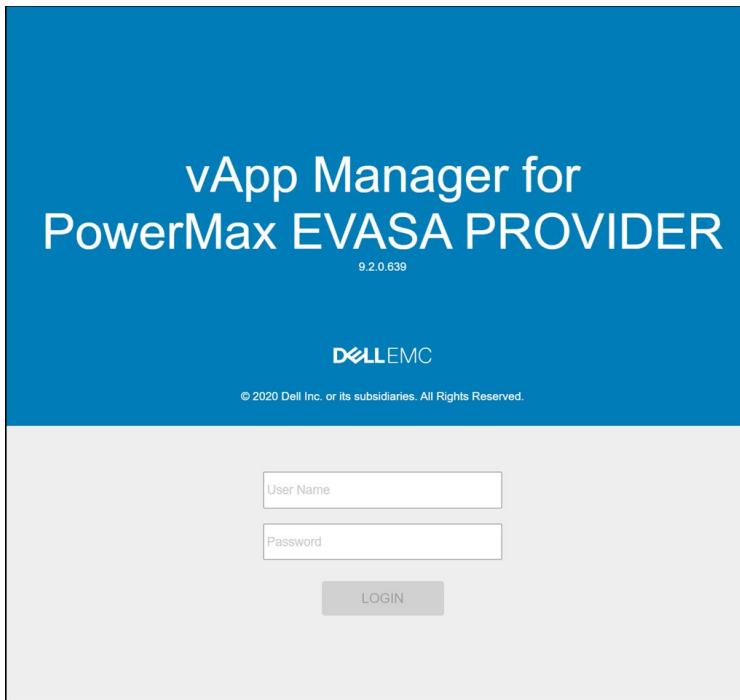
It is not possible to back up the VASA Provider.

4.2 Management

The management of the VASA Provider differs between the PowerMax platforms. Each is covered separately below.

4.2.1 PowerMax 2000/8000

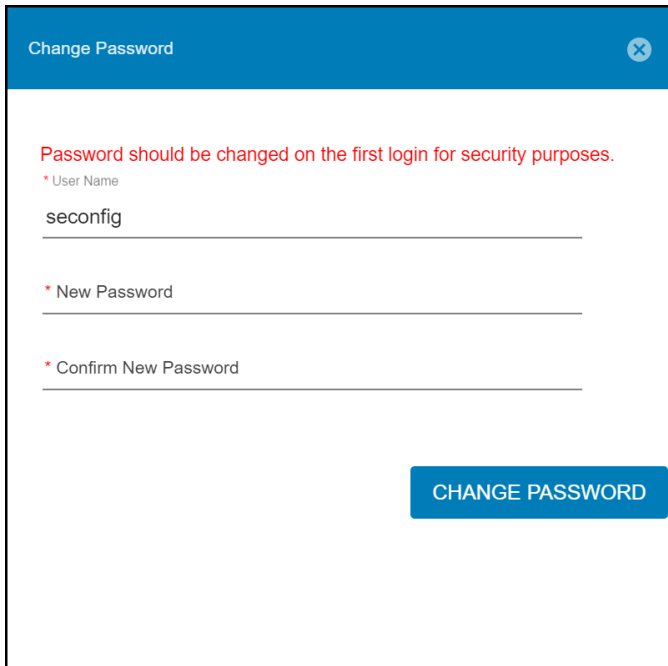
During initial login to the embedded VASA Provider, the user will be instructed to change the password. Navigate to **<https://<FQDN>:5480>** which is the login screen for the vApp Manager seen in [Figure 10](#):



The image shows the login interface for vApp Manager for PowerMax EVASA PROVIDER. The header is blue with the title "vApp Manager for PowerMax EVASA PROVIDER" in white, followed by the version "9.2.0.639" and the Dell EMC logo. Below the header, on a light gray background, are two input fields labeled "User Name" and "Password", and a "LOGIN" button.

Figure 10. VP vApp login

The initial login is *seconfig/seconfig* but a password change will immediately be requested as in Figure 11.



The image shows a "Change Password" dialog box. It has a blue header with the title "Change Password" and a close button. The main area is white and contains a red message: "Password should be changed on the first login for security purposes." Below this are three input fields, each with a red asterisk and a label: "* User Name" (containing "seconfig"), "* New Password", and "* Confirm New Password". A blue "CHANGE PASSWORD" button is at the bottom right.

Figure 11. VP vApp password change

Upon login, the dashboard presents the system information, network, SYMMETRIX array(s), authority status, and disk usage which is seen in Figure 12.

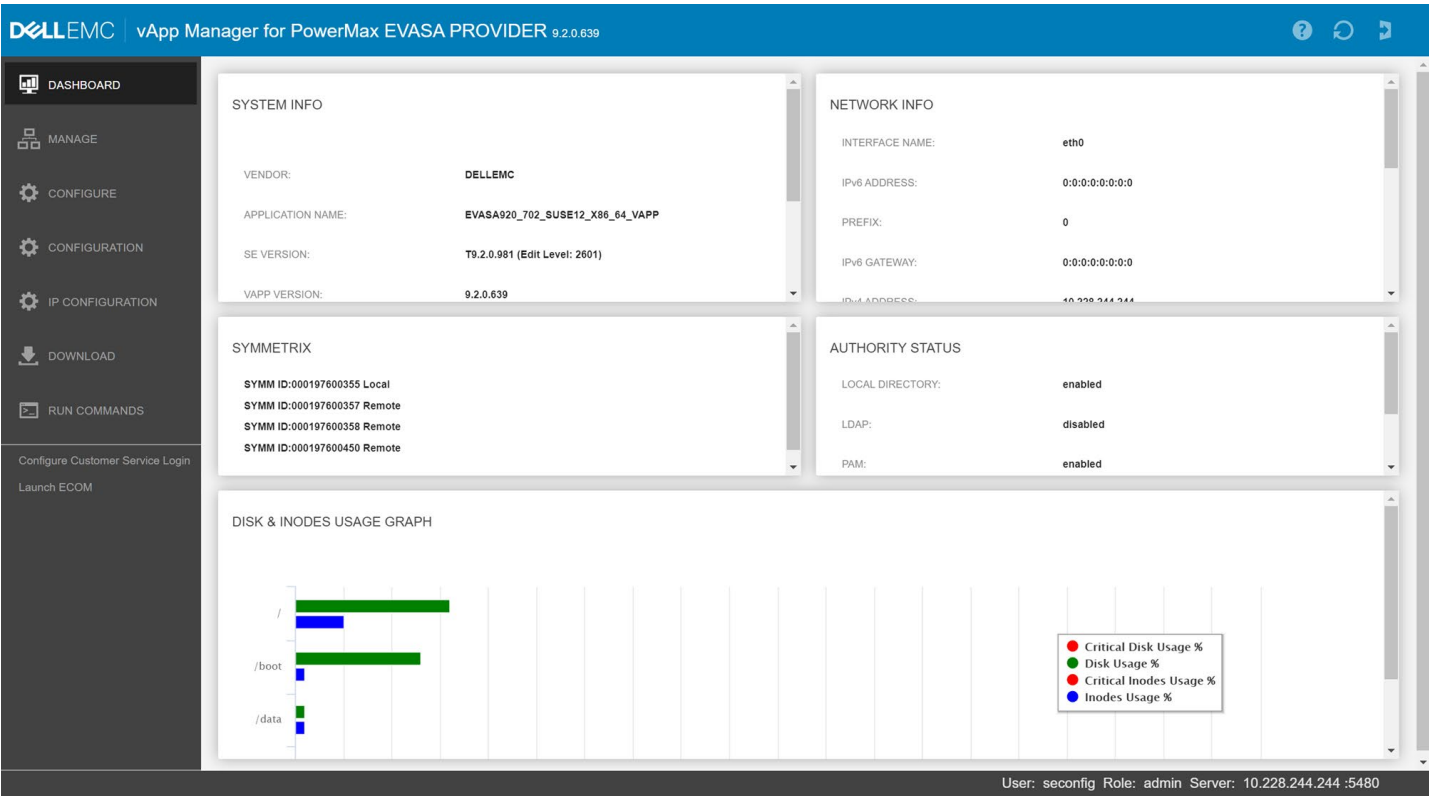


Figure 12. VP vApp dashboard

To view the status of the VASA Provider, navigate to **MANAGE-> DAEMONS** in Figure 13.

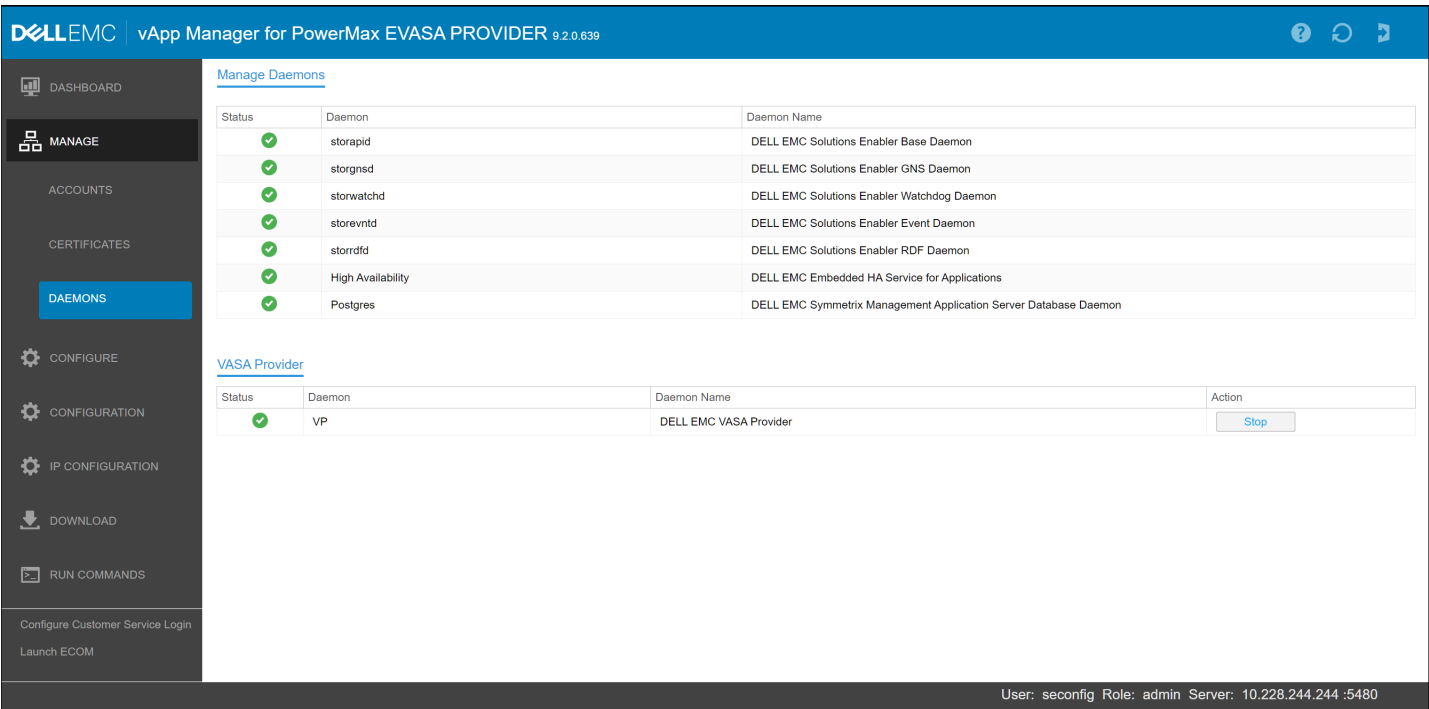


Figure 13. Manage Daemons in VP

4.2.1.1 Logs

If, for some reason, support requires logs from EVASA, they can be downloaded from within the interface. Navigate to the **DOWNLOAD/Daemon Logs** screen and select the *EVASA* Daemon as in Figure 14.

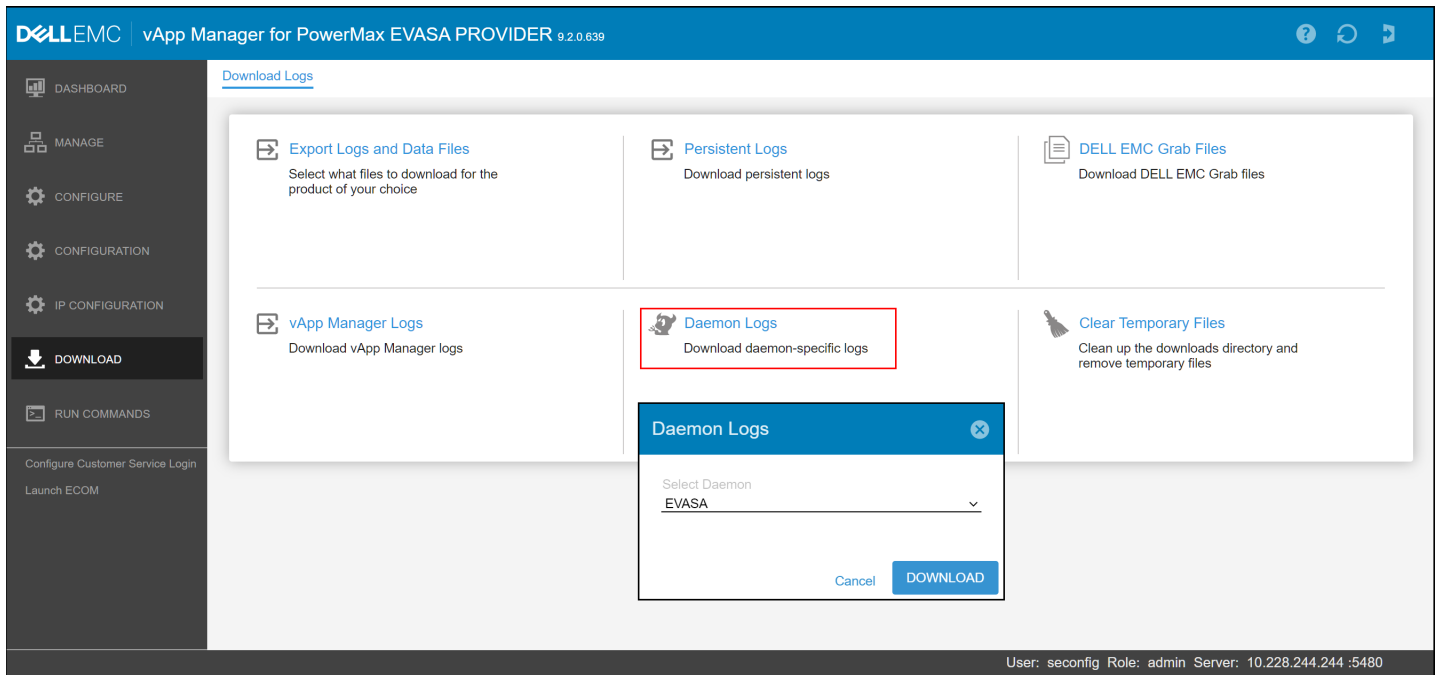


Figure 14. Downloading VASA logs

4.2.1.2 VP parameters

Though the vApp offers some parameters to change the VP configuration in the **VP CONFIGURATION** screen off the **CONFIGURE** menu, seen in Figure 15, these should be left at their default settings unless otherwise instructed by Dell Support, or if using multiple vCenters (see Security Certificates.)

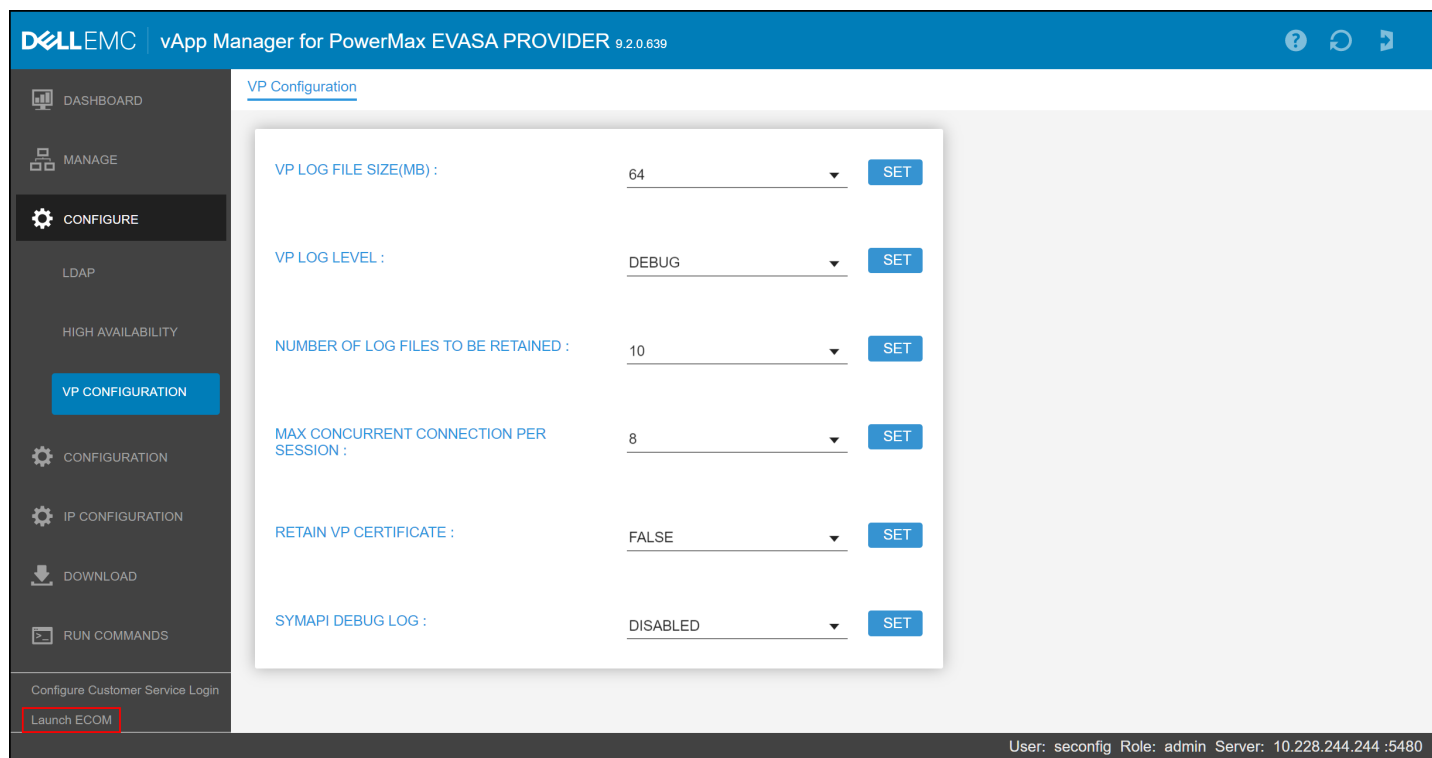


Figure 15. VP Configuration parameters

4.2.2 PowerMax 2500/8500

On the V4 array, management is now embedded in Unisphere in an area called Serviceability. The Serviceability menu is shown in Figure 16. Note that the menu is only available on the embedded Unisphere instance and will not show in an external Unisphere implementation.

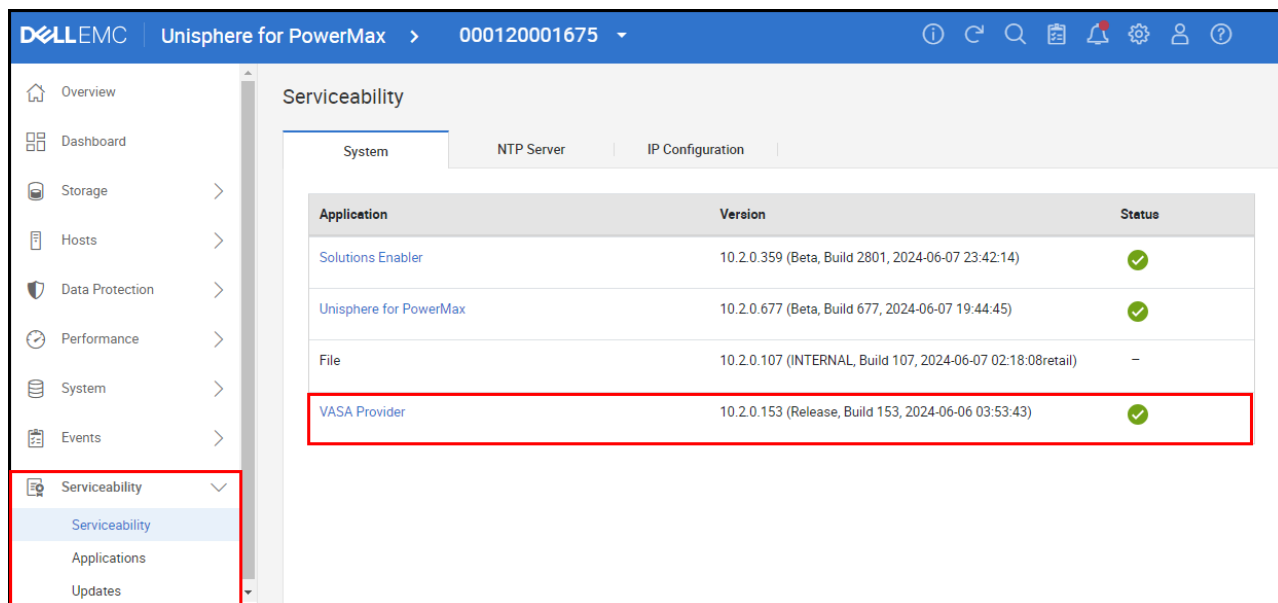


Figure 16. Serviceability in Unisphere for PowerMax

From the menu, select the **VASA Provider** hyperlink. This will change focus to the **Applications** menu and the **VASA** tab. There are three sub-tabs: **System**, **SE Settings**, **VP Settings** and **Virtual Hosts** (only in VASA 5) in Figure 17. From here, the VASA Providers can be stopped and started, the URLs copied for registration in vCenter, and ECOM launched.

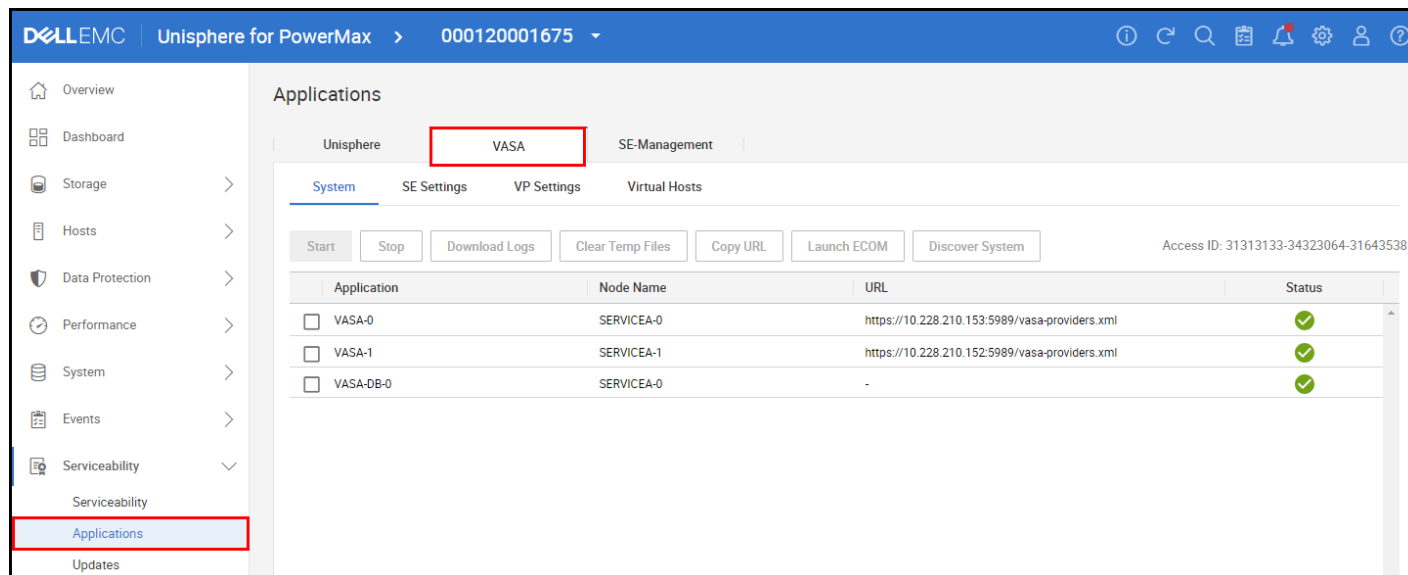


Figure 17. VP dashboard

4.2.2.1 Logs

If, for some reason, support requires logs from the VASA Provider, they can be downloaded from within this interface. Select the active VASA Provider (determined from the vVol Dashboard or the vCenter) and select **Download Logs** as in Figure 18. Note that the pop-up screen does not contain the check box in VASA 3 and 4.

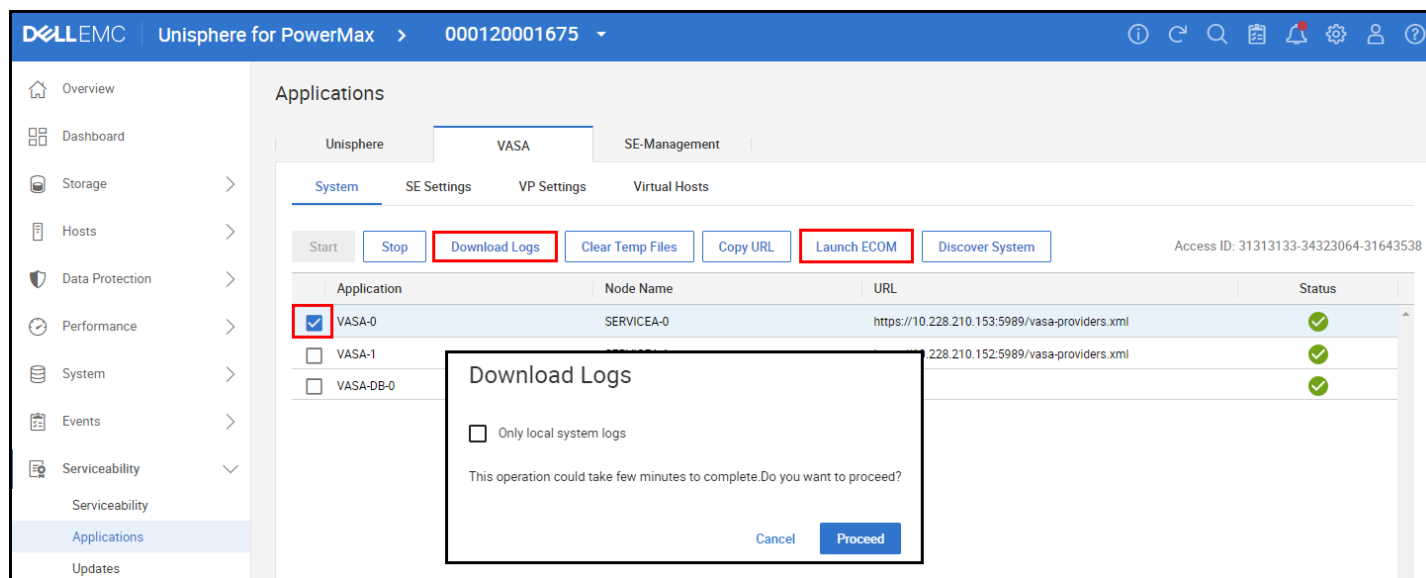


Figure 18. Downloading VASA logs

4.2.2.2 VP parameters

Though there are some parameters available in the VP Settings tab, shown in [Figure 19](#), these should be left at their default settings unless otherwise instructed by Dell Support, or if using multiple vCenters in VASA 3 and 4.

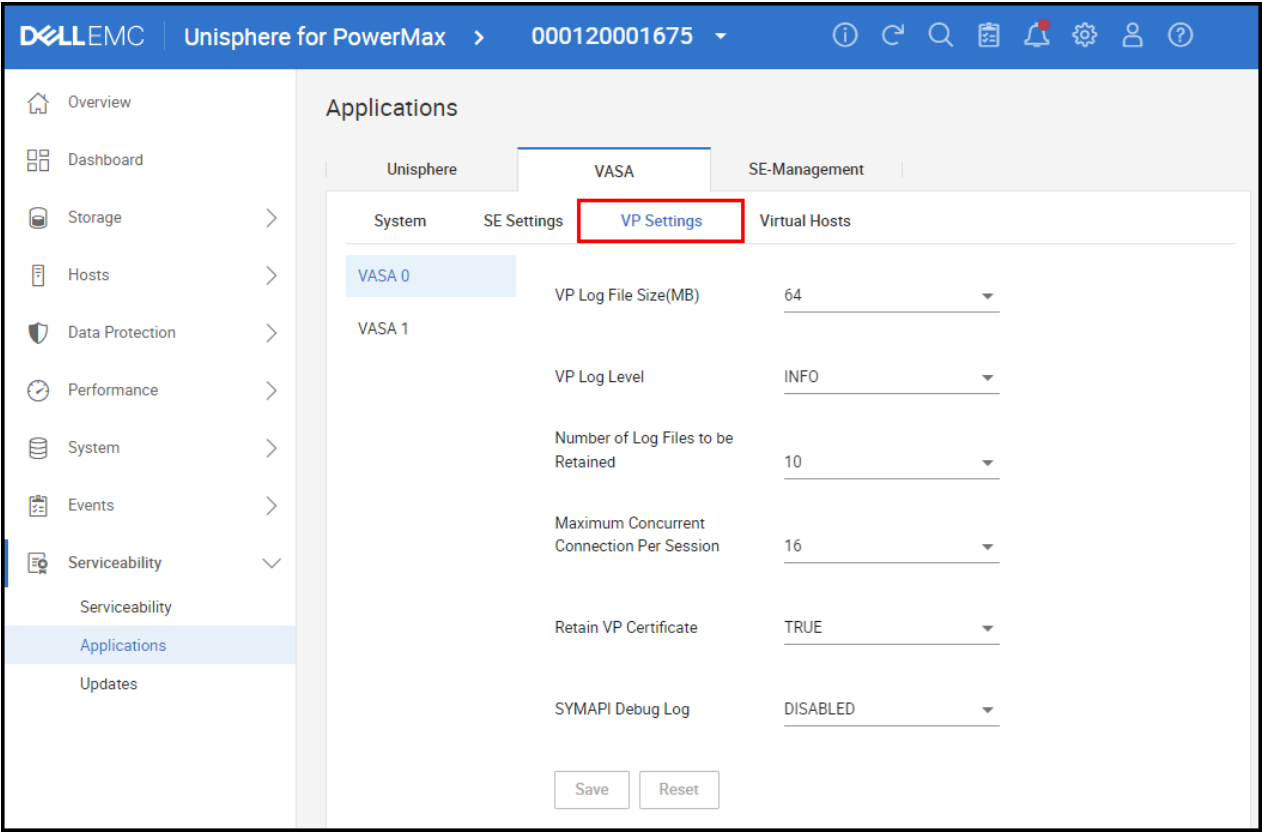


Figure 19. VP Configuration parameters

5 Security Certificates

VASA 3 and 4 use a different certificate mechanism than VASA 5. Before registering the VASA Providers, review the following sections.

5.1 VASA 3 and 4

By default, vSphere includes the VMware Certificate Authority (VMCA). The VMCA creates all internal certificates used in vSphere environment. Communication with the VASA Provider is protected by SSL certificates which are either from the VASA Provider or the VMCA.

When the VASA Provider is registered, VMCA follows these steps automatically:

1. When a VASA Provider is first added to the vCenter Server storage management service (SMS), it produces a self-signed certificate.
2. SMS first verifies the certificate and then requests a Certificate Signing Request (CSR) from the VASA Provider.
3. SMS then validates the CSR and presents it to the VMCA on behalf of the VASA Provider, requesting a CA signed certificate.

Note: The VMCA can be configured to function as a standalone CA, or as a subordinate to an enterprise CA. When the VMCA is configured as a subordinate CA, the VMCA signs the CSR with the full chain.

4. The signed certificate with the root certificates is passed to the VASA Provider. The VASA Provider can then authenticate all future secure connections.

5.2 VASA 5

In VASA 5, the VASA Provider uses a new authentication method which is stricter than VASA 3 and 4. It requires ESXi to be authenticated in context of the vCenter Server. The new model has the following attributes:

- For each vCenter Server that registers with the array, the VASA Provider creates a dedicated virtual host (see Virtual host). The VASA client in vCenter relies on certificate-based authentication and authorization to access its dedicated virtual host created on the array. All certificates get registered with the virtual host. This creates an isolation for each vCenter server.
- Each vCenter provisions a certificate for the VASA Provider, which is managed through the unique virtual host. All ESXi hosts in that vCenter use that virtual host.
- vCenter provisions the VASA client certificate for each ESXi host and synchronizes the public key of the certificate with the VASA Provider. The VASA Provider identifies and authorizes an individual client using the public key.
- vSphere does not trust self-signed certificates for TLS communications. The only exception is during a short period when the VASA provider gets registered and for backward compatibility. An administrator can use a Custom CA certificate for the VASA provider to override the self-signed certificate at the array for purposes of backward compatibility and bootstrapping.

6 Registering the VASA Provider in vCenter

The VASA Provider (VP) is the orchestration entity behind vVols. It enables most functions related to vVols including creation, deletion, powering on/off, etc. The primary function of a VM, IO, however, does not require the VP once the vVol is bound and the VM running. Both VASA Providers must be registered in the vSphere vCenter so VMware can communicate with them. One will be labeled “Active” the other “Standby”. The user does not control the role. The array will adjust the role if required.

The registration differs slightly between VASA 3 and 4 versus VASA 5. For clarity, both are included.

6.1 User authentication for registering the VASA Provider

Before registering the VASA Provider in vCenter a few important configuration changes should be made:

1. Change the default admin password (forced in VASA 5)
2. Create a new user account for VASA (admin privilege). If multiple vCenters will be used in VASA 5, Dell recommends different VASA users for each vCenter.

It is strongly recommended to change the default administrative password and create a separate user account for vCenter access to the VASA Provider.

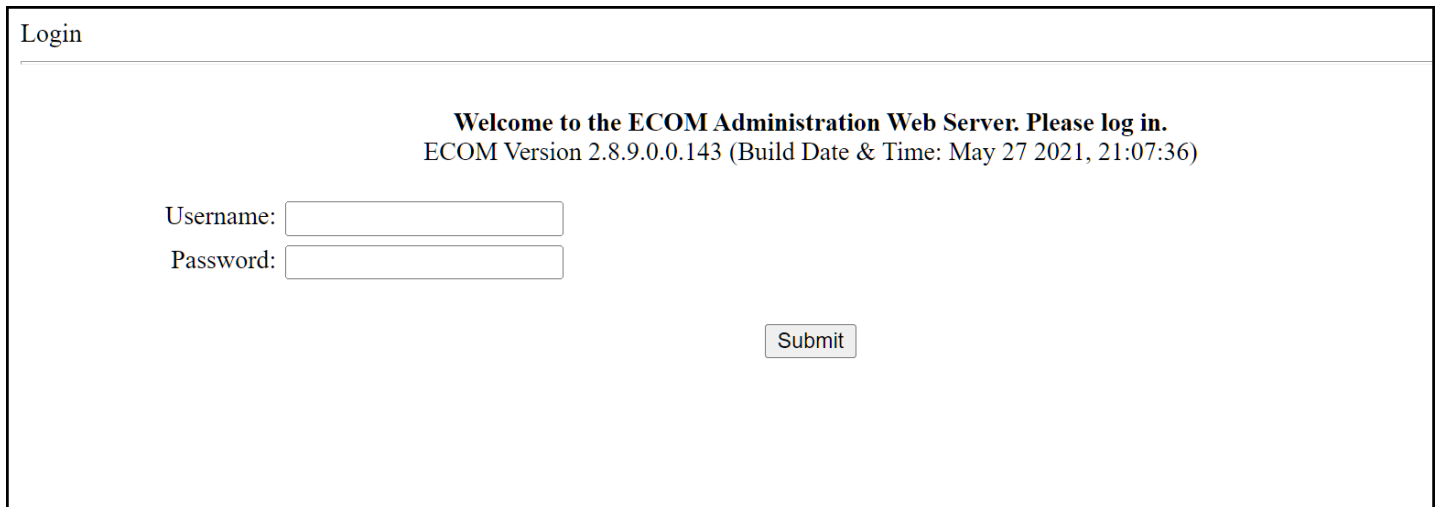
The VASA Provider offers a web-based interface for creating and managing user accounts. The default username and password as well as the direct URL for management access are listed below:

Username: admin

Password: #1Password

Direct management URL: https://<FQDN or IP of the VASA Provider>:5989/ecomconfig

The ECOM page is most easily accessed from the link at the bottom left of the vApp shown in the red box in [Figure 15](#) for the V3, and in [Figure 18](#) for the V4. The logon page for the web-based management interface is shown in [Figure 20](#). Note the interface is the same for all VASA versions. Only the ECOM version is different.



Login

Welcome to the ECOM Administration Web Server. Please log in.
ECOM Version 2.8.9.0.0.143 (Build Date & Time: May 27 2021, 21:07:36)

Username:

Password:

Figure 20. Logon page for ECOM web-based management

The first thing that the user should do is to change the default admin password from #1Password to a unique and complex password to prevent unauthorized access to the VASA Provider. This process is shown in steps in [Figure 21](#) for VASA 3 and 4. Note that all screens are shown together, though when changing the password, the menu on the left disappears once an option is selected. This should be done for both VASA Providers, using a different password for each.

ECOM Administration

ECOM Version 2.8.9.0.0.143 (Build Date & Time: May 27 2021, 21:07:36)
Logged in as admin

Logging:
[Display Log File](#)
[Logging Options](#)

Security:
[Add User](#)
[Modify User](#)
[Change Password](#) 1
[Set Password quality](#)
[Delete User](#)
[List Users](#)
[Display Security Log File](#)
[Client IP Filtering](#)
[Local IP Filtering](#)
[SSL Certificate Management](#)
[LDAP Configuration](#)
[OSLogin Configuration](#)

[Dynamic Settings](#)
[List Running Providers](#)
[Simple CIM Browser](#)

[Logout](#)

Change Password:

User Name:

Current Password:

New Password:

Re-Enter New Password:

2

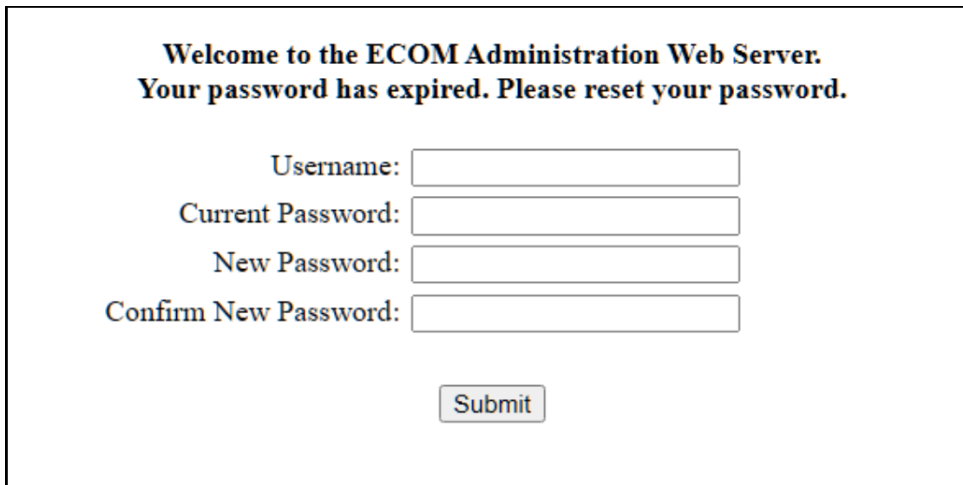
[Back to ECOM Config Page](#)

3

Password changed for user "admin"

Figure 21. Changing the default admin password VASA 3 and 4

For VASA 5, the user will be forced to change the password upon login as in [Figure 22](#).



The screenshot shows a web form titled "Welcome to the ECOM Administration Web Server. Your password has expired. Please reset your password." The form contains four input fields: "Username:", "Current Password:", "New Password:", and "Confirm New Password:". Below these fields is a "Submit" button.

Welcome to the ECOM Administration Web Server.
Your password has expired. Please reset your password.

Username:

Current Password:

New Password:

Confirm New Password:

Figure 22. Changing the default admin password in VASA 5

In addition to changing the admin password, Dell recommends creating a new user dedicated for VASA authentication from vCenter. Administrative access is required for VASA Provider registration with vCenter for vVols. [Figure 23](#) shows the creation of a user account named "vvoluser" with the role type of "administrator".

ECOM Administration

ECOM Version 2.8.9.0.0.143 (Build Date & Time: May 27 2021, 21:07:36)
Logged in as admin

Logging:
[Display Log File](#)
[Logging Options](#)

Security: 1
[Add User](#)
[Modify User](#)
[Change Password](#)
[Set Password quality](#)
[Delete User](#)
[List Users](#)
[Display Security Log File](#)
[Client IP Filtering](#)
[Local IP Filtering](#)
[SSL Certificate Management](#)
[LDAP Configuration](#)
[OSLogin Configuration](#)

[Dynamic Settings](#)
[List Running Providers](#)
[Simple CIM Browser](#)

[Logout](#)

Add User:

User Name:

Password:

Re-Enter Password:

Role: 2

Scope:

Account settings:

☐ User must change password at next logon

☒ User cannot change password

☒ Password never expires

☐ User is disabled

[Back to ECOM Config Page](#)

User "vvoluser" added to Ecom User List 3

Figure 23. Creating a user for the VASA Provider registration

Again, do this for both VASA Providers, using different passwords for the user on each one.

6.2 VP registration wizard in VASA 3 and 4

Start by highlighting the vCenter name in the left-hand panel in step 1. Now in step 2 select the **Configure** tab on the right, and then **Storage Providers** on the left-hand side menu shown in [Figure 24](#). Next select **+ Add** to open the dialog for the VASA Provider.

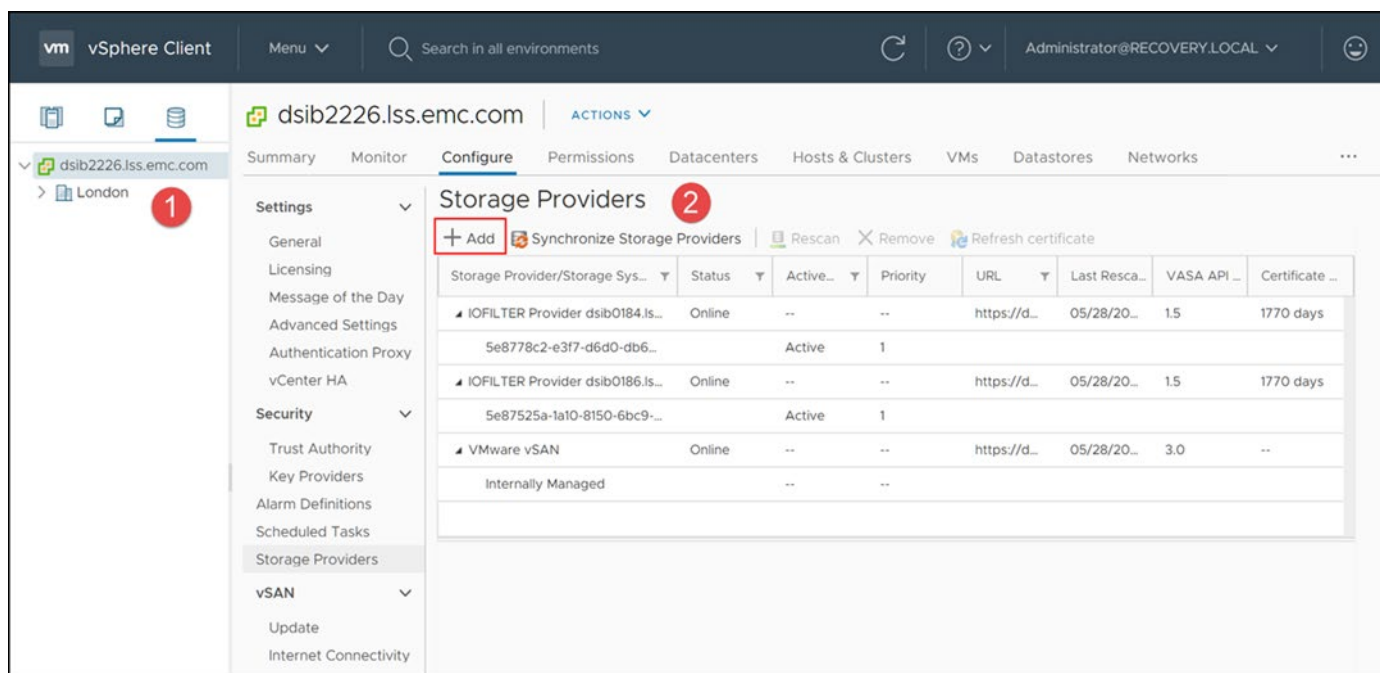


Figure 24. Registering the VASA Provider - step 2

In [Figure 25](#) enter the VP information in the dialog box. The dialog box has 4 fields:

- **Name:** Any descriptive name
- **URL:** This is the VP URL for the deployed appliance. The URL can be easily retrieved from the initial Appliance Info screen seen in [Figure 12](#) or [Figure 17](#). In the Operations portion of the Dashboard of the V3 the URL can be copied directly (step 3) and pasted into the URL field; or on the V4 use the Copy URL option after selecting the VASA Provider.
- **User name:** The user must be one with administrative privilege. In this example the newly created user, “voluser”, is used.
- **Password:** The user password.

Once the fields are filled, select OK.

Note: An additional checkbox is available in the dialog box to register the storage provide which is titled “Use a storage provider certificate”. Dell does not support using this checkbox if the user wants a non-default certificate. To use an existing signed certificate or CA certificate, it must be imported into ECOM. The process is covered in the section [Certificates](#).

Figure 25. Registering the VASA Provider – step 3

In step 4, [Figure 26](#), VMware returns a certificate error indicating the VP host is not trusted. Select Yes to accept the certificate. This error is expected. This error can be avoided by importing the vCenter certificate into ECOM before registration, though it is unnecessary if you trust the host.

Figure 26. Registering the VASA Provider - step 4

Repeat the process for the second VP. Once the registrations succeed, the VP will show that the array is online. It will appear like [Figure 27](#). If they do not show as online immediately, a **Rescan** can be run, though a short wait will resolve things. Note the different roles of the two VPs, one Active, one Standby. The array will control the roles of the VASA Providers. VMware will use the Active VP.

Note: Registration of the VASA Provider is not blocked if the PE or vPE is not presented to the host, so a successful registration should not be used as an indication that the PE or vPE is available.

The screenshot shows the vSphere Client interface for the host 'dsib2226.lss.emc.com'. The 'Configure' tab is selected, and the 'Storage Providers' section is active. The left-hand navigation pane shows 'Settings' and 'Security' sections. The main pane displays a table of storage providers. A red box highlights the 'VASA_O_358' and 'VASA_1_358' providers, which are both 'Online' and have a status of 'Active' or 'Standby'. A red circle with the number '5' is in the top right corner of the main pane.

Storage Provider/Storage Sys...	Status	Active/...	Priority	URL	Last Rescan Ti...	VASA API ...	Cer
5e87525a-1a10-8150-6bc9-...	Active	1					
▲ VASA_O_358	Online	--	--	https://10.228.2...	05/31/2020,...	3.0	338
000197600358 (2/2 online)	Active	255					
▲ VASA_1_358	Online	--	--	https://10.228.2...	--	3.0	338
000197600358 (2/2 online)	Standby	255					
▲ VMware vSAN	Online	--	--	https://dsib222...	05/28/2020,...	3.0	--
Internally Managed	--	--					

The 'Certificate Info' tab is also visible, showing details for the VASA-O certificate, including the subject, issuer, version, kind, valid from, valid to, algorithm, and public key algorithm.

Figure 27. Registering the VASA Provider – completion

6.3 VP registration wizard in VASA 5

With VASA 5, VMCA certificates are used by default. Start by highlighting the vCenter name in the left-hand panel. Now in step 2 select the **Configure** tab on the right, and then **Storage Providers** on the left-hand side menu shown in Figure 24. Next click **ADD** to open the dialog for the VASA Provider.

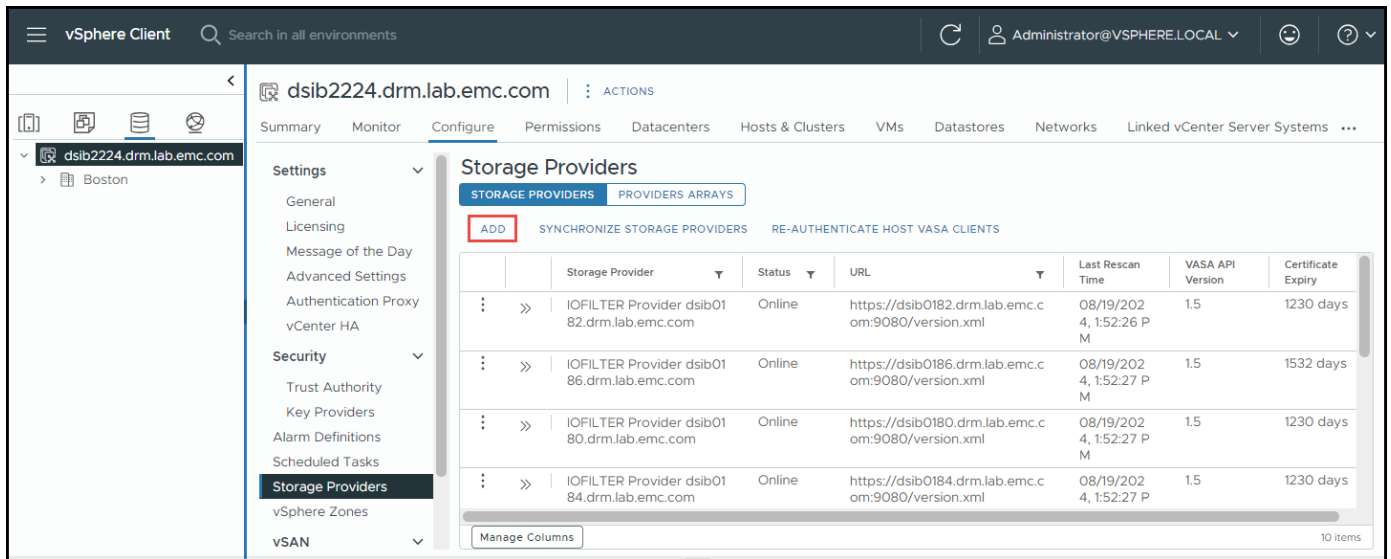


Figure 28. Registering the VASA Provider – Step 1

In Figure 25 enter the VP information in the dialog box. The dialog box has four fields:

- **Name:** Any descriptive name
- **URL:** This is the VP URL for the deployed appliance. The URL can be easily retrieved from the **Applications** screen under the **VASA** tab by checking the box next to either **VASA-0** or **VASA-1** and select **Copy URL**.
- **User name:** The user must be one with administrative privilege. In this example the newly created user, “vvoluser”, is used.
- **Password:** The user password.

Once the fields are filled, select OK.

Note: An additional checkbox is available in the dialog box to register the storage provide which is titled “Use a storage provider certificate”. The process for VASA 5 is covered in the section [Certificates](#).

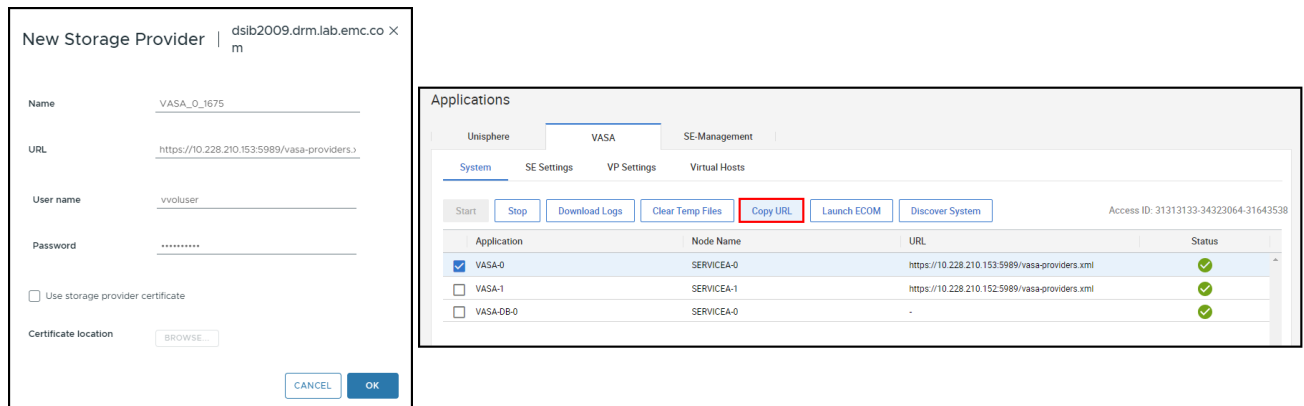


Figure 29. Registering the VASA Provider – Step 2

In step 3, Figure 26, VMware returns a certificate error indicating the VP host is not trusted. Select Yes to accept the certificate. This error is expected.

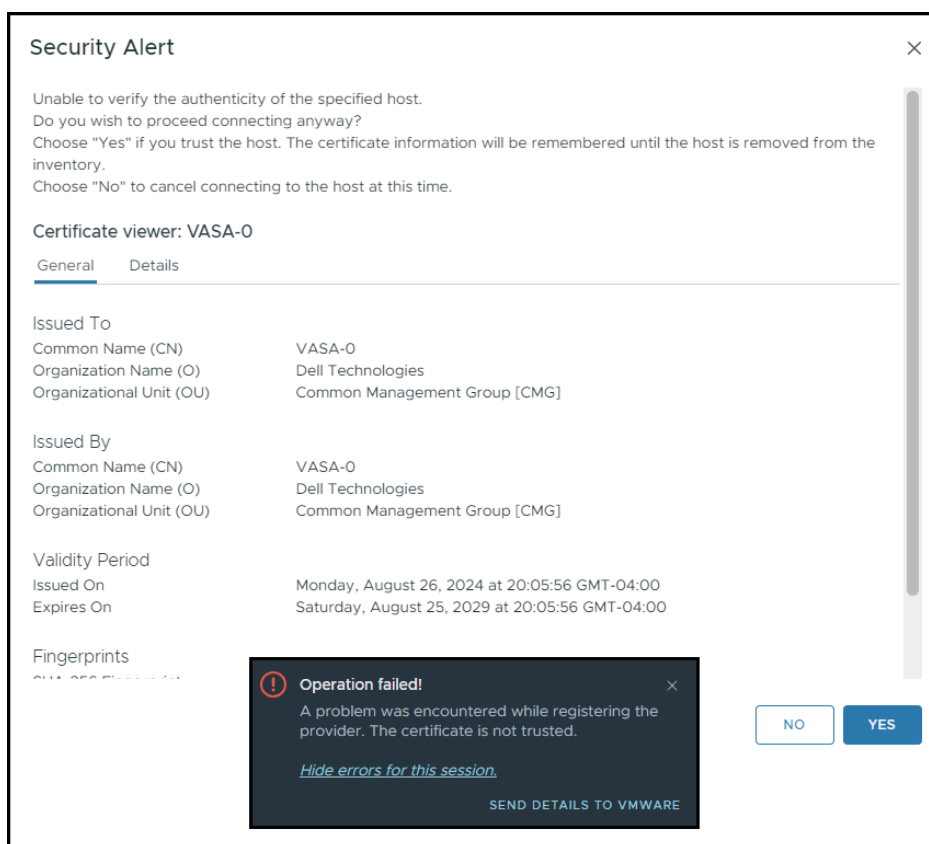


Figure 30. Registering the VASA Provider – Step 3

Repeat the process for the second VP. Once the registrations succeed, the VP will show that the array is online. It will appear like [Figure 27](#). If they do not show as online immediately, a **Rescan** can be run, though a short wait will resolve things. Note the different roles of the two VPs, one Active, one Standby. The array will control the roles of the VASA Providers. VMware will use the Active VP. Note that the active container can be either VASA-0 or VASA-1.

Note: Registration of the VASA Provider is not blocked if the PE or vPE is not presented to the host, so a successful registration should not be used as an indication that the PE or vPE is available.

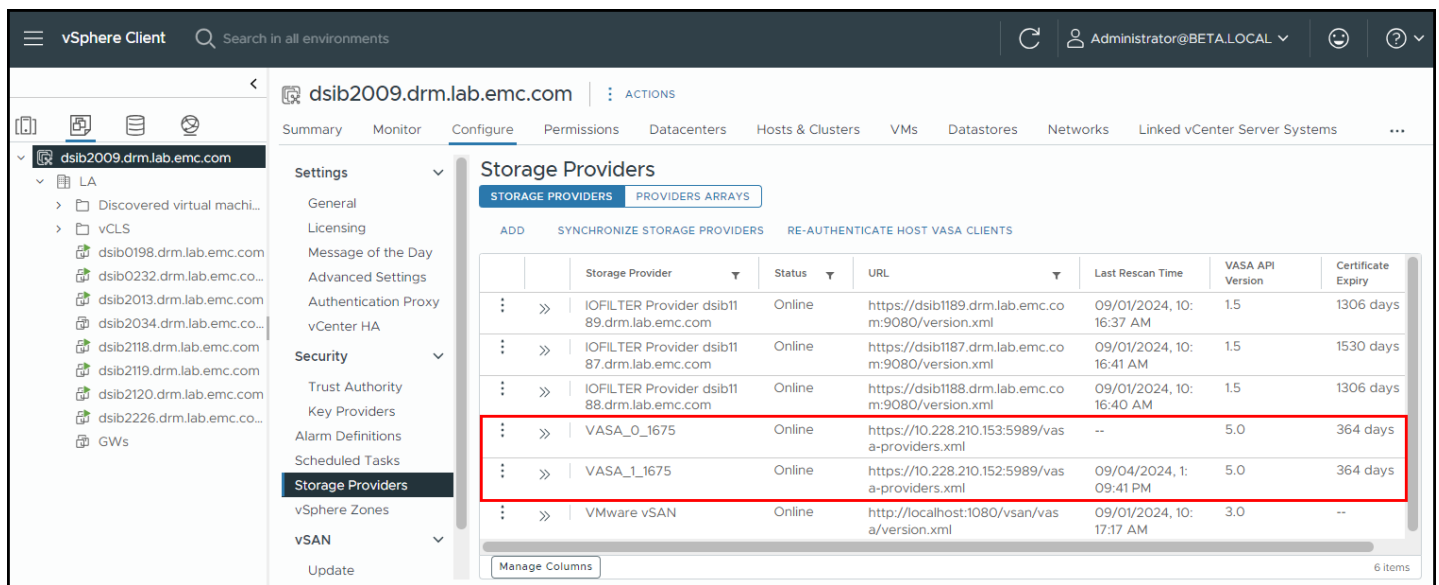


Figure 31. Registering the VASA Provider – completion

6.3.1 Virtual host

Upon registration of the VASA Provider, a virtual host is automatically created on the array. The virtual host's name consists of the vCenter GUID with the suffix **.emc.com**. An example is shown in Figure 32 in Unisphere for PowerMax. The same name is used for both VASA 0 and VASA 1. Note that the **Serviceability** screen is only accessible from embedded Unisphere and not an external instance.

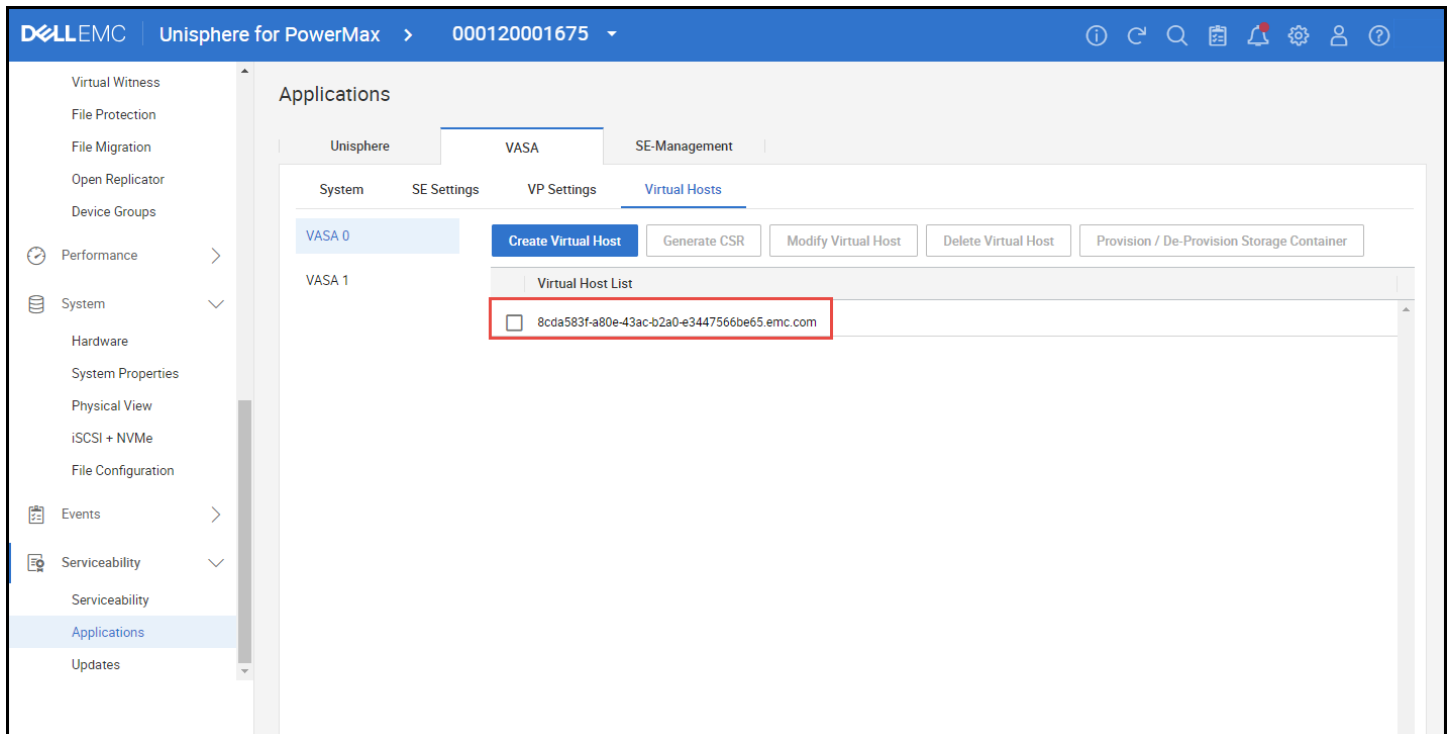


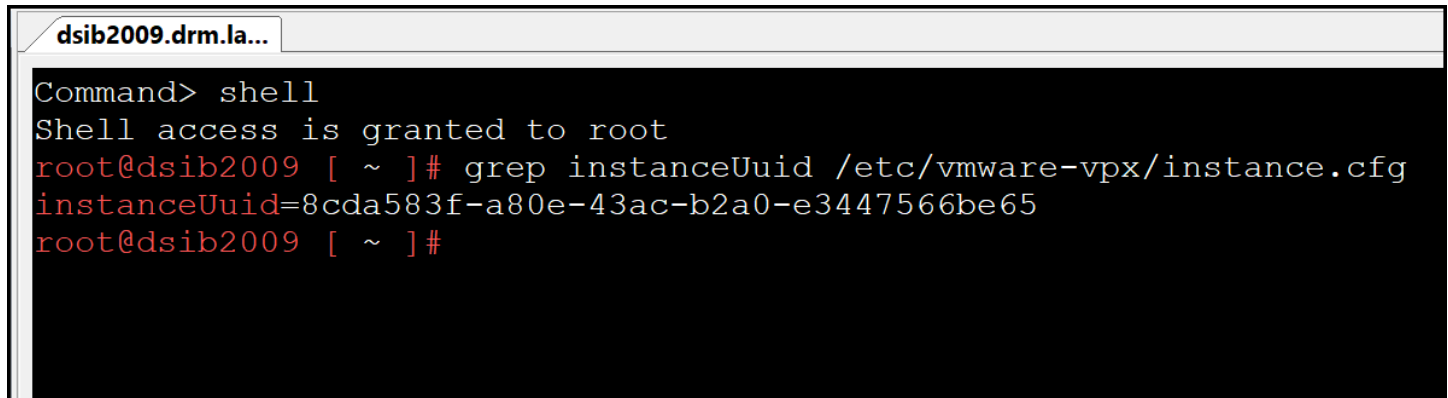
Figure 32. Virtual host

When the VASA Provider creates the virtual host during registration, the only functions in Unisphere available are: **Provision/De-Provision Storage Container** and **Delete Virtual Host**.

6.3.1.1 vCenter GUID

The easiest way to obtain the vCenter GUID is by accessing the vCenter Appliance shell, and running the following command:

```
grep instanceUuid /etc/vmware-vpx/instance.cfg
```



The screenshot shows a terminal window with a title bar that reads "dsib2009.drm.la...". The terminal content is as follows:

```
Command> shell
Shell access is granted to root
root@dsib2009 [ ~ ]# grep instanceUuid /etc/vmware-vpx/instance.cfg
instanceUuid=8cda583f-a80e-43ac-b2a0-e3447566be65
root@dsib2009 [ ~ ]#
```

Figure 33. vCenter GUID

When using the default certificate management of the VASA Providers, it is unnecessary to obtain this GUID. Only when a customer wishes to use their own certificate, is the GUID required. This is covered in the [Certificates in VASA 5](#) section.

6.3.2 Storage container provisioning and de-provisioning

One final aspect of VASA 5 that sets it apart from VASA 3 and 4 is the provisioning and de-provisioning of storage containers to the virtual host. The following section assumes that a storage container is already available. Please see the section [Creating the Storage Container in Unisphere](#) if one does not yet exist and create it prior to continuing. Note that the process of provisioning and de-provisioning is implemented according to the VMware specification in VASA 5.

In VASA 5, after the creation of the virtual host and the registration of the VASA Providers, the vCenter will not see any storage containers on which to create a vVol datastore. The user must first grant the virtual host access to the storage container, as it is blocked by default. This mechanism helps to keep multiple vCenters from seeing the same storage container(s) if that is undesirable. The following section explains how to provision and de-provision storage containers to a particular virtual host.

6.3.2.1 Provisioning

In embedded Unisphere navigate to **Serviceability -> Applications -> VASA -> Virtual Hosts**. Click the checkbox next to the virtual host of the desired vCenter and click the button **Provision / De-Provision Storage Container** in [Figure 34](#).

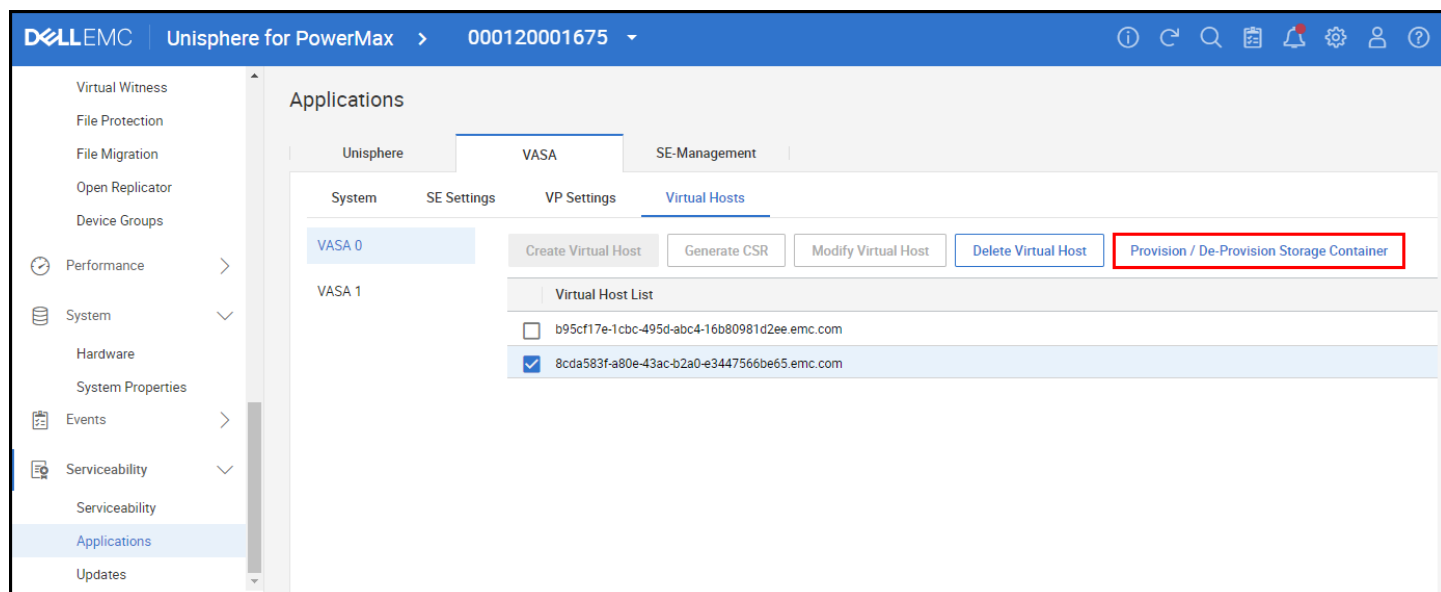


Figure 34. Provision / De-Provision Storage Container – Step 1

In the subsequent dialog box, click the checkbox next to the chosen storage container(s) and use the arrow to push it over to the right. Another dialog box will immediately pop-up indicating success in [Figure 35](#).

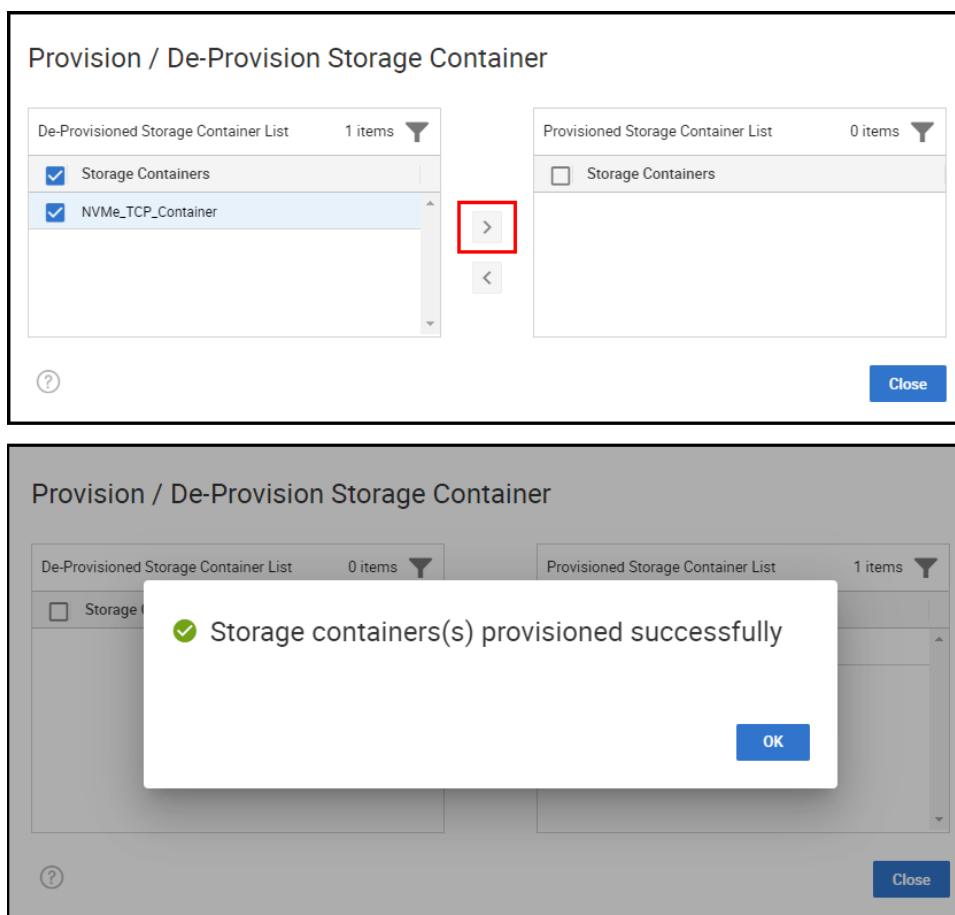


Figure 35. Provision Storage Container – Step 2

It is only possible to provision a storage container to the active VASA Provider. It will automatically give the standby VASA Provider access.

6.3.2.2 De-Provisioning

Storage containers can also be de-provisioned in the same manner. Prior to de-provisioning the storage container, take the following steps:

1. Shutdown or migrate any running VMs in the vVol datastore that is backed by the storage container.
2. Unmount the datastore.

Once complete, use the same dialog to move the storage container from the right-hand side to the left using the arrow as in [Figure 36](#).

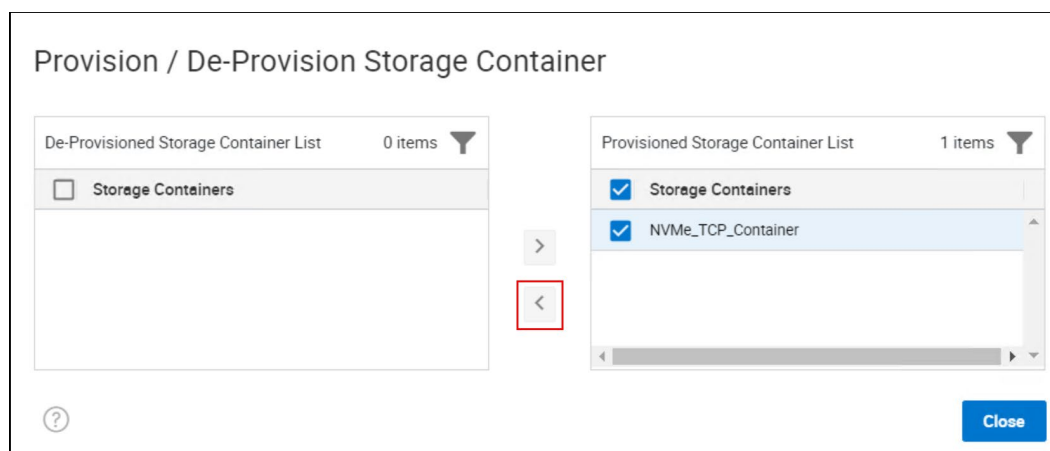


Figure 36. De-Provision Storage Container – Step 1

If the user completed the previously outlined steps, the following dialog box in [Figure 37](#) is presented.

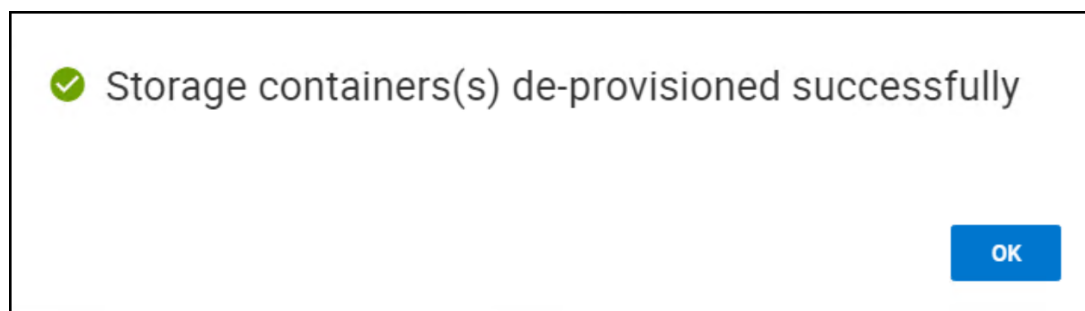


Figure 37. De-Provision Storage Container – Step 2

After de-provisioning, run a rescan of the active VASA Provider from vCenter to complete the task. This ensure that the storage container is unavailable if the user attempts to create a new vVol datastore.

If the user fails to complete the steps above, however, and the storage container remains in use, the warning in [Figure 38](#) is shown indicating there are vVol bindings.

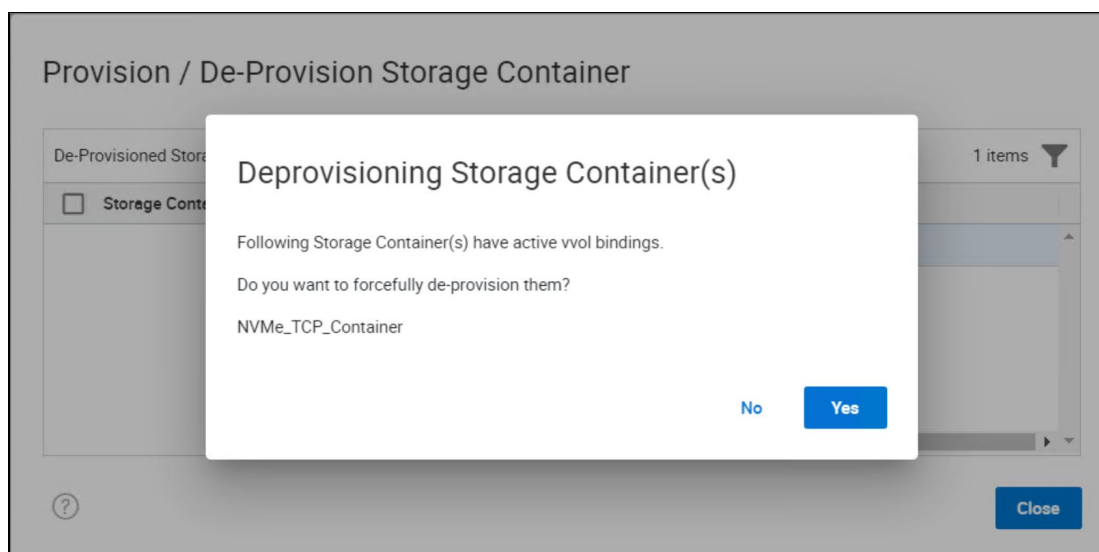


Figure 38. De-Provision Storage Container – Step 2

It is preferable to select **No** in this dialog box and complete the previous two steps before re-attempting. If, however, the user forces the de-provisioning, the task will complete. In such an instance, the vVol datastore will remain available and fully functional until the user completes the steps and rescans the VASA Provider.

Note: It is critical to understand that unless the vVol datastore is unmounted and the VASA Provider rescanned, the vCenter will retain access to the storage container even if it is de-provisioned. De-provisioning has no capability to forcibly remove access to a mounted vVol datastore. Furthermore, failure to rescan the VASA Provider after de-provisioning will allow the vCenter to re-use the storage container. Rescanning is not an automated process in vCenter as of publication of this paper.

7 Certificates in VASA 3 and 4

The following sections cover using customer certificates and the changes required when using two vCenters with the same VASA 3 or 4 Providers.

7.1 Signed certificates

If the customer needs to use their own signed certificate or CA certificate, it must be imported into ECOM because, as noted, there is no support for the checkbox in the register storage provider dialog box. To do this, first login to ECOM as shown in [Figure 20](#). Next, navigate to SSL Certificate Management in [Figure 39](#).



Figure 39. SSL Certificate Management menu

In the next screen in [Figure 40](#), there are three options available. Option #1 offers the ability to generate a signed certificate and then import it. Option #3 offers the ability to directly import a CA certificate.

SSL Certificate Management

SSL Certificate Management

Option #1: Signed Certificate
 Sign a ECOM generated Certificate Signing Request (CSR).

1. Generate a Certificate Signing Request.
2. Export the PKCS#10 formatted Certificate Signing Request.
3. Get Certificate Signing Request signed by your Certificate Authority (CA).
4. Import the signed certificate as PEM data.

 See Wikipedia for more information about [Certificate Signing Requests](#).

Generate a Certificate Signing Request

Import Signed Certificate

Option #2: Self Signed Certificate
 Select this option to auto generate a self-signed certificate. A newly generated self-signed certificate will replace the currently used one. The certificate's Subject information will contain the defaults.

Generate a Self-Signed Certificate

Option #3: Import CA certificate file
 Select this option to import the CA certificate file (as PEM data).

Import CA certificate file

Back

Figure 40. SSL Certificate Management Options

After importing the appropriate certificate, the ECOM process must be stopped and started. If using the vApp, stop and start VP in [Figure 13](#). If using the embedded VASA Provider, stop and start VASA-0 or VASA-1 as seen in [Figure 17](#). Be sure to do the import and stop/start for each VASA.

7.2 Default certificate

If it is necessary to return to a default certificate, follow these steps:

1. If any VASA Providers are registered in any vCenters, remove all of them. Restart the VP daemon/VASA Provider seen in [Figure 13](#) or [Figure 17](#) for both active and standby VASA Providers.
2. Log in to the ECOM Config page as an admin user in [Figure 20](#).
3. Go to SSL Certificate Management in [Figure 39](#) and then select Option #2 Generate Self-Signed Certificate in [Figure 40](#).
4. The Self-Signed Certificate page is displayed as shown in [Figure 41](#).

Self-Signed Certificate

See Wikipedia for more information about [X.509 Subject fields](#).
 See Wikipedia for more information about [SAN - Subject Alternative Name](#).

Common Name (additional semicolon separated list of SAN DNS names):

Country Name (2 letter code):

State or Province Name:

Locality Name:

Organization Name:

Organizational Unit Name:

Serial Number:

SAN Email Address:

SAN IP (semicolon separated list)

SAN URI (semicolon separated list)

Key Usage (semicolon separated list) (If empty, the value from XML setting 'ExtKeyUsage' will be used.)

CA (Certificate Authority) flag ☐

Figure 41. Self-Signed Certificate generation

Enter the following information into the fields:

- Common Name: <hostname of VASA Provider>, e.g., VASA-0
 - Country: <Country>
 - State: <state>
 - Locality: <locality>
 - Organization name: <Org name>
 - Organization Unit name: <Org unit name>
 - Serial Number: <must not be 0 and be unique>, e.g., array SID
 - SAN Email Address: <empty>
 - SAN IP: <Vasa Provider IP address>
 - SAN URI: <empty>
 - Key Usage: <default>
 - CA: <leave this unchecked>
5. Click Generate a Self-Signed Certificate.
 6. Restart the VP daemon/VASA Provider as previously covered in step 1.
 7. Register the VASA Provider in each vCenter. If already registered, unregister and register again.
 8. Repeat for the standby VASA Provider if using the embedded version.

7.3 Multiple vCenters and certificates

If the VASA Provider will be registered in two vCenters (max supported)⁴, a couple changes are required. These changes need to be made before registering the VPs in the vCenters, otherwise connectivity issues will ensue due to the certificate. The steps are outlined below and should be followed in order.

Note: Stopping and starting the VASA Providers can take a few minutes. Do not interrupt the process.

Note: VASA 5 does not require any changes to register multiple vCenters.

7.3.1 Unregister the VASA Providers

If either VASA Provider from the array is registered in either vCenter, unregister it. There can be no registered VASA Providers from the chosen array before beginning this process.

Once the VASA Providers are unregistered, stop each VASA Provider (e.g., VASA-0 and VASA-1). When both are stopped, restart them, always beginning with the first one (e.g., VASA-0).

7.3.2 Retain VP Certificate

The first step is to change a flag in the VASA Provider configuration seen in [Figure 15](#) or [Figure 19](#). In this screen use the drop-down box next to “RETAIN VP CERTIFICATE” to change it from FALSE to TRUE and select “SET” or “Save” depending on the interface. The user may be alerted that the ECOM service must be restarted. Now repeat the process of stopping and starting the VASA Providers as explained in the previous section.

7.3.3 Self-signed certificate

When the Retain VP Certificate parameter is set to TRUE, it is not possible to use the default self-signed ECOM certificate. This is because VMware requires that the key *CertificateSign* has a value in the certificate. The ECOM one does not. Therefore, a new certificate must be generated for each VASA Provider. Follow these steps to generate the new certificate on each VASA Provider, beginning with VASA-0.

1. Log in to the ECOM Config page as an admin user in [Figure 20](#).
 2. Go to SSL Certificate Management in [Figure 39](#) and then select Option #2 Generate Self-Signed Certificate in [Figure 40](#).
 3. The Self-Signed Certificate page is displayed as shown in [Figure 41](#). Enter the following details. It is critical that the Serial Number, SAN IP, and Key Usage fields are correctly inputted. In particular, note that the Key Usage field is different than what is the default setting.
 - Common Name: <hostname of VASA Provider>, e.g., VASA-0 or VASA-1
 - Country: <Country>
 - State: <state>
-

⁴ If more than two vCenters are desired, an RPQ is required.

- Locality: <locality>
- Organization name: <Org name>
- Organization Unit name: <Org unit name>
- Serial Number: <must not be 0>, e.g., use array SID for both VASA-0 and VASA-1
- SAN Email Address: <empty>
- SAN IP: <VASA Provider IP address from either VASA-0 or VASA-1>
- SAN URI: <empty>
- Key Usage:
DigitalSignature;NonRepudiation;KeyEncipherment;KeyAgreement;EncipherOnly;DecipherOnly;CertificateSign
- CA: <leave this unchecked>

An example is show in [Figure 42](#).

See Wikipedia for more information about [X.509 Subject fields](#).
 See Wikipedia for more information about [SAN - Subject Alternative Name](#).

Common Name (*additional semicolon separated list of SAN DNS names*):

Country Name (2 letter code):

State or Province Name:

Locality Name:

Organization Name:

Organizational Unit Name:

Serial Number:

SAN Email Address:

SAN IP (*semicolon separated list*)

SAN URI (*semicolon separated list*)

Key Usage (*semicolon separated list*)

CA (Certificate Authority) flag ☐

[Figure 42. Certificate example](#)

4. Click Generate a Self-Signed Certificate.
5. Repeat steps 1-4 for VASA-1. Be sure the Common Name is VASA-1 and that the SAN IP is the IP of VASA-1 and not VASA-0.
6. Stop and start both VASA Providers as previously covered.

7. Register VASA-0 and VASA-1 VASA Providers in the first vCenter. VMware will still generate the certificate error which should be acknowledged. When finished, check which VASA Provider is shown as Active and which is Standby (active should be VASA-0 but check in case it is VASA-1). Whichever VASA Provider is Active, register that one first in the second vCenter. Then register the other VASA Provider.

Note: While it is possible to mount the same vVol container to both registered vCenters, it is not recommended for production use, just as VMware does not recommend mounting the same VMFS datastore to multiple vCenters.

A couple rules govern multiple vCenters which are:

- The user must register the current active VASA Provider first in the second vCenter. If the standby VASA Provider is registered first, then VP will fail the request with ActivateFailed error. The current active VASA Provider can be obtained from the first vCenter.
- If either of the vCenters lose communication with the current active VASA Provider, then the array activates the standby VASA Provider. In such scenarios, the other vCenter is also switched over to the standby VASA Provider (now the active one).

7.3.3.1 Process

When a VASA Provider is registered with a vCenter, the following actions take place:

1. The VASA Provider presents the vCenter with a self-signed certificate. The vSphere user validates the thumbprint and accepts the certificate.
2. After the validation, VASA Provider sends a CSR with its public key.
3. vCenter responds with a VMCA signed certificate that is stored in the VASA Provider trust store.

Note: The above steps are only applicable when using a VMCA as certificate authority. The VASA Provider uses VMCA signed certificates to communicate with ESXi and vCenter servers. This works because all the components (vCenter/ESXi/VP) use VMCA signed certificates.

In a multi-vCenter environment, when a VP is registered with the first vCenter the steps mentioned above are performed and VP receives a VMCA certificate from the first vCenter. When the same VASA Provider is registered again with second vCenter, and if the RETAIN_VP_CERTIFICATE flag is set to FALSE (not enabled) then the certificate is overwritten and the communication with the first vCenter is broken and only second vCenter communicates with VASA Provider.

If the RETAIN VP CERTIFICATE flag is set to TRUE, this makes sure that the VP trust store is not overwritten, hence the communication with first vCenter is not broken, but communication with second vCenter will not be established.

If user wants to use both vCenters, there are two options:

1. A common VMCA is required, where both vCenters are part of a common VMCA.
2. Use Self Signed Certificate along with the extension CertificateSign to be present in the keyUsage (In this case the RETAIN VP CERTIFICATE flag must be set to true).

The VMCA can be deployed as a standalone or as a subordinate to another CA. If both vCenters use a standalone VMCA, the configuration is not supported because they are two different CA instances, and the VP can retain only one.

7.3.4 CA certificate

To use the CA certificate in the VP environment, the `RETAIN_VP_CERTIFICATE` flag should be unset (marked `FALSE`). In this scenario, at VP registration time, the vCenter ignores the certificate that is sent by the VP and sends the vCenter certificate as the new certificate. This certificate is stored by the VP for all future handshakes.

8 Certificates in VASA 5

The following sections cover how a customer can pre-create the virtual host and use their own third-party CA signed certificate, rather than the default VMCA in VASA 5.

8.1 Signed certificates

In VASA 5, a customer must pre-create the virtual host in Unisphere in order to use third-party certificates. Before beginning, obtain the vCenter GUID by using the process explained in section [vCenter GUID](#). Then in Unisphere, navigate to **Serviceability -> Applications -> VASA -> Virtual Hosts** in [Figure 43](#).

Note: This procedure must be followed for both VASA 0 and VASA 1.

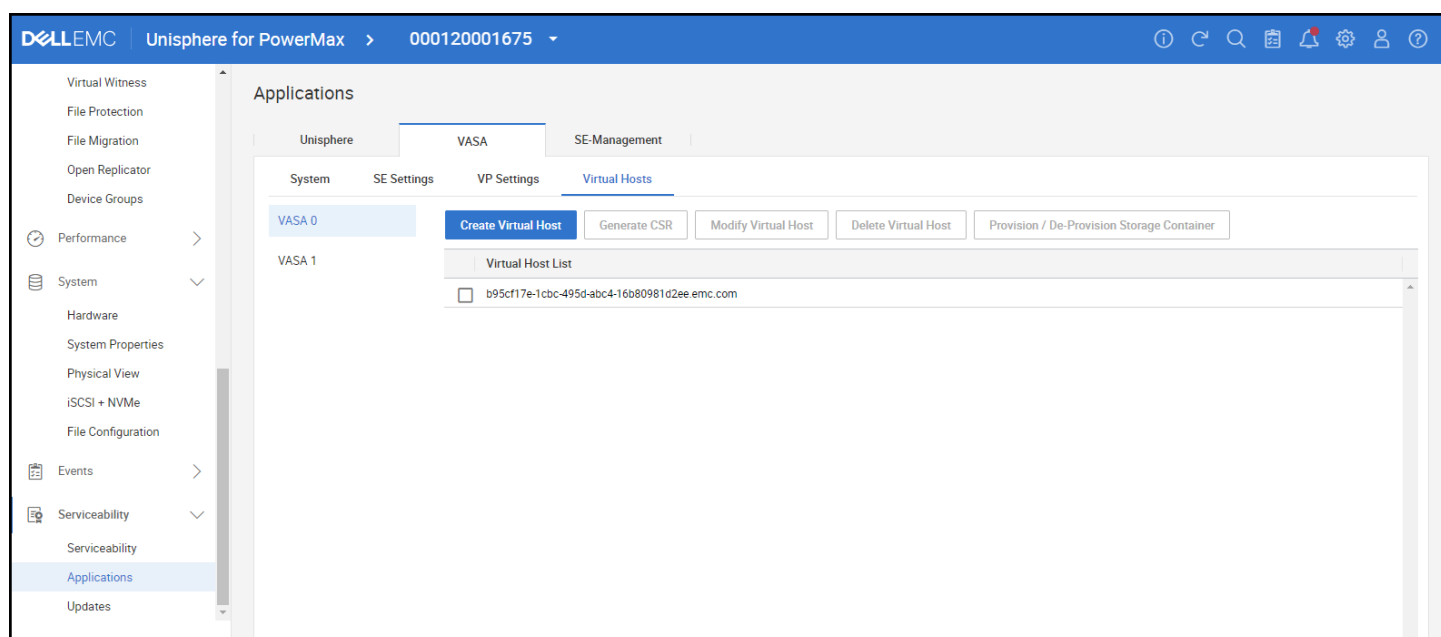


Figure 43. Pre-create a virtual host – Step 1

Next, click on **Create Virtual Host** and type or paste the vCenter GUID. Click **Create** in [Figure 44](#).

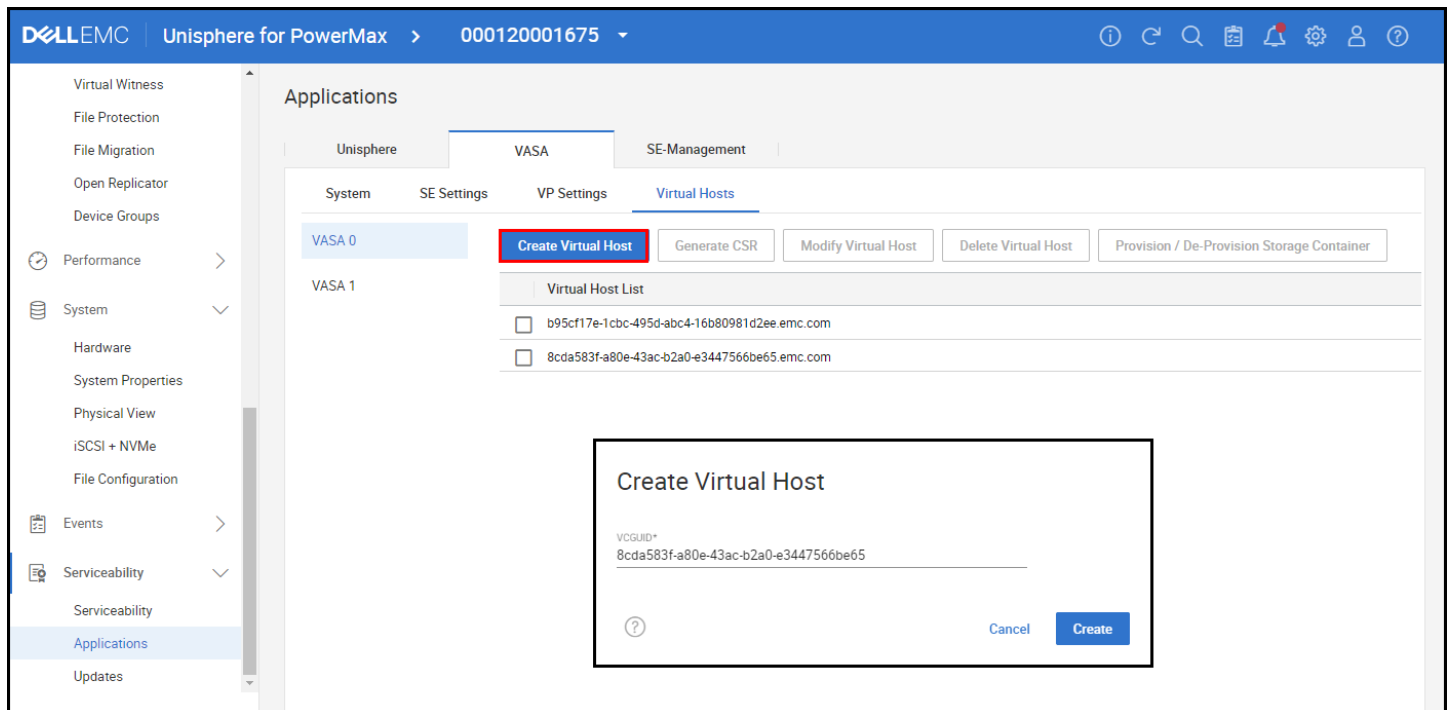


Figure 44. Pre-create a virtual host – Step 2

The user must now update the virtual host with third-party signed root CA certificates. Begin by generating a CSR certificate. In the same screen in Unisphere, click the checkbox next to the newly added virtual host and click the **Generate CSR** button. In the pop-up dialog click on **Download** and save the .pem file locally. This is shown in Figure 45.

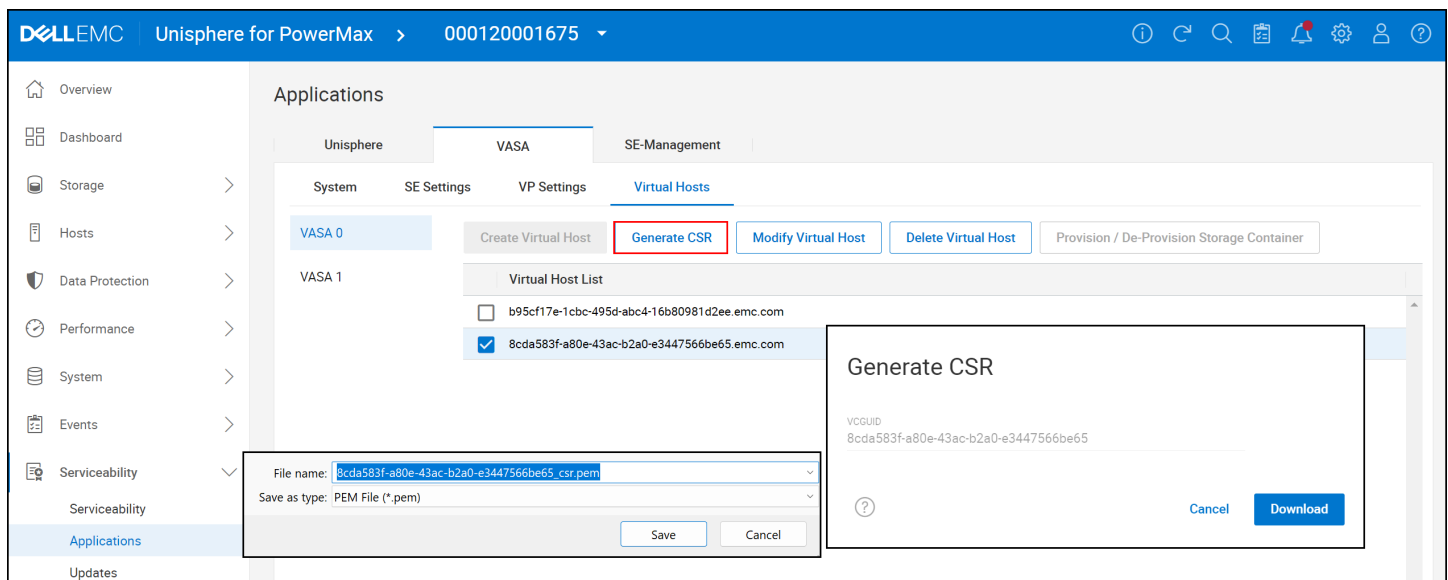


Figure 45. Pre-create a virtual host – Step 3

Use the CSR file to sign the third-party CA root certificate. Once complete, the user must modify the virtual host. Return to the same Unisphere screen, check the box next to the virtual host and click

on **Modify Virtual Host**. The user will be prompted to select the newly signed certificate in [Figure 46](#). Repeat all steps for VASA 1.

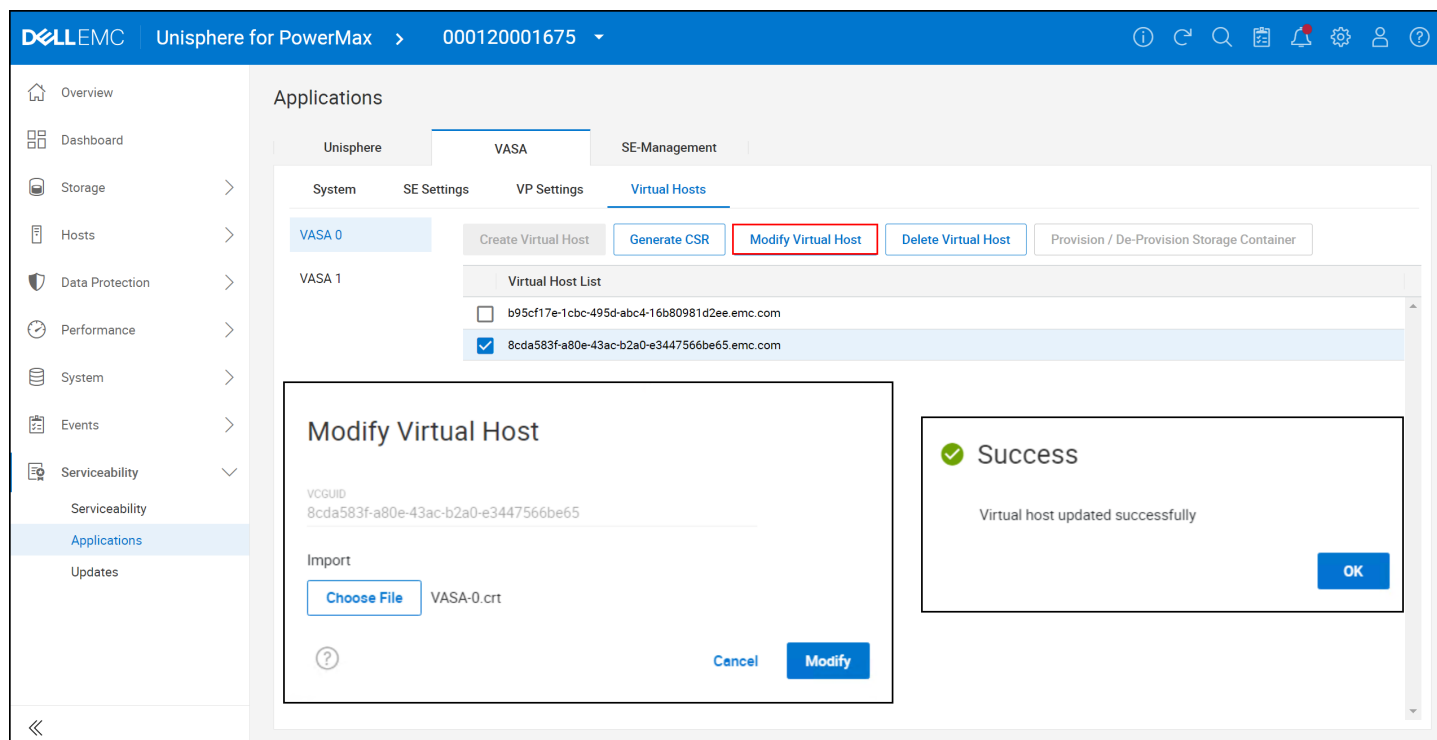


Figure 46. Pre-create a virtual host – Step 4

8.1.1 VC trust store

Before registering the VASA Provider, the root CA certificates, one at a time, should be added to the VC trust store. In vCenter, navigate to **Administration -> Certificates -> Certificate Management**. Click on **ADD TRUSTED ROOT CERTIFICATE** and supply the certificate. Click **ADD** as in [Figure 47](#).

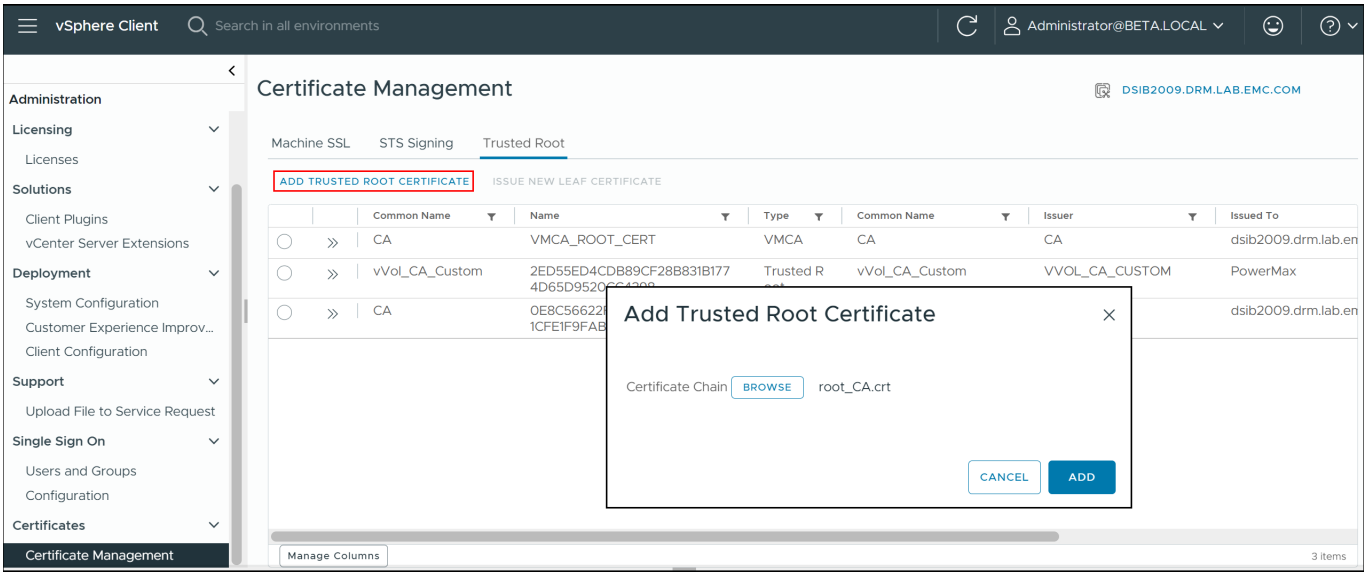


Figure 47. Add root CA to vSphere

Now proceed to register the VASA Providers.

8.2 Multiple vCenters

The certificate changes and use of the virtual host in VASA 5 mean that the user is no longer required to make any modifications to the VASA settings on the PowerMax to enable multiple vCenters. The virtual host functionality takes the place of sharing certificates as in VASA 3 and 4 as explained in the section [Multiple vCenters and certificates](#). When registering additional vCenters, start with the active VASA Provider then the standby.

VASA 5 supports a maximum of five (5) vCenters per PowerMax array.

9 Configuring Virtual Volumes on the PowerMax

The following sections will detail how the SAs can manage their vVol tasks through the GUI and the CLI. As there are only cosmetic differences between Unisphere for PowerMax versions 9.x and 10.x, different sections will use different versions.

Note: Dell does not support Virtual Volumes on external storage attached to a PowerMax.

9.1 Using Unisphere with Virtual Volumes

vVols has been integrated into Unisphere for PowerMax through a dashboard: vVol Dashboard. This dashboard is the central location for managing vVols in a PowerMax environment. From here, the storage administrator can create storage containers with the required storage resources, provision protocol endpoints to the ESXi hosts, and configure VASA Replication Groups. The vVols dashboard appears in [Figure 48](#).

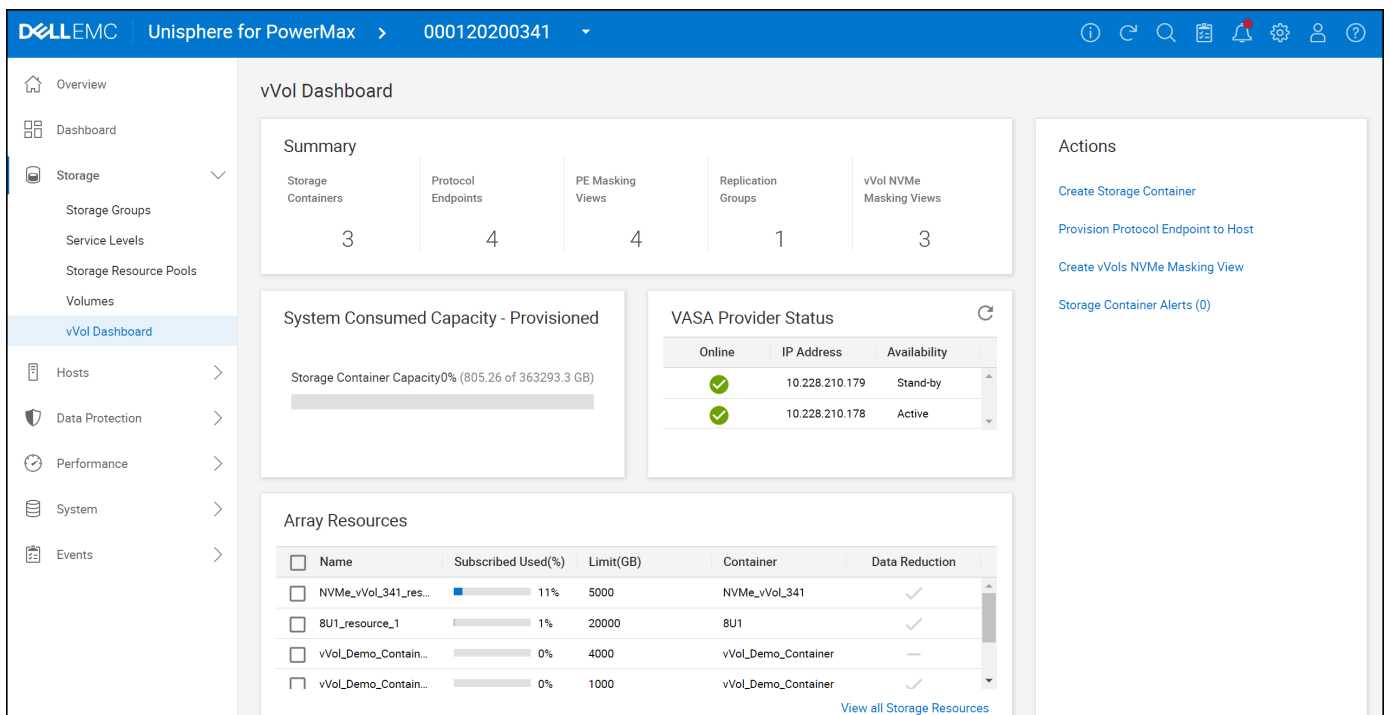


Figure 48. Unisphere for PowerMax vVol dashboard

9.1.1 VASA Provider Status

Unisphere will automatically recognize if an array has the embedded VASA and list both instances in the dashboard in the widget "VASA Provider Status". Note in [Figure 48](#) that in addition to the general status of the VASA Provider the roles are included, Active and Standby. The roles are also available in the vCenter once the VASA Provider is registered.

9.2 Creating the Storage Container in Unisphere

As previously explained, a storage container (SC) is a logical construct on the array that partitions space based on Service Level (SL). The SA creates storage resources which at a high level are a combination of a service level and a storage size. A storage container may only have a single

storage resource for each SL. An SC, for instance, may not have two storage resources with an SL of Optimized. By default, data reduction is enabled on newly created storage resources.

Note: Dell supports sixteen (16) total storage containers on a PowerMax array. Multiple containers may be desired, for example, to separate test and dev environments from production, or limit storage for a particular business unit. Multiple containers do not change performance, however. Only the service level of a storage resource impacts performance.

The following walks through the SC wizard in individual steps.

9.2.1 SC creation wizard

From the vVol dashboard, [Figure 49](#), access the **CREATE STORAGE CONTAINER** option in the Actions menu on the right-hand side.

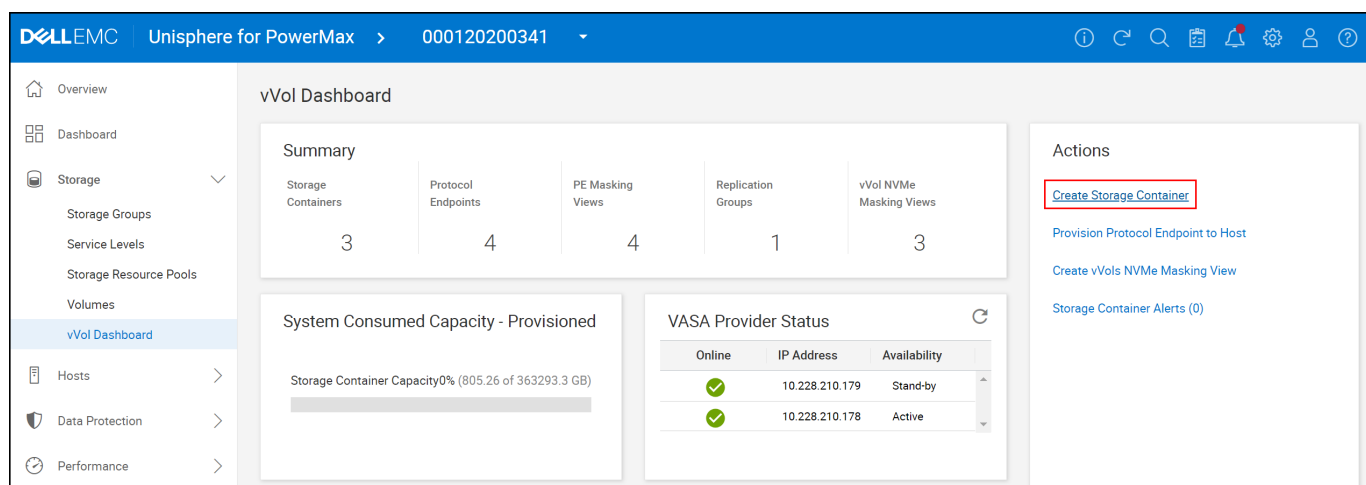


Figure 49. Create Storage Container – step 1

In the second step the user must choose the type of storage container based on the protocol being used. For iSCSI or FC select **SCSI**, for NVMe/TCP select **NVMe**. Note in particular the warning when selecting NVMe, that the vPE (via masking view) must be presented to the host before creating the VMware vVol datastore. Even though there is no such warning for SCSI, the same principle applies.

Enter a name for the SC and, if desired, a description as in [Figure 50](#). If both SCSI and NVMe/TCP are used on the same ESXi host, it is beneficial to include the protocol in the name. The description is only available to the SA. The VMware admin will not see it so it should not be used in the hopes it can provide important information within vSphere. Select **Next**. Note that SCSI will be used for the rest of the wizard.

Create Storage Container

Storage Container
Storage Resources
Summary

Storage Container

Storage Container Name *

NVMe_vVol_Container

Description

Container for NVMe protocol

☐ SCSI
☒ NVMe

Important

Please ensure NVMe masking view is created for the relevant host before mounting the storage container to avoid intermittent BIND failures.

?
Cancel
Next

Create Storage Container

Storage Container
Storage Resources
Summary

Storage Container

Storage Container Name *

SCSI_vVol_Container

Description

Container for SCSI protocol

☒ SCSI
☐ NVMe

?
Cancel
Next

Figure 50. Create Storage Container – step 2

Add as many Storage Resources as desired for the SC [Figure 51](#). Data reduction is on by default, though the box can be unchecked if desired. Each time a line is complete, use the plus symbol on the right of the resource to add another line. The plus symbol will only appear when the cursor hovers near the end of the line. Unisphere will automatically generate a new name for each resource based on the SC name. If desired, change the name. All fields are required. An SC may have storage resources from multiple SRPs if they exist. The Limit represents the total storage in GB available for that SL. The wizard will prevent the user from having two storage resources with the same SL by removing it from the drop-down box. As the PowerMax does not have the concept of a workload, only one SL per container is possible. If more than one storage resource with the same SL is desired, create a second storage container. Note that the compression attribute does

not change the limitation of one SL type per storage container. For example, if a diamond SL resource with and without compression is desired, two storage containers are needed.

Note: Data reduction cannot be modified on an existing storage resource. A new storage resource in the existing or new container is required with the desired data reduction setting. Then migrate the VMs to the new resource.

Note: The Total Resource Subscribed Limit is a logical limit. Storage resources do not reserve space in the SRP. Consumers using the SRP outside of the vVol paradigm are not prevented from allocating beyond the size of free space available in the SC. However, vVol storage groups are included in the [Storage Group Demand Report](#) so actual usage is always reflected properly.

Create Storage Container

Storage Container

Storage Resources

Summary

Storage Resources

Name	SRP	Service Level		Limit (GB)	Data Reduction
SCSI_vVol_Container_resource_1	SRP_1	Bronze	i	1000	<input type="checkbox"/>
SCSI_vVol_Container_resource_2	SRP_1	Diamond	i	5000	<input checked="" type="checkbox"/>

Total Resource Subscribed Limit 6000GB Total Resources 2

?

Cancel

Back

Next

Figure 51. Create Storage Container – step 3

Review the final screen in [Figure 52](#) and select **Run Now** to create the SC.

Storage Container

Storage Resources

Summary

Summary

Storage Container Name
NVMe_vVol_Container

Description
Container for NVMe protocol

Storage Container Type
NVMe

Name	SRP	Service Level	Workload	Limit (GB)	Data Reduction
SCSI_vVol_Con...	SRP_1	Bronze	—	1000	✗
SCSI_vVol_Con...	SRP_1	Diamond	—	5000	✓

?

Cancel

Back

Run Now

Figure 52. Create Storage Container – step 4

Each time a task is run in Unisphere, a dialog box (Figure 53) appears providing the details of the task.

Task in progress

✓ Success

Starting Tasks...
Creating Storage Container [NVMe_vVol_Container]...
Succeeded
Add Storage Resource [SCSI_vVol_Container_resource_1] to Storage Container [NVMe_vVol_Container]...
Succeeded
Add Storage Resource [SCSI_vVol_Container_resource_2] to Storage Container [NVMe_vVol_Container]...
Succeeded

OK

Figure 53. Create Storage Container – completion

9.2.1.1 VASA 5

If using VASA 5, the user must provision the storage container to a virtual host(s) before it is available for vVol datastore creation in vCenter. This step was covered in the section [Storage container provisioning and de-provisioning](#).

9.3 Creating the VASA Replication Group in Unisphere

Replication with vVols is enabled at the storage container level, not the storage resource. Once a VASA Replication Group is created between two storage containers on two arrays, any storage resource in that container can support replication. In order to replicate a VM, a VM Storage Policy is required. It is not possible to simply create a VM in a vVol datastore without a storage policy if the VM is to be replicated.

Note: If desiring to migrate from one protocol to another, it is possible to setup a VASA replication group between two different types of containers and then run a planned migration in VMware Live Site Recovery⁵.

A VASA replication group (VRG) represents a relationship between two storage containers on the two PowerMax arrays involved in SRDF. In order to replicate a VM between two arrays, a VRG is required. Prior to VRG creation, the two PowerMax arrays should be configured for SRDF (e.g., networking, zoning, etc.) and there must be an existing, traditional SRDF group created, otherwise the VRG wizard will not recognize the remote array.

To create a new VRG, navigate to **vVol Dashboard -> Storage Containers**. Check the box next to the container to be used in the SRDF relationship and select **Create VASA Replication Group**.

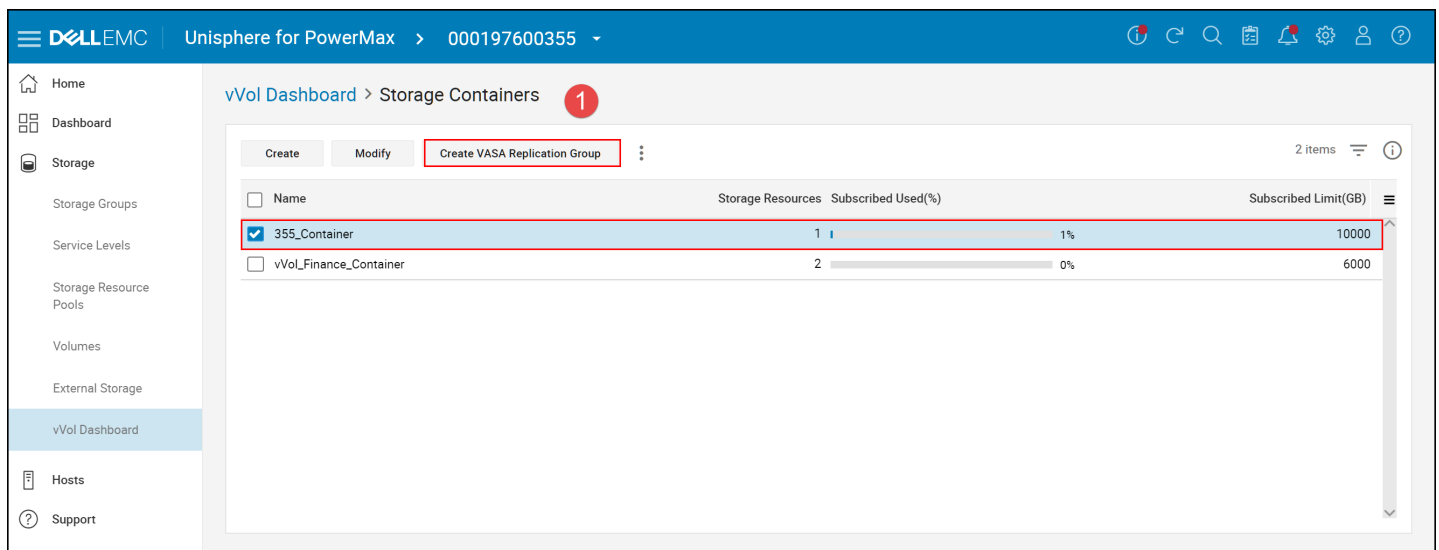


Figure 54. Creating a VASA Replication Group – step 1

Begin in [Figure 55](#) by selecting the Remote Array ID if there is more than one and add a Replication Group Label. This label will be displayed in vCenter when replicating a VM so the naming convention is important. Once the remote array is chosen, all remote storage containers on that array will be displayed. The remote storage container need not have the same storage

⁵ VMware Live Site Recovery is the new name in version 9.x for vCenter Site Recovery Manager (SRM). For backward compatibility, this paper will continue to use VMware SRM.

resources as the local storage container, nor does it have to be the same size; however, if SRM is going to be part of the solution, the remote container will be used for all test devices. It is not unreasonable, therefore, for the remote container to be twice the size of the local container. Remember, the size of the container is a logical amount, storage is only allocated as needed, but a test will fail if the total device size of the test exceeds the subscribed limit of the storage resource and/or storage container.

1 Select Remote

2 Configure Local

3 Configure Remote

4 Review Summary

Communication Protocol

FC

Remote Array ID

000197600358

Scan

Replication Group Label *

VP_355

Remote Storage Container

Name	Storage Resources	Subscribed Used(%)	Subscribed Limit(GB)
<input checked="" type="checkbox"/> 358_Container	1	1%	10000

?

CANCEL

NEXT

Figure 55. Creating a VASA Replication Group – step 2

In step 3 select a Replication Group Number and then the RDF ports. Dell recommends four ports, though in this example in Figure 56 only two are available.

1 Select Remote

2 Configure Local

3 Configure Remote

4 Review Summary

Replication Group Number *

17

3

Select Local Port

Select Local Port

✓ RF-1E:7

✓ RF-2E:7

Advanced Options

?

BACK

CANCEL

NEXT

Figure 56. Creating a VASA Replication Group – step 3

Similarly configure the remote array Replication Group Number and remote ports in [Figure 57](#).

1 Select Remote

2 Configure Local

3 Configure Remote

4 Review Summary

Remote Replication Group Number *

17

4

Select Remote Port

Select Remote Port

✓ RF-1G:10

RF-1G:11

✓ RF-2G:10

RF-2G:11

Advanced Options

?

BACK

CANCEL

NEXT

Figure 57. Creating a VASA Replication Group – step 4

Finally, in step 5 in [Figure 58](#), choose whether to enable hardware or software compression (optional), and then select Run Now.

Create VASA Replication Group

1 Select Remote
2 Configure Local
3 Configure Remote
4 Review Summary

Review Replication Group Summary

Replication Group Label	VP_355
Local Array	000197600355
Replication Group Number	17
Local Storage Container	355_Container
Local Ports	RF-1E:7, RF-2E:7
Remote Array	000197600358
Remote Replication Group Number	17
Remote Storage Container	358_Container
Remote Ports	RF-1G:10, RF-2G:10
Communication Protocol	FC
<input type="checkbox"/> Use Software Compression	
<input type="checkbox"/> Use Hardware Compression	

5

BACK CANCEL Run Now Add to Job List

Figure 58. Creating a VASA Replication Group – step 5

To view the details of a VRG, navigate from the vVol Dashboard to VASA Replication Groups. Select the checkbox next to the desired VRG. A VRG has a special attribute, **VASA Async**, which denotes it for use with Virtual Volumes highlighted in [Figure 59](#).

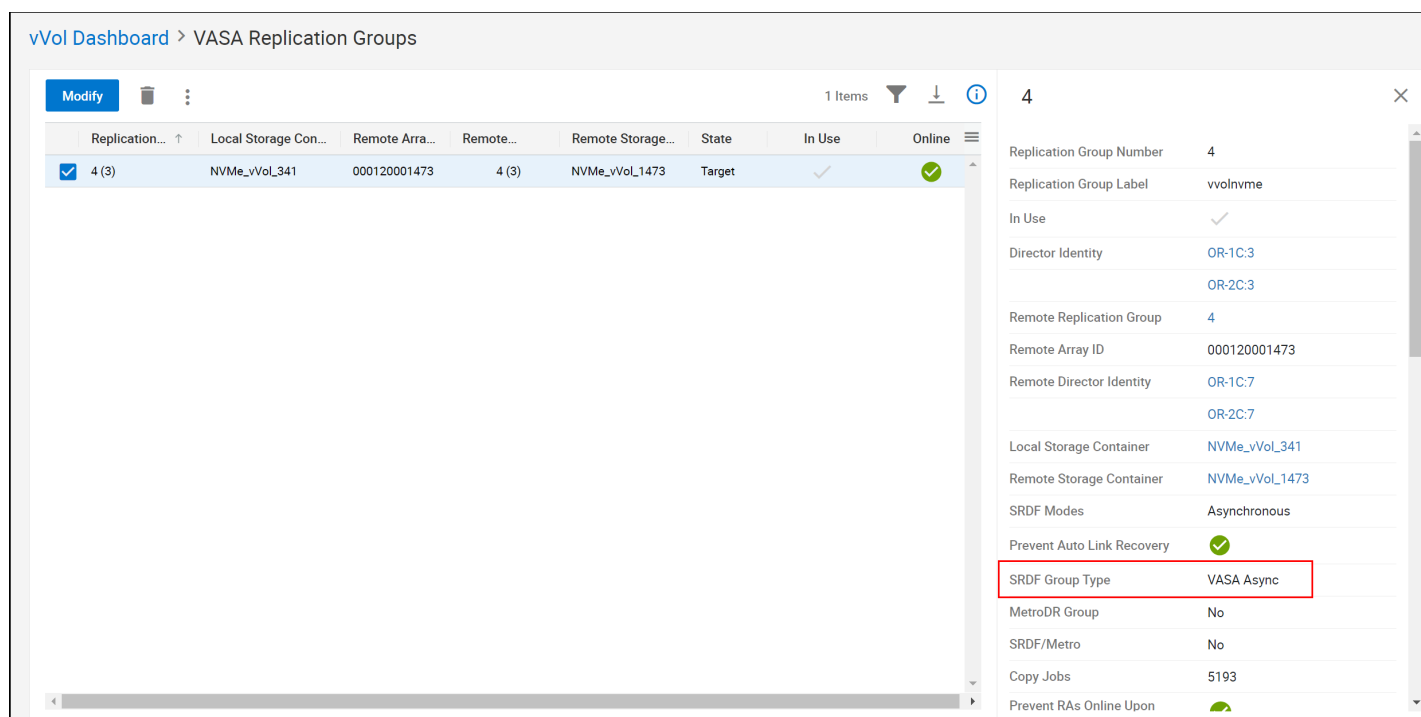


Figure 59. A VASA Replication Group

9.3.1 VRG State

A VRG has a “State” on both the local and remote side to designate its function. On the array it is created, the State will be “Source” while on the remote array, “Target”. The State may also change to reflect an activity taking place, such as an SRM test. The different possible states are:

- Source – the R1 site
- Target – the R2 site
- InTest - the VASA Replication Group is currently in a testfailover (shows on R2)
- FailedOver - the VASA Replication Group is failed over (shows on R2)

A VRG will reverse Source/Target during a reprotect operation.

9.4 Provisioning the Protocol Endpoint in Unisphere

The Protocol Endpoint is a small, physical device for SCSI (PE) and a virtual one for NVMe (vPE), that is used to enable IO between vSphere and the vVols on the array. vVols are bound and unbound to the PE/vPE by the VASA Provider, but once a vVol is bound to the PE/vPE, the VP is not required for IO to take place between the VM and the array. Replication is also never impacted as that is conducted between arrays.⁶ This means even if the VP crashes, IO and replication continues. Each ESXi host must be presented a unique PE/vPE to support vVols on the PowerMax. Each ESXi host in a cluster may not see any PE/vPE but the one uniquely presented to it.

⁶ The VP is required to conduct any type of management activity such as adding new vVols, snapshots, powering on or off the VM, or deleting the VM.

It should be noted that the SCSI PE, like vVols, uses a different World-Wide Name (WWN) than a traditional device. The format is known as a mobility safe ID. [Figure 60](#) shows a traditional WWN and [Figure 61](#) shows a mobility safe ID. One of the hallmarks of the traditional WWN is it has the array SID in the naming, while mobility safe ID does not.

Properties	Paths	Partition Details
General		
Name		EMC Fibre Channel Disk (naa.60000970000197700103533030303031)
Identifier		naa.60000970000197700103533030303031
Type		disk
Location		/vmfs/devices/disks/naa.60000970000197700103533030303031
Capacity		5.63 MB
Drive Type		HDD
Hardware Acceleration		Supported
Transport		Fibre Channel
Owner		NMP
Sector Format		512n
Multipathing Policies		
Path Selection Policy		Round Robin (VMware)
Storage Array Type Policy		VMW_SATP_SYMM

Figure 60. Traditional WWN format

Properties	Paths	Partition Details
General		
Name		EMC Fibre Channel Disk (naa.600009700bcbb733289000e9000000000)
Identifier		naa.600009700bcbb733289000e9000000000
Type		disk
Location		/vmfs/devices/disks/naa.600009700bcbb733289000e9000000000
Capacity		3.75 MB
Drive Type		HDD
Hardware Acceleration		Supported
Transport		Fibre Channel
Owner		NMP
Sector Format		512n
Multipathing Policies		
Path Selection Policy		Round Robin (VMware)
Storage Array Type Policy		VMW_SATP_SYMM

Figure 61. Mobility safe ID WWN format

Note: PE devices with mobility safe ID will not advertise ALUA support. This ensures the pathing software will not recognize the PE as an ALUA device.

Because the vPE is virtual, it will not be visible in vCenter (8.0.2 and higher) until a VM is powered on.

9.4.1 PE wizard

There are two different wizards for provisioning the protocol endpoint to account for the differences between SCSI and NVMe. While the SCSI wizard provisions a physical device to the selected host, the NVMe wizard creates an NVMe masking view which associates the host with a virtual, nonvisible protocol endpoint. Each wizard is covered below with the SCSI PE using the 10.0.1.0 Unisphere wizard and NVMe using the 10.1.0.0 Unisphere wizard.

Note: During the creation of the port group, if less than two ports are used or less than two directors, Unisphere will warn the user. This does not prevent creation.

9.4.1.1 SCSI PE

From the vVol dashboard, [Figure 62](#), navigate to **PROVISION PROTOCOL ENDPOINT TO HOST** in the **Actions** section.

The screenshot shows the Unisphere for PowerMax vVol Dashboard. The left sidebar contains navigation options: Overview, Dashboard, Storage (selected), Storage Groups, Service Levels, Storage Resource Pools, Volumes, vVol Dashboard (highlighted), Hosts, Data Protection, Performance, System, Events, and Support. The main content area is titled 'vVol Dashboard' and includes a 'Summary' section with four metrics: Storage Containers (2), Protocol Endpoints (4), PE Masking Views (4), and Replication Groups (1). Below this are two sections: 'System Consumed Capacity - Provisioned' showing a progress bar at 3% (5593.02 of 164202.1GB) and 'VASA Provider Status' with a table of online providers. At the bottom is the 'Array Resources' table. On the right, the 'Actions' panel is visible, with the 'PROVISION PROTOCOL ENDPOINT TO HOST' option highlighted by a red box and a red circle with the number 1.

Online	IP Address	Availability
✓	10.228.209.50	
✓	10.228.209.49	

Name	Subscribed Used(%)	Limit(GB)	Container	Data Reduction
ORACLE_302_PCLI...	28%	1000	ORACLE_302_PCLI	—
ORACLE_302_PCLI...	28%	1000	ORACLE_302_PCLI	✓
ORACLE_302_PCLI...	25%	20000	ORACLE_302_PCLI	✓

Figure 62. Provision SCSI PE – step 1

Start by selecting the correct initiator group (host) for the ESXi host to which the PE will be provisioned. Be sure the group is comprised of only FC or iSCSI initiators from a single host. Initiators from multiple hosts in a single group or cascaded initiator groups are not supported. A warning message indicates some of the restrictions in the dialog. The step is shown in [Figure 63](#).

Provision Protocol Endpoint to Host

Select Host

Select Port Group

Summary

Select Host

Enter name of Host or Host Group to filter

☐ dsib0041_iscsi_ig
☐ dsib0049_32_ig
☐ dsib0049_iscsi_ig
☐ dsib0049_tcp_ig
☒ dsib0051_32_ig 2
☐ dsib0051_iscsi_ig
☐ dsib0051_tcp_ig
☐ dsib0078_32_ig
☐ dsib0078_iscsi_ig
☐ dsib0078_tcp_ig

Create Host

Create Host Group

Important

Each ESXi host used with Virtual Volumes must have only one unique Protocol Endpoint. Protocol Endpoints may not be shared across ESXi hosts and ESXi hosts may not have multiple Protocol Endpoints from the same array. However, the same ESXi host may have a PE from each array in an environment.

?

Cancel

Next

Figure 63. Provision SCSI PE – step 2

In step 3, [Figure 64](#), select the port group which contains the zoned ports. The wizard offers the user the chance to create a new port group if desired. If selected, Unisphere will show the ports which are currently zoned and allow the user to choose which ones to add to the group.

Provision Protocol Endpoint to Host

Select Host ✓

Select Port Group

Summary

Select Port Group

☒ New ☐ Existing

Port Group Type

SCSI FC

Port Group Name *

dsib0051_32_pe_pg

3

<input checked="" type="checkbox"/> Dir-Port ↑	Identifier	Initiators	PGs	Mappings	% Busy
<input checked="" type="checkbox"/> OR-1C:0	50000972100...	54	1	36	0.7204
<input checked="" type="checkbox"/> OR-2C:0	50000972100...	37	1	36	0.7209

☐ Include ports not visible to host

?

Cancel

Back

Next

Figure 64. Provision SCSI PE – step 3

In the final review, Unisphere will automatically generate a name for the masking view and storage group based on the initiator group name. In [Figure 65](#) they have been renamed to better reflect the purpose of the masking view. When the job is run, Unisphere will create a PE device and add it to the new storage group, create the port group if needed, and finally create the masking view.

Provision Protocol Endpoint to Host

Select Host ✓

Select Port Group ✓

Summary

4

Summary

Masking View *

dsib0051_32_ig_MV

Storage Group *

dsib0051_32_ig_SG

Host

dsib0051_32_ig

Port Group

dsib0051_32_pe_pg

1 x Protocol Endpoint

?

Cancel

Back

Run Now

Figure 65. Provision SCSI PE – step 4

In the vSphere Client the PE device appears as a 3.75 MB TDEV shown in [Figure 66](#). Note the mobility safe ID.

Storage Devices									
Refresh Attach Detach Rename...									
Name	LUN	Type	Capacit...	Operational...	Hardware Accelera...	Drive T...	Transport		
Local USB Direct-Access (mpx.vmhba32:C0:T0:L1)	1	disk	0.00 B	Attached	Not supported	HDD	Block Adapter		
Local USB CD-ROM (mpx.vmhba32:C0:T0:L0)	0	cdrom		Attached	Not supported	HDD	Block Adapter		
Local USB Direct-Access (mpx.vmhba32:C0:T0:L2)	2	disk	0.00 B	Attached	Not supported	HDD	Block Adapter		
Local HL-DT-ST CD-ROM (mpx.vmhba65:C0:T0:L0)	0	cdrom		Attached	Not supported	HDD	Block Adapter		
EMC Fibre Channel Disk (naa.600009700bcb733289000e900000000)	10	disk	3.75 MB	Attached	Supported	HDD	Fibre Channel		
EMC Fibre Channel Disk (naa.60000970000197700103533030303031)	0	disk	5.63 MB	Attached	Supported	HDD	Fibre Channel		
EMC Fibre Channel Disk (naa.6000097000019770006253303030303643)	10	disk	5.63 MB	Attached	Supported	HDD	Fibre Channel		
EMC Fibre Channel Disk (naa.6000097000019760018853303030303031)	0	disk	5.63 MB	Attached	Supported	HDD	Fibre Channel		
EMC Fibre Channel Disk (naa.6000097000019760018853303030313732)	2	disk	5.63 MB	Attached	Supported	HDD	Fibre Channel		

Figure 66. PE device in vSphere

There is also a specific screen for Protocol Endpoints which can be accessed in [Figure 67](#).

Note: Note that the PE will not appear in its designated location ([Figure 67](#)) in the vSphere Client until the vVol datastore is created.

Protocol Endpoints

Name	Type	Storage array	Location	LUN	Operational state
EMC Fibre Channel Disk (naa.600009700bcb73328...	SCSI	VmaxVVolVasaProvider:6000...	/vmfs/devices/disks/...	10	Accessible

Properties

Paths

Datastores

General

Runtime name

Type

Identifier

Location

LUN

Operational state

Transport

Owner

Storage array

EMC Fibre Channel Disk (naa.600009700bcb733289000e900000000)

SCSI

naa.600009700bcb733289000e900000000

/vmfs/devices/disks/naa.600009700bcb733289000e900000000

10

Accessible

Fibre Channel

NMP

VmaxVVolVasaProvider:60000970000197900083F00000000000

Multipathing Policies

Path Selection Policy

Storage Array Type Policy

Round Robin (VMware)

VMW_SATP_SYMM

Figure 67. Protocol Endpoints sub-tab

9.4.1.2 NVMe vPE

Prior to creating an NVMe masking view, complete the following two tasks for all ESXi hosts:

1. Add at least one TCP software adapter and connect it to a virtual port on the array.
2. Create the vVol initiator group for the ESXi host.

If the user is unfamiliar with these tasks, they are covered in Appendix: Virtual Volume Operational Detail at the end of the document. Once created, proceed.

From the vVol dashboard, [Figure 68](#), navigate to **Create vVols NVMe Masking View** in the **Actions** section.

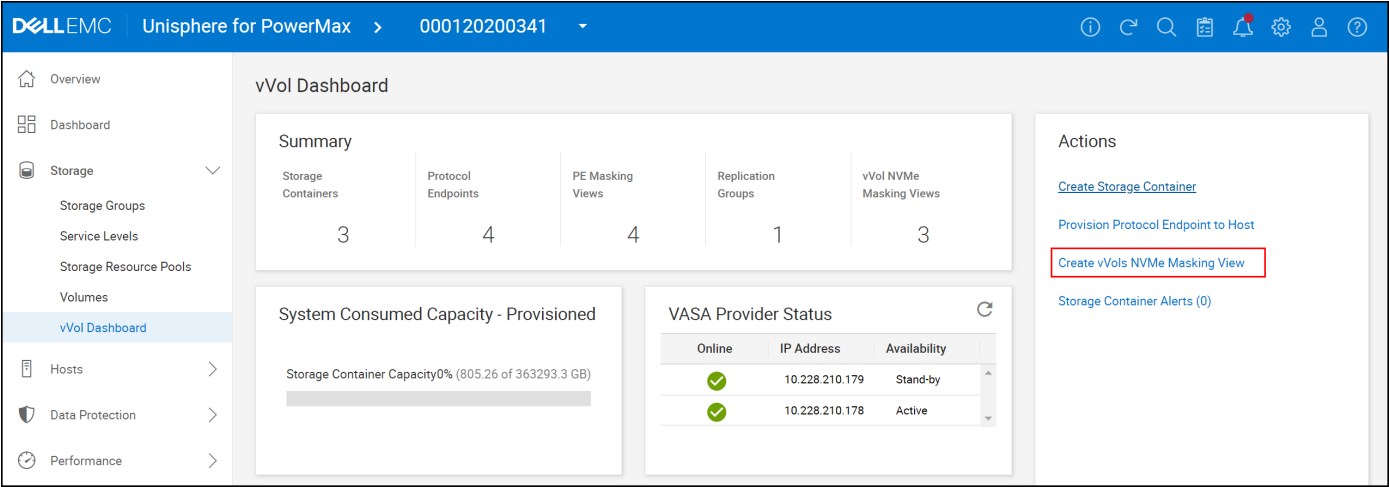


Figure 68. Provision NVMe vPE – step 1

Start by selecting the correct initiator group (host) for the ESXi host to which the vPE will be provisioned. Be sure the group is comprised of only NVMe vVol initiators from a single host. Initiators from multiple hosts in a single group or cascaded initiator groups are not supported. A warning message indicates some of the restrictions in the dialog. The step is shown in [Figure 69](#).

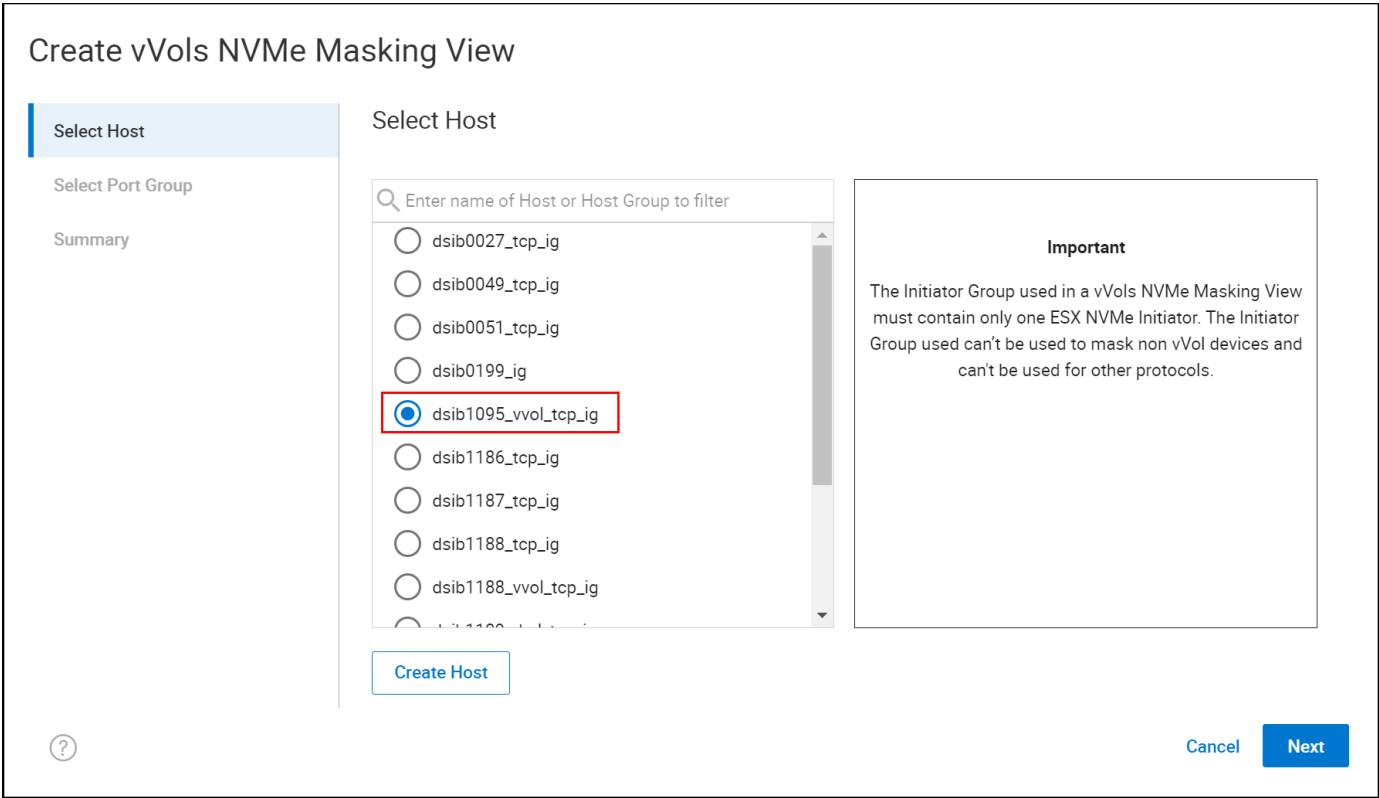


Figure 69. Provision NVMe vPE – step 2

In step 3 in [Figure 70](#), select an existing port group, or as in this case create a new one. Enter in a port group name and then check the box labeled **Include ports not visible to host**. The ports will not show up otherwise. Select the ports and then **Next**.

Create vVols NVMe Masking View

Select Host

Select Port Group

Summary

Select Port Group

New

Existing

Port Group Type

Port Group Name *

NVMe/TCP

dsib1095_vvol_tcp_pg

Dir-Port	Identifier	Initiators	PGs	Mappings	% Busy
OR-1C:000	—	0	1	0	—
OR-2C:000	—	0	1	0	—

Include ports not visible to host

Cancel

Back

Next

Figure 70. Provision NVMe vPE – step 3

In the final review, Unisphere will automatically generate a name for the masking view and storage group based on the initiator group name. In [Figure 71](#) they have been renamed to better reflect the purpose of the masking view. When the job is run, Unisphere will create a hidden storage group, the port group if needed, and finally create the masking view.

Create vVols NVMe Masking View

Select Host

Select Port Group

Summary

Summary

Masking View *

Host

Port Group

dsib1095_vvol_tcp_mv

dsib1095_vvol_tcp_ig

dsib1095_vvol_tcp_pg

Cancel

Back

Run Now

Figure 71. Provision NVMe vPE – step 4

Navigate to the vVol dashboard and select **vVol NVMe Masking Views**. Here one can see the vPE masking views in [Figure 72](#).

vVol Dashboard > Masking Views

[Rename](#) [View Path Details](#) 3 Items

<input type="checkbox"/>	Name ↑	Host/Host Group	Port Group	Storage Group
<input checked="" type="checkbox"/>	dsib1095_vvol_tcp_mv	dsib1095_vvol_tcp_ig	dsib1095_vvol_tcp_pg	_VVOLS_INTERNAL
<input type="checkbox"/>	dsib1188_vvol_tcp_mv	dsib1188_vvol_tcp_ig	tcp_pg	_VVOLS_INTERNAL
<input type="checkbox"/>	dsib1189_vvol_tcp_mv	dsib1189_vvol_tcp_ig	tcp_pg	_VVOLS_INTERNAL

Figure 72. vVol NVMe masking views

Unlike the SCSI PE, even if the user creates a datastore from the NVMe container, the vPE will not be visible. The user must first create and power on a VM. After this is done, navigate to the datastores screen in the left-hand panel, select the NVMe vVol datastore and in the right-hand panel select **Configure -> NVMe Protocol Endpoints**. The vPE is visible in Figure 73.

vSphere Client Search in all environments Administrator@BETA.LOCAL

dsib2009.drm.lab.emc.com

- LA
 - 341_DEV_23A
 - 341_DEV_26A
 - 341_DEV_26B
 - ALL_HOSTS_NO_RDF
 - ALL_HOSTS_NO_RDF_2
 - datastore1
 - NVMe_vVol_341
 - NVMe_vVol_Datastore**
 - PF_341_TEMP_3_1

NVMe_vVol_Datastore ACTIONS

Summary Monitor **Configure** Permissions Files Hosts VMs

Alarm Definitions
Scheduled Tasks
General
Connectivity with Hosts
NVMe Protocol Endpoints
Capability sets
Default profiles

NVMe Protocol Endpoints

Target NGN	Storage array	Online
nqn.1988-11.com.dell:PowerMax_8500:00:000120200341	VmaxVVolVasaProvider:60000970000120200341F00000000000	true

1 item

Figure 73. vPE in vCenter

9.5 Using Solutions Enabler with Virtual Volumes

In addition to Unisphere, vVols may be managed through Solutions Enabler. Almost all capabilities that exist in Unisphere for vVols are available through the CLI, though not in “wizard” format. The exception is the VASA Provider Status which has no associated command. The following will detail the commands available to the user through SE for vVols. For each creation statement a deletion statement follows in parentheses. The CLI examples use Solutions Enabler 10.1.0.0.

9.5.1 Creation of the Storage Container

There are two parts to creating the storage container: the container object and the storage resource object. A storage container by itself has no actual storage associated with it, rather it is a logical grouping of storage resources. Storage resource objects, represented by an SL and size, are added to the container to provide the storage from which to provision vVols.

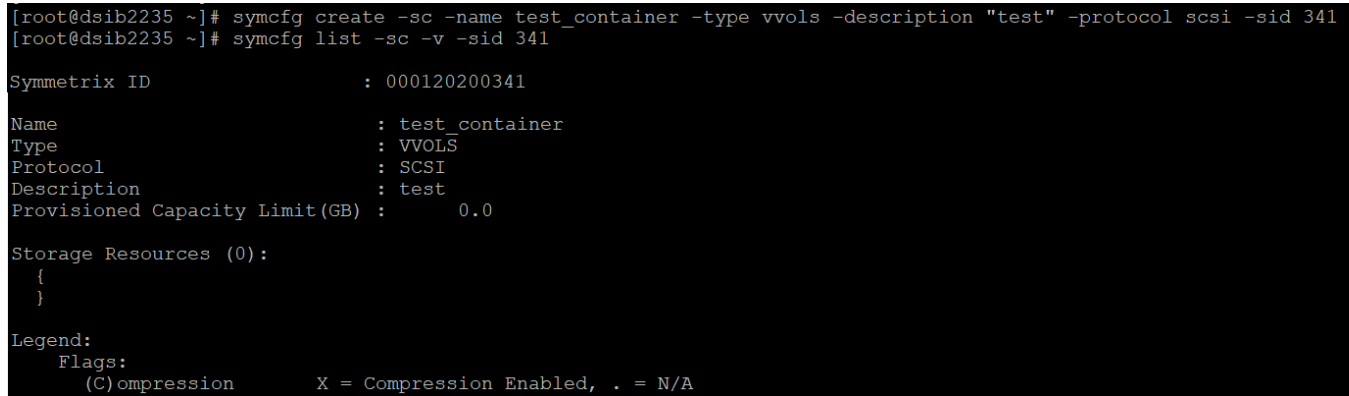
9.5.1.1 Storage Container

To create the container use the following syntax. Supply the protocol as scsi or nvme:

```
symcfg create -sc -name test_container -type vvols -description "test" -
protocol scsi -sid 341
(symcfg -sid 341 delete -sc -sc_name test_container)
```

There is no response to the command. To list the container:

```
symcfg list -sc -v -sid 341
```



```
[root@dsib2235 ~]# symcfg create -sc -name test_container -type vvols -description "test" -protocol scsi -sid 341
[root@dsib2235 ~]# symcfg list -sc -v -sid 341

Symmetrix ID          : 000120200341
Name                  : test_container
Type                  : VVOLS
Protocol              : SCSI
Description            : test
Provisioned Capacity Limit(GB) : 0.0

Storage Resources (0):
{
}

Legend:
Flags:
(C)ompression      X = Compression Enabled, . = N/A
```

Figure 74. Create and list storage container

Note: If a storage container is associated with a VASA Replication Group but otherwise empty, it can be deleted. This disassociates the replication group with the container and effectively makes it unusable. The VRG can still be deleted at that point, but Dell recommends removing the VASA Replication Group before removing the container. This ensures there are no RDF pairs that might still be in the storage container.

Note: If a storage container is provisioned to a virtual host, it does not prevent the deletion of it.

9.5.1.2 Storage Resources

Once the container is created, storage resources can be added, one for each SRP and SL combination as required. The subscribed maximum defaults to GB.

```
symcfg -sid 341 -sc -sc_name test_container add -sresource Gold_Resource -slo
Gold -subscribed_max 1024
```

```
(symcfg -sc -sc_name test_container remove -sresource Gold_Resource -sid 341)
```

Included in the screenshot below is first an attempt to add a service level with a workload. The PowerMax does not support workloads so the error shown will be generated. Following that is the correct syntax on the PowerMax where the Workload column will always be "<none>".


```
[root@dsib2235 ~]# symcfg create -sc -name test_container -type vvols -description "test" -protocol scsi -sid 341
[root@dsib2235 ~]# symcfg -sid 341 -sc -sc_name test_container add -sresource Gold_Resource -slo Gold -subscribed_max 1024
[root@dsib2235 ~]# symcfg list -sc -v -sid 341
```

Symmetrix ID : 000120200341

Name : test_container
Type : VVOLS
Protocol : SCSI
Description : test
Provisioned Capacity Limit(GB) : 1024.0

Storage Resources (1):

```
{
  -----
  Name                Flg Service Level      SRP          Capacity
                   C  Name          Workload Name  Limit Prov
                   -----
  Gold_Resource       X  Gold          <none>      SRP_1
                                     1024.0
  Total
  }
                                     1024.0
```

Legend:
Flags:
(C)ompression X = Compression Enabled, . = N/A

Figure 75. Add storage resources to container

If the `--detail` flag is added to the list command, the subscribed usage is displayed as in [Figure 76](#).

```
[root@dsib2235 ~]# symcfg list -sc -v -sid 341 -detail
```

Symmetrix ID : 000120200341

Name : 8U1
Type : VVOLS
Protocol : SCSI
Description :
Provisioned Capacity Limit(GB) : 41500.0
Provisioned Capacity (GB) : 263.3
Provisioned Capacity (%) : 0

Storage Resources (3):

```
{
  -----
  Name                Flg Service Level      SRP          Capacity
                   C  Name          Workload Name  Limit Prov
                   -----
  8U1_resource_1     X  Optimized    <none>      SRP_1
                                     20000.0  263.0  1  3.3:1
  8U1_resource_2     X  Diamond     <none>      SRP_1
                                     20000.0   0.0  0  N/A
  8U1_resource_3     .  Gold        <none>      SRP_1
                                     1500.0   0.3  0  N/A
  Total
  }
                                     41500.0  263.3  0  N/A
```

Legend:
Flags:
(C)ompression X = Compression Enabled, . = N/A

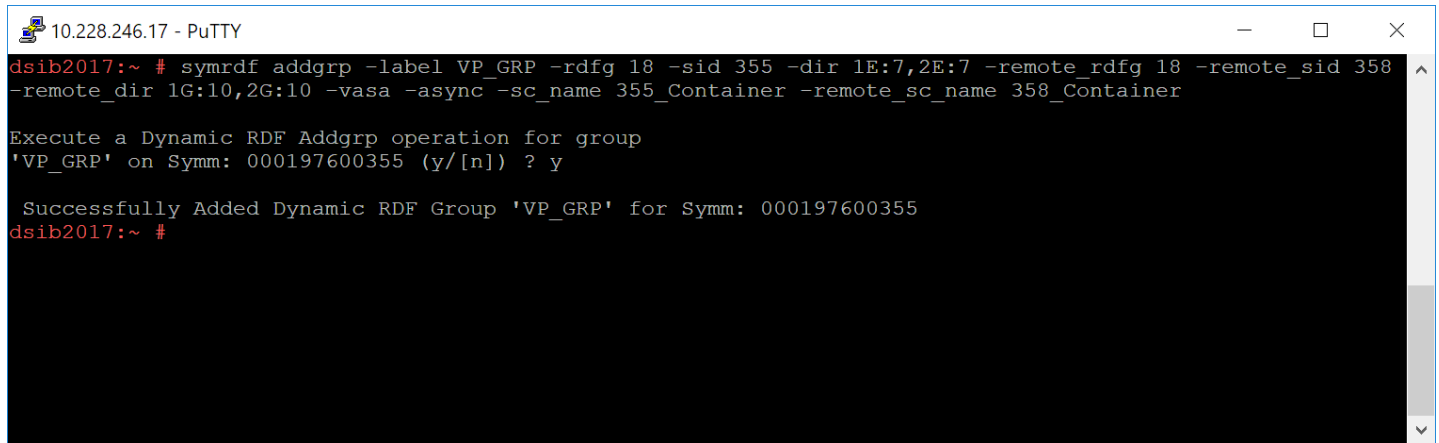
Figure 76. Listing storage usage of container

9.5.2 Creation of the VASA Replication Group

To create a VASA Replication Group from the command line:

```
symrdf addgrp -label VP_GRP -rdfg 18 -sid 355 -dir 1E:7,2E:7 -remote_rdfg 18
-remote_sid 358 -remote_dir 1G:10,2G:10 -vasa -async -sc_name 355_Container -
remote_sc_name 358_Container
```

```
(symrdf removegrp -rdfig 18 -sid 355)
```



```
10.228.246.17 - PuTTY
dsib2017:~ # symrdf addgrp -label VP_GRP -rdfig 18 -sid 355 -dir 1E:7,2E:7 -remote_rdfg 18 -remote_sid 358
-remote_dir 1G:10,2G:10 -vasa -async -sc_name 355_Container -remote_sc_name 358_Container

Execute a Dynamic RDF Addgrp operation for group
'VP_GRP' on Symm: 000197600355 (y/[n]) ? y

Successfully Added Dynamic RDF Group 'VP_GRP' for Symm: 000197600355
dsib2017:~ #
```

Figure 77. VASA Replication Group creation

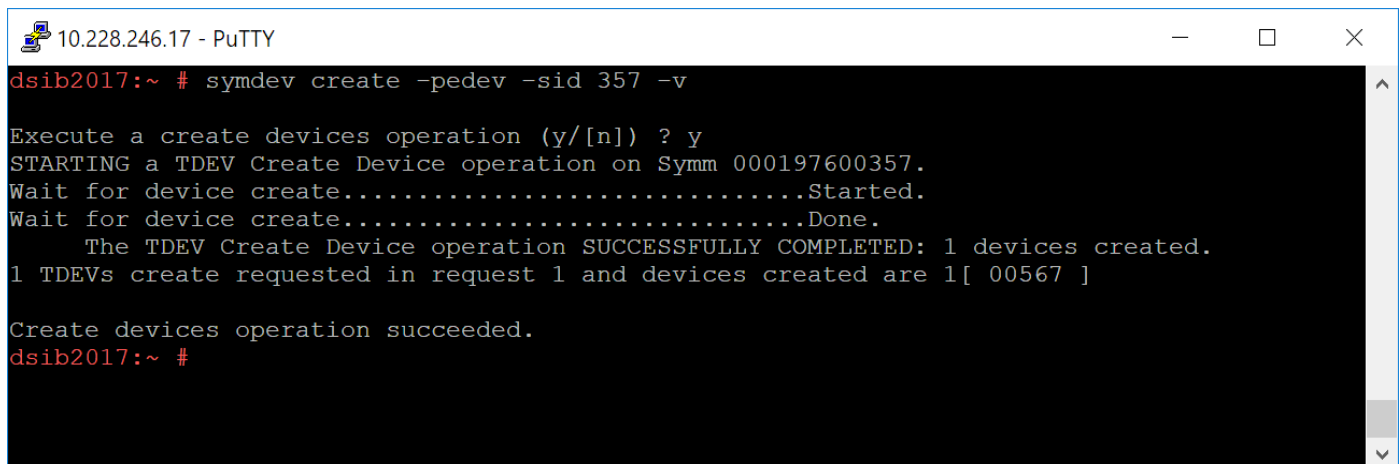
9.5.3 Creation of the SCSI PE

To create a PE device:

```
symdev create -pedev -sid 357 -v
```

```
(symdev delete 58 -sid 357)
```

Note the `-v` (verbose) flag is not required, however without it the device ID will not be returned. Figure 78 shows the output.



```
10.228.246.17 - PuTTY
dsib2017:~ # symdev create -pedev -sid 357 -v

Execute a create devices operation (y/[n]) ? y
STARTING a TDEV Create Device operation on Symm 000197600357.
Wait for device create.....Started.
Wait for device create.....Done.
The TDEV Create Device operation SUCCESSFULLY COMPLETED: 1 devices created.
1 TDEVs create requested in request 1 and devices created are 1[ 00567 ]

Create devices operation succeeded.
dsib2017:~ #
```

Figure 78. PE device creation

A PE device is like any other TDEV – once created it can be added to a storage group and then presented to an ESXi host. Remember that each ESXi host must see its own unique PE device and the initiator group for that masking view may only contain FC or iSCSI initiators of that host (no cascading or initiators from other ESXi hosts).

9.5.4 Creation of the NVMe vPE

The vPE is not created like the PE in CLI since there is no device to create. Instead, the masking view can be created with the “vasa” flag. This is the syntax:

```
symaccess create view -name nvme_vvol_mv -ig dsib1095_vvol_tcp_ig -pg
dsib1095_vvol_tcp_pg -vasa -sid 341
```

Note that all the same autoprovisioning restrictions apply for the vPE as the PE – one masking view per ESXi host and the initiator group may only contain the vVol initiator from one host.

9.6 Host IO Limits/Storage IO Control (SIOC)

Host IO Limits is a feature of the PowerMax that allow users to place limits on the front-end bandwidth and the IOPS consumed by applications. Currently, vVols do not support the use of Host IO Limits. While it is possible to set a Host IO limit on the storage group that contains the Protocol Endpoint, it will have no bearing on the vVol IOs.

vVols support the use of VMware Storage I/O Control at the Storage Policy level, providing a way to limit IO to virtual machines. [Figure 79](#) demonstrates the functionality.

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 Host based services**
- 4 VmaxVVOLProvider rules
- 5 Storage compatibility
- 6 Review and finish

Host based services

Create rules for data services provided by hosts. Available data services could include encryption, I/O control, caching, etc. Host based services will be applied in addition to any datastore specific rules.

Encryption **Storage I/O Control**

☐ Disabled

☒ Use storage policy component **Low IO shares allocation**

Storage policy component	Low IO shares allocation
Description	Storage policy component for Low SIOC controls
Provider	VMware Storage IO Control
VMware Storage I/O Control	
IOPS limit	1,000
IOPS reservation	10
IOPS shares	500

☐ Custom

CANCEL BACK NEXT

Figure 79. Storage I/O Control with vVols

10 Creating a vVol Datastore

Creating the vVol datastore is the final step in the vVol setup. It relies both on successful registration of the VASA Provider and presentation of the Protocol Endpoint to the host. If the VASA Provider is registered but the Protocol Endpoint is not presented, vVol datastore creation will succeed, but then the datastore will enter an inaccessible status.

Though datastore creation through vCenter is a common task, it is covered herein since vVols are a new paradigm.

10.1 vVol datastore wizard

Start by accessing the Storage icon in the **Home** page of the vSphere Client. Then in [Figure 80](#) below highlight the datacenter, select the **Datastores** tab on the right menu, then from the **ACTIONS** menu, navigate to **Storage -> New Datastore**.

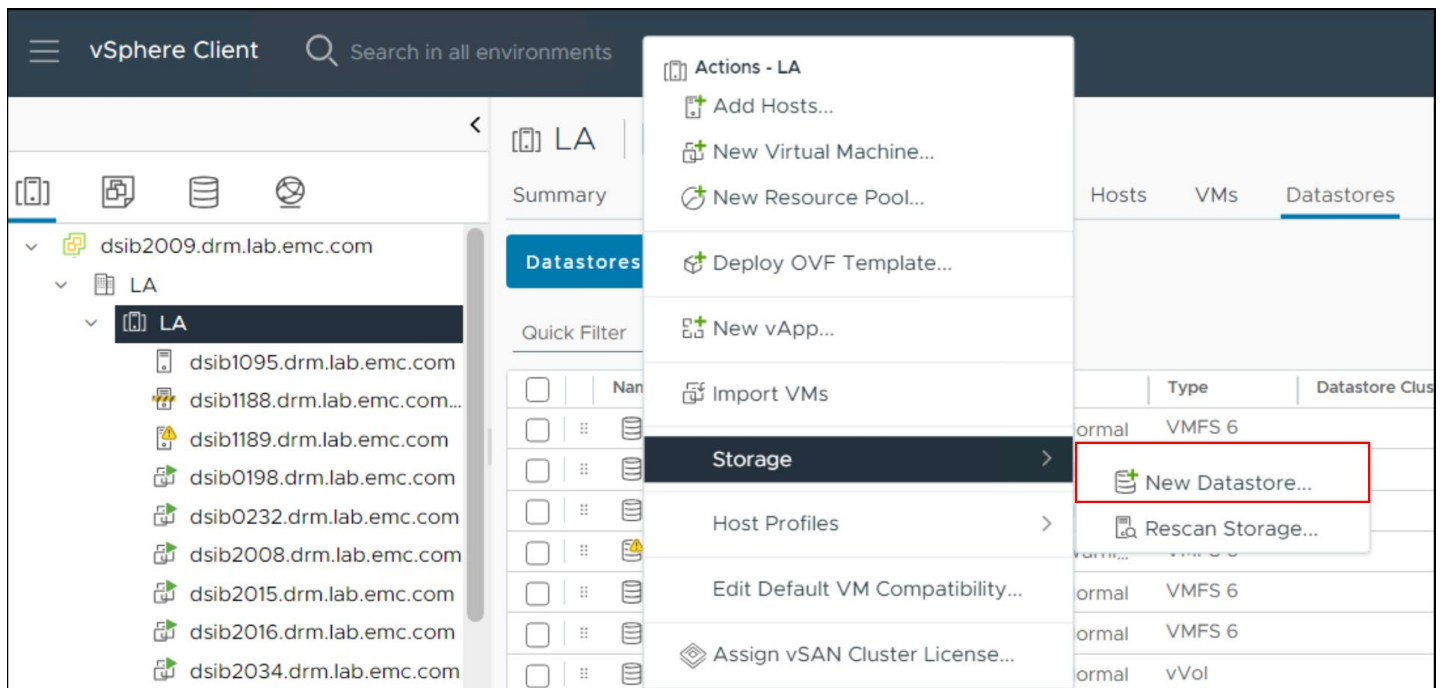


Figure 80. Creating a vVol datastore – step 1

Step 2 prompts for the type of datastore. Select **vVol** and **NEXT**.

New Datastore

- 1 Type**
- 2 Name and container selection
- 3 Hosts accessibility
- 4 Ready to complete

Type

Specify datastore type.

☐ VMFS
Create a VMFS datastore on a disk/LUN.

☐ NFS
Create an NFS datastore on an NFS share over the network.

☒ **vVol**
Create a Virtual Volumes datastore on a storage container connected to a storage provider.

CANCEL NEXT

Figure 81. Creating a vVol datastore - step 2

Select a name for the datastore in step 3 and the storage container that is to be associated with that datastore. Recall that the vVol datastore is the vSphere representation of the PowerMax storage container in the vCenter. Figure 82 shows the three storage containers from which to choose. If using VASA 5 these containers would have been provisioned to the virtual host (vCenter). As mentioned, there is no indication whether a container is SCSI or NVMe which is why the naming is critical when using both protocols. Note that the column “Maximum Disk Size” will always show as 16 TB. This refers to the maximum size of a single vVol (vmdk) that Dell imposes, even though VMware supports 62 TB. It has no association with the size of the storage container. The size of the container will not be visible until after datastore creation, and that has no imposed limit.

New Datastore

- 1 Type
- 2 Name and container selection**
- 3 Ready to complete

Name and container selection

Specify datastore name and backing storage container.

Name NVMe_vVol_Datastore

Backing Storage Container

	Name	Identifier	Maximum Disk Size	Exist
<input type="radio"/>	NVMe_vVol_341	vvol:60000970072a1c95-4dee014c0000f...	16 TB	--
<input checked="" type="radio"/>	NVMe_vVol_Container	vvol:60000970072a1c95-4dee014c0001e...	16 TB	--
<input type="radio"/>	SCSI_vVol_Container	vvol:60000970072a1c95-4dee014c0001e...	16 TB	--

6 items

Warning: For SCSI-backed vVol datastores, PE LUNs need to be configured manually. Configure SCSI PE LUNs before creating a datastore. If the datastore is created without configuring PE LUNs, the ESXi host marks corresponding vVol datastore as inaccessible.

Backing Storage Container Details

Storage array(s) 000120200341

Storage provider(s) VASA_0

CANCEL BACK NEXT

Figure 82. Creating a vVol datastore - step 3

In step 4 select on to which hosts the datastore should be mounted. In Figure 83 the three available hosts are chosen. Be sure any selected hosts have a PE or vPE depending on the container type.

New Datastore

- Type
- Name and container selection
- Hosts accessibility
- Ready to complete

Hosts accessibility

Specify which hosts will have access to the datastore.

COMPATIBLE (3 HOSTS) **INCOMPATIBLE (0 HOSTS)**

Quick Filter

Host	Cluster
<input checked="" type="checkbox"/> dsib1095.drm.lab.emc.com	LA
<input checked="" type="checkbox"/> dsib1188.drm.lab.emc.com	LA
<input checked="" type="checkbox"/> dsib1189.drm.lab.emc.com	LA

☒ 3 3 items

CANCEL **BACK** **NEXT**

Figure 83. Creating a vVol datastore - step 4

In step 5 in [Figure 84](#) review the chosen options and when ready select Finish. The datastore will be created and mounted to the host.

New Datastore

- Type
- Name and container selection
- Hosts accessibility
- Ready to complete

Ready to complete

Review your selections before finishing the wizard

▼ **Name and container selection**

Datastore name: NVMe_vVol_Datastore

Datastore type: vVol

Storage container name: NVMe_vVol_Container

Storage container UUID: vvol:60000970072a1c95-4dee014c0001e78a

Storage array(s): 000120200341

Storage provider(s): VASA_0

▼ **Hosts accessibility**

Hosts

- dsib1095.drm.lab.emc.com
- dsib1188.drm.lab.emc.com
- dsib1189.drm.lab.emc.com

CANCEL **BACK** **FINISH**

Figure 84. Creating a vVol datastore – step 5

Once created, the vVol datastore's storage capacity as viewed in vSphere ([Figure 85](#)) is the sum of the subscribed capacity for the storage container's storage resources. In this case there is two storage resources totaling just under 6 TB as the actual capacity assigned to the container is somewhat less in vSphere due to rounding and metadata.

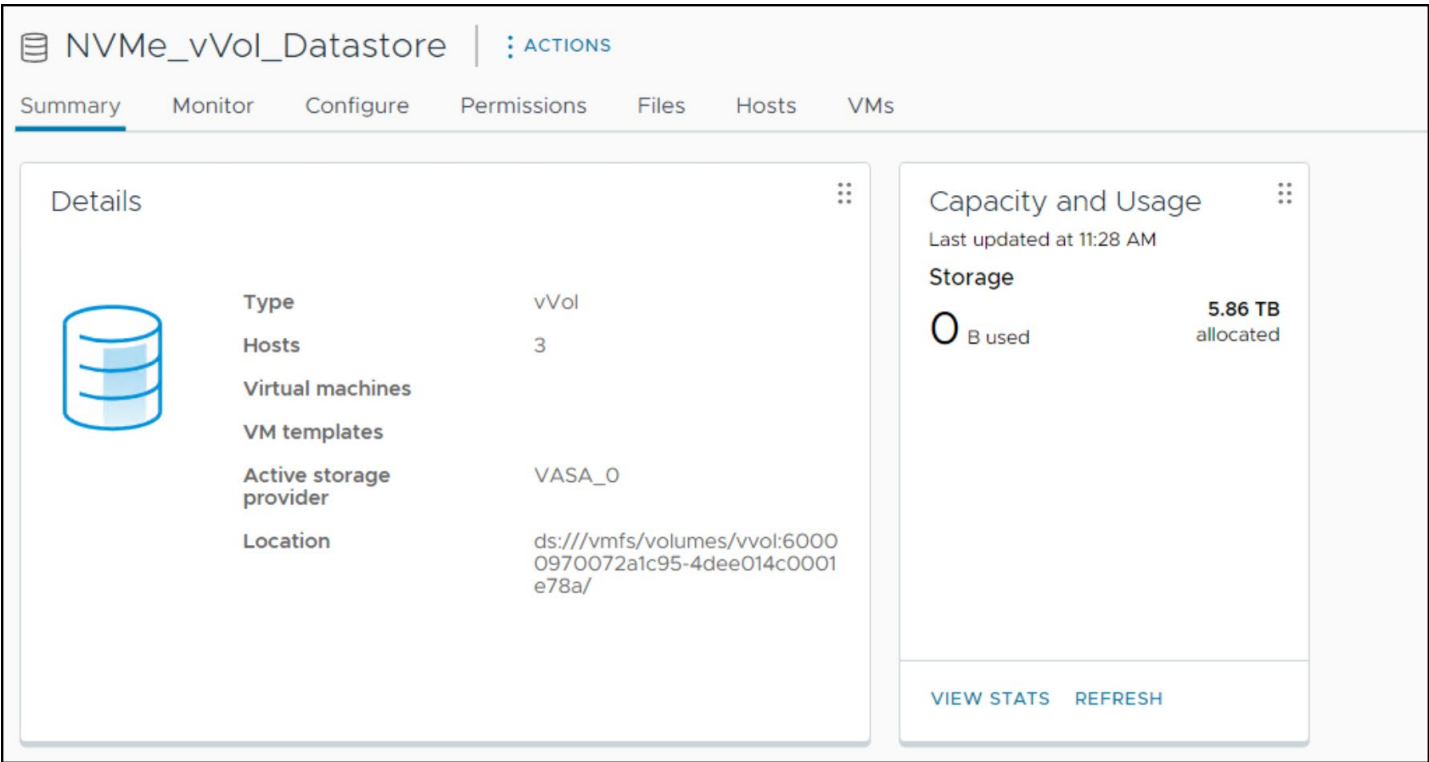


Figure 85. vVol datastore capacity

Note: vCenter will allow the creation of both SCSI and NVMe vVol datastore even if the masking view(s) is not present; however, the storage will show as zero (0) and the datastore will become inaccessible.

10.2 Modifying the Storage Container in Unisphere

If an existing storage resource in a storage container is changed – space added or removed – those changes will be reflected in vSphere upon refresh of the datastore. The following section provides an example.

Note: Storage containers and resources cannot be renamed. If names must be adjusted, existing VMs would have to be moved into a new storage resource, and/or container with the correct name.

At any time, a storage container may be modified by the storage administrator. A container may have a new storage resource added or removed or space added or removed to existing storage resources. The service level of an existing storage resource may not be changed. In order to see the changes in the vCenter, one of two actions will be required. If a storage resource is added, it is necessary to rescan the VASA Provider in the vCenter. If an existing storage resource is modified, however, a simple refresh of the vVol datastore will show the changes.

10.2.1 SC modification steps

From the vVol dashboard, [Figure 86](#), select the **Storage Containers** icon.

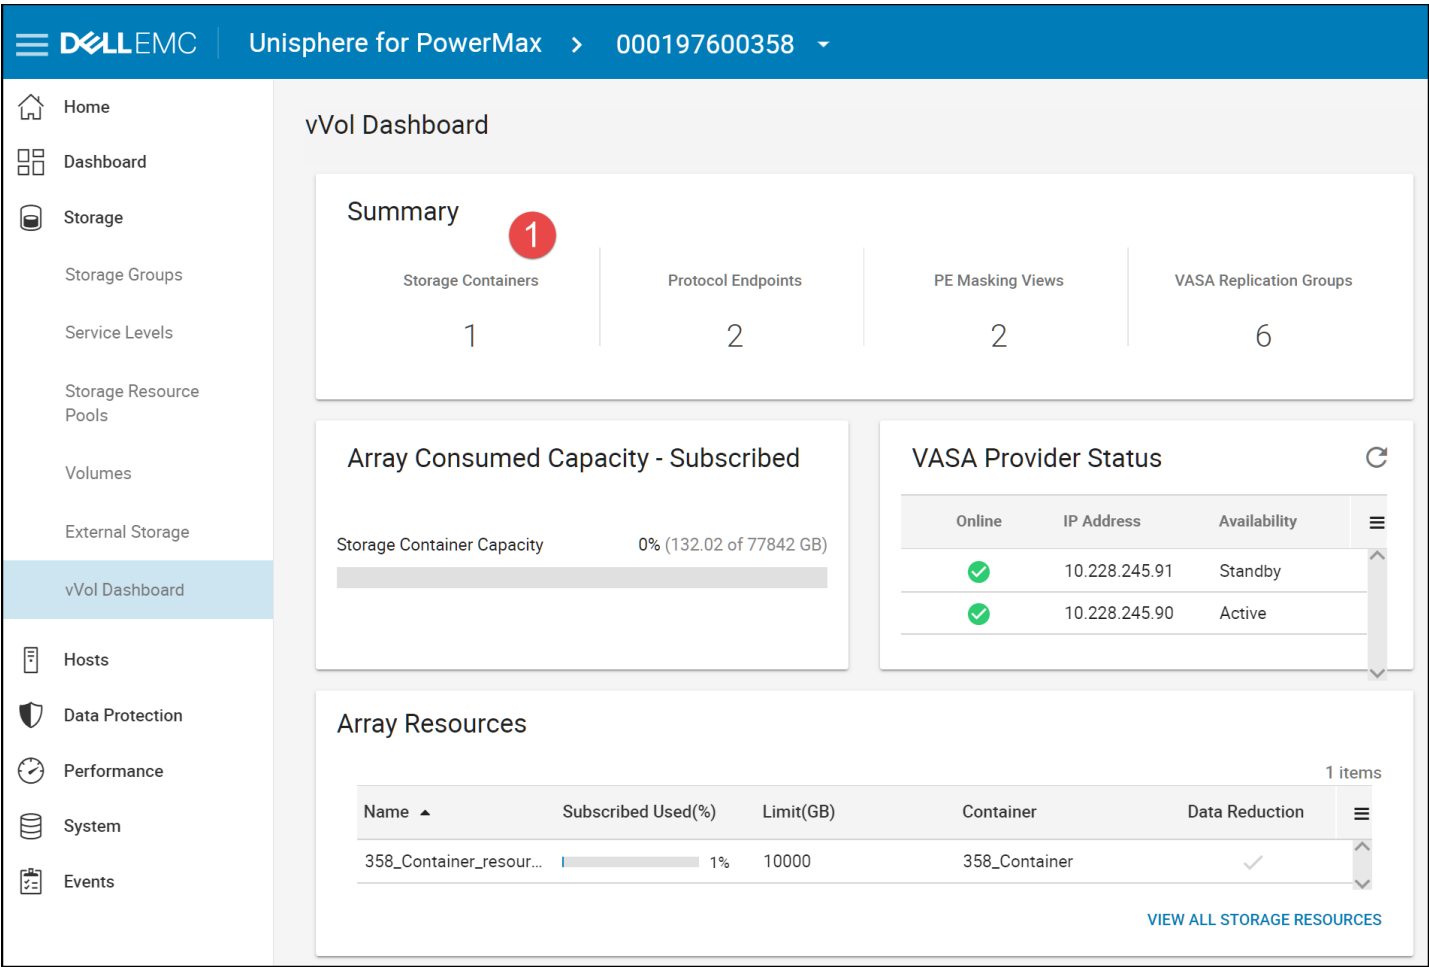


Figure 86. Modify Storage Container – step 1

In step 2 highlight the storage container and double-click to reveal the details in the panel on the right-hand side shown in Figure 87. Then in step 3 click on the hyperlink for the **Storage Resources** in the panel.




Figure 87. Modify Storage Container – step 2-3

In step 4 highlight one of the storage resources and select **Modify**. This is shown in Figure 88.

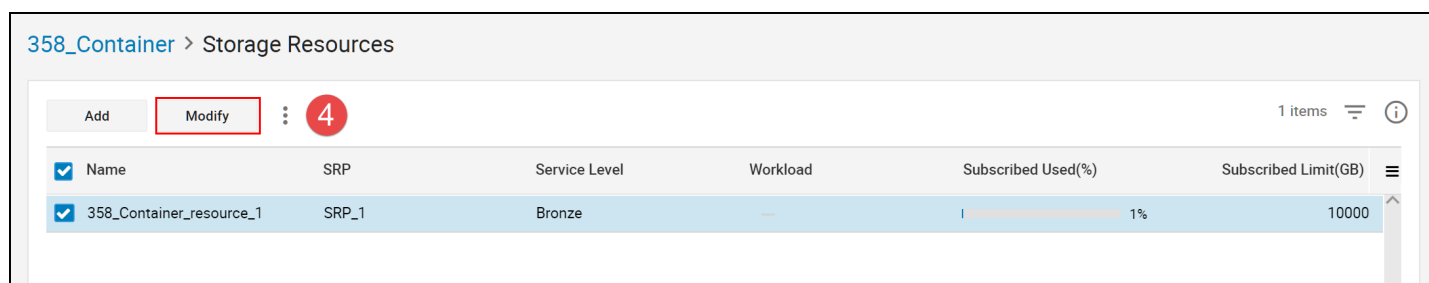


Figure 88. Modify Storage Container – step 4

In step 5 modify the Subscribed Limit to the new desired value. In this example the storage resource is increased from 10000 GB to 15000 GB as in Figure 89 and select OK.

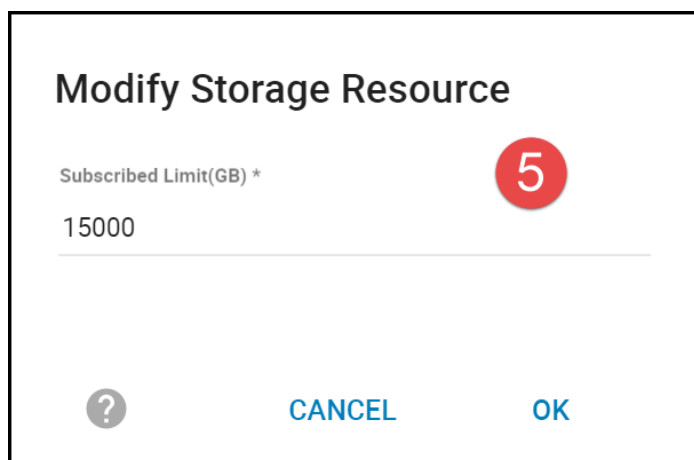


Figure 89. Modify Storage Container – step 5

The final result is shown in Figure 90.

358_Container > Storage Resources

AddModify

1 items

<input checked="" type="checkbox"/>	Name	SRP	Service Level	Workload	Subscribed Used(%)	Subscribed Limit(GB)
<input checked="" type="checkbox"/>	358_Container_resource_1	SRP_1	Bronze	—	<div><div></div></div> 1%	15000

Figure 90. Modify Storage Container – completion

There are limitations as to what is possible when modifying the storage container. For instance, a storage resource cannot be resized below its used capacity. A storage resource also cannot be removed from a storage container while vVols are bound to it.

10.2.2 Recognizing new SC size in vCenter

After the storage resource is modified, it is necessary to refresh the information in the vCenter. Simply select the **REFRESH** button in the vVol datastore screen to reflect the changes. This can be done from a variety of screens in the vCenter. Figure 91 demonstrates the change in the vVol datastore capacity when 2000 GB is added to the storage resource.

358_vVol_Datastore ACTIONS

SummaryMonitorConfigurePermissionsFilesHostsVMs

Alarm Definitions
Scheduled Tasks
General
Connectivity with Hosts
Protocol Endpoints
Capability sets
Default profiles

Properties

Name358_vVol_Datastore

TypevVol

Capacity

Total Capacity9.77 TB

Provisioned Space132.01 GB

Free Space9.64 TB

Default Storage Policy

PolicyVVol No Requirements Policy

REFRESH

Capacity

Total Capacity14.65 TB

Provisioned Space132.01 GB

Free Space14.52 TB

REFRESH

Figure 91. Refresh vVol datastore

If, however, a storage resource is added to the container rather than an existing one modified, as in Figure 92, the VASA Provider must be rescanned.

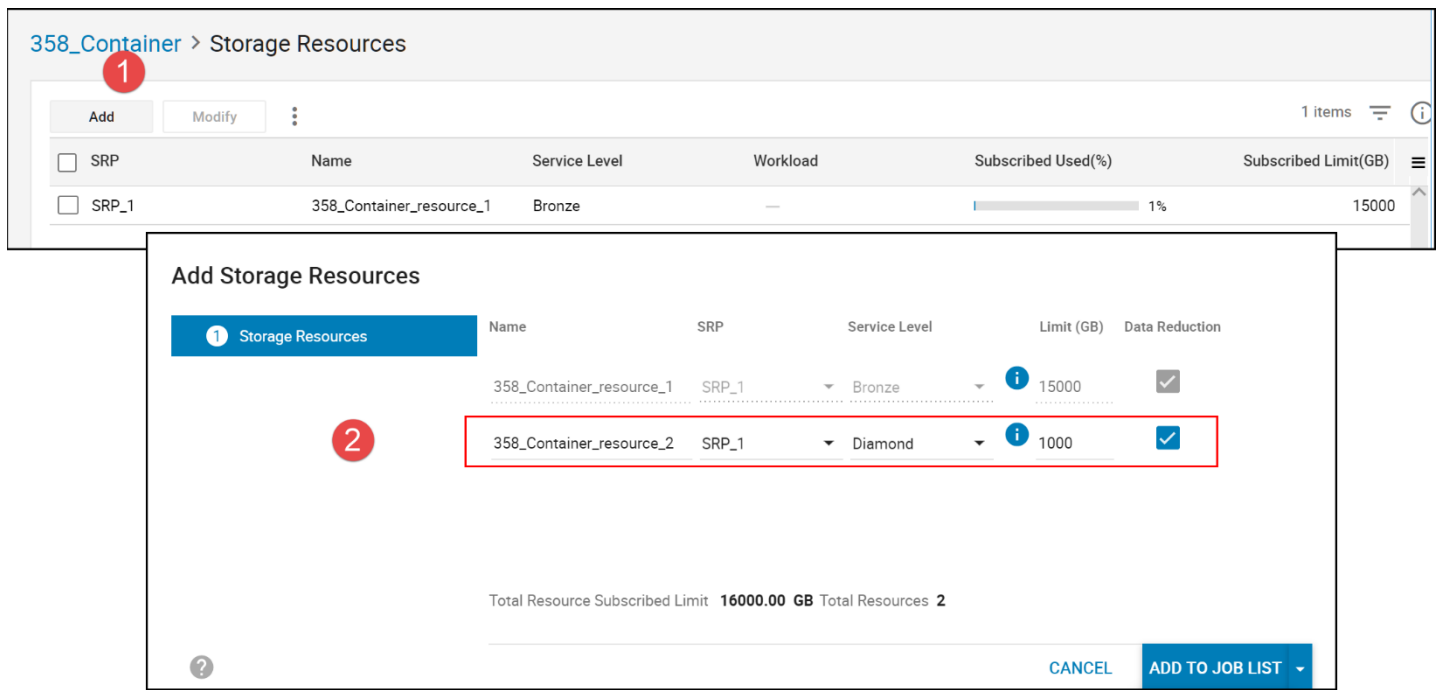


Figure 92. Adding a storage resource to a storage container

Then highlight the Dell VASA Provider in step 1 and select the **Rescan** option in step 2 as demonstrated in Figure 93. Note that in vSphere 8 the rescan option is to the left of the VASA Provider by select three buttons.

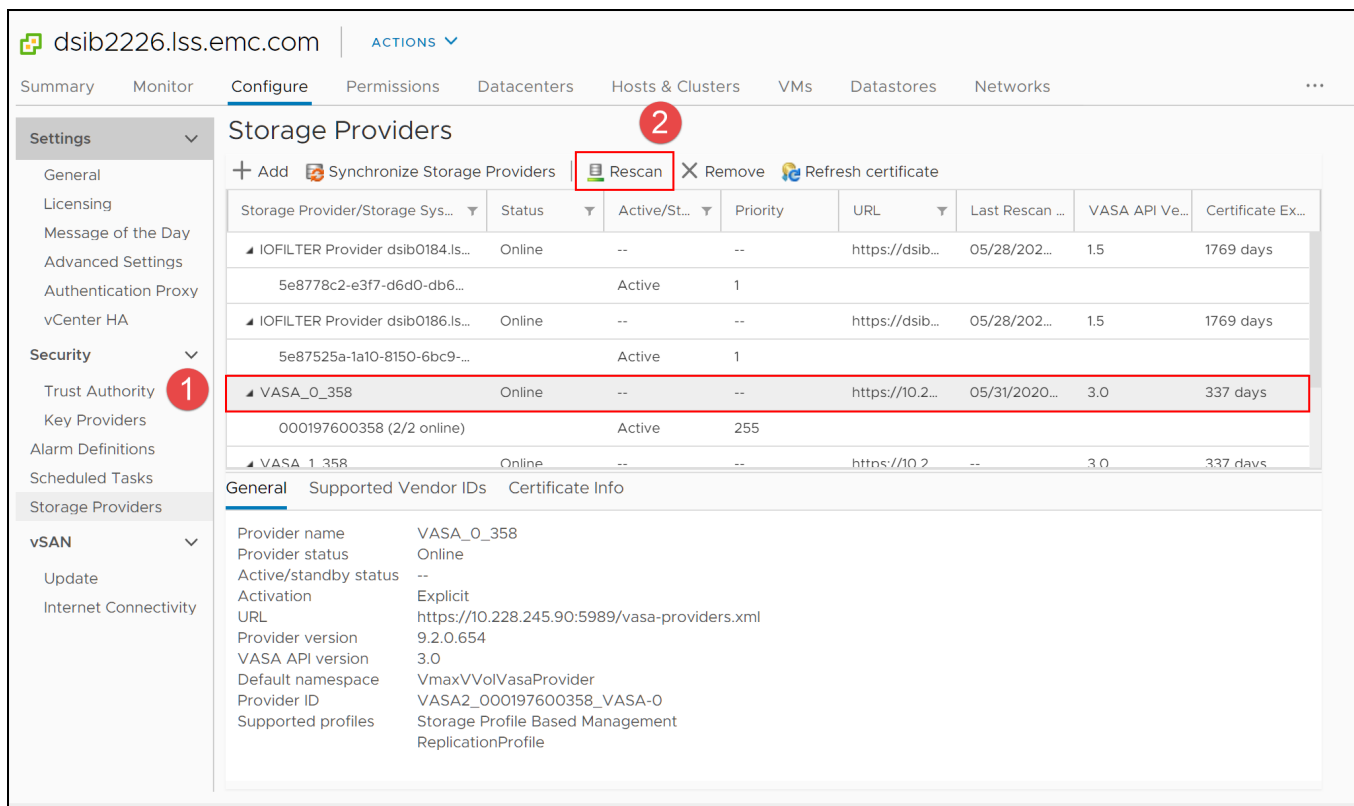


Figure 93. Rescan the VASA Provider

Once rescanned, the additional capacity will show for the vVol datastore below in [Figure 94](#).

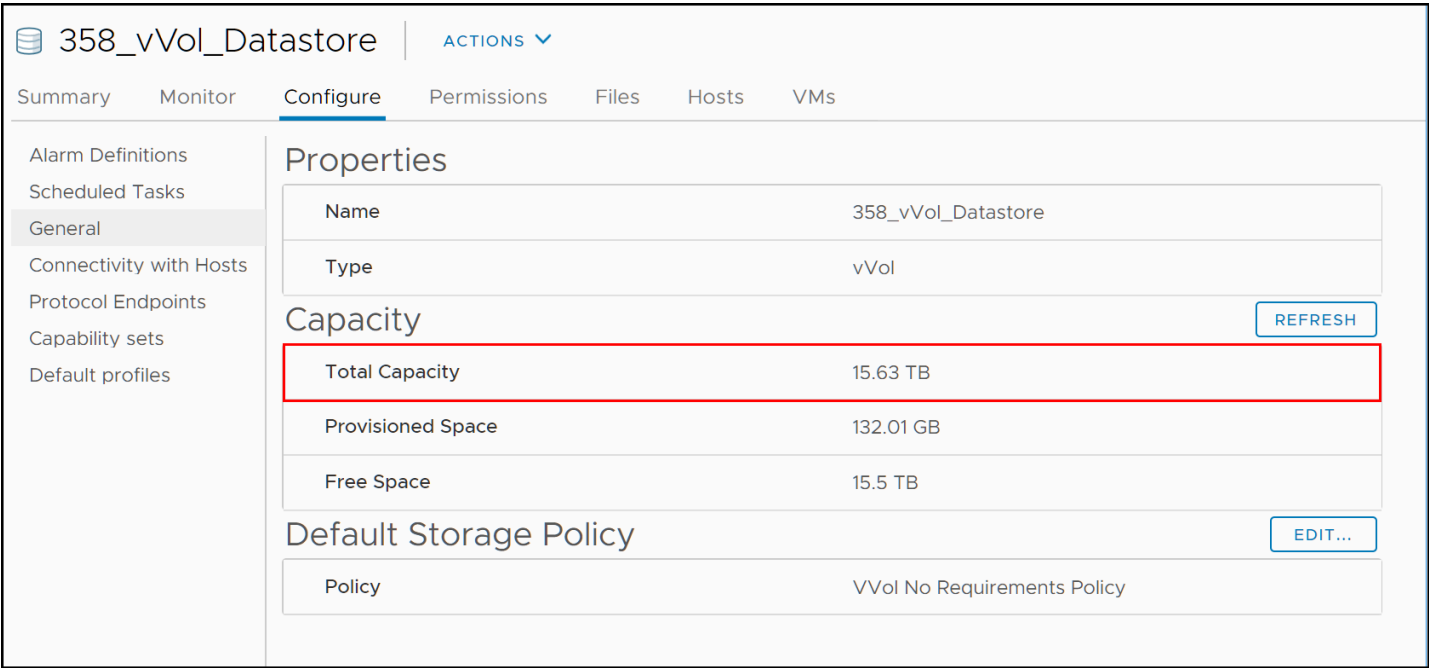


Figure 94. Total capacity includes newly created storage resource

10.3 Out of space errors

If space is exhausted in a storage resource and an activity is attempted in vCenter that requires additional storage, VMware will report that the datastore itself has run out of space, even if the vVol datastore has space remaining. For instance, in [Figure 95](#) there is an attempt to take a VM snapshot but it fails with “Insufficient disk space on datastore”.

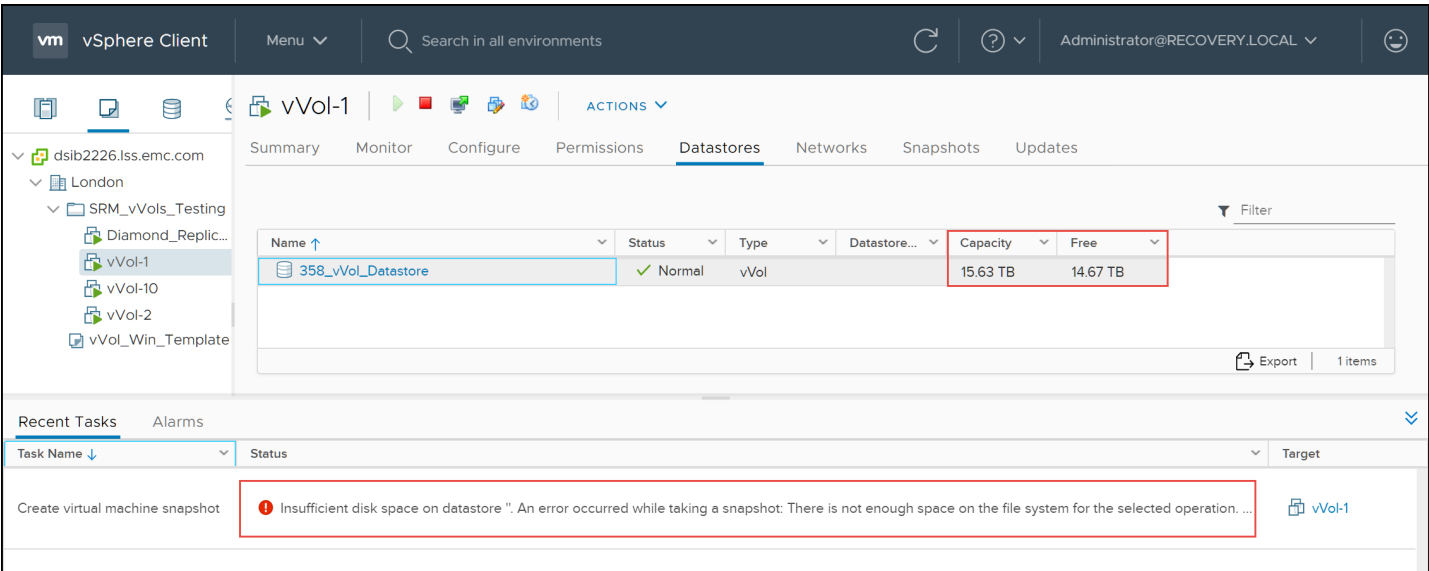
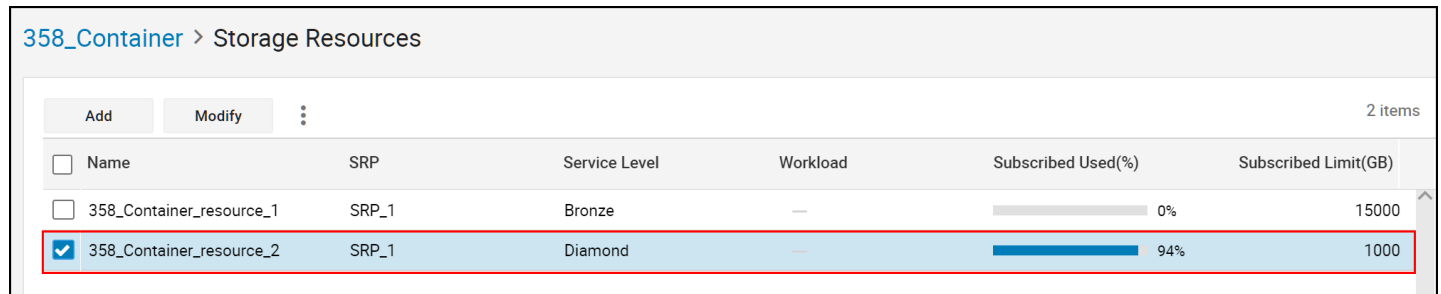


Figure 95. Total capacity includes newly created storage resource

However, there is 14.67 TB free in the datastore so it seems counterintuitive that the task would fail. The reason is that the datastore is comprised of storage resources, so if the resource in which the

VM resides is full, the task will fail. In this example, the Diamond storage resource is exhausted in [Figure 96](#)



358_Container > Storage Resources

2 items

<input type="checkbox"/>	Name	SRP	Service Level	Workload	Subscribed Used(%)	Subscribed Limit(GB)
<input type="checkbox"/>	358_Container_resource_1	SRP_1	Bronze	—	0%	15000
<input checked="" type="checkbox"/>	358_Container_resource_2	SRP_1	Diamond	—	94%	1000

[Figure 96](#). Total capacity includes newly created storage resource

VMware is unaware of the storage resources; it simply knows the VASA Provider rejected the call due to lack of space. Once space is added to the resource, the task will complete. Be sure to refresh the datastore capacity before re-attempting. Alternatively, the VM could be assigned a different storage policy to move it to the Bronze service level where space already exists.

11 Creating a VM Storage Policy for vVols

VMware utilizes Storage Policy Based Management in conjunction with vVols. The PowerMax advertises its capabilities to vSphere. The user creates policies that map to those capabilities so that when the user provisions a VM, a policy can be selected that will filter the datastores so that the appropriate one is selected. The capabilities are passed to the VASA Provider so that the vVols are created with the proper Service Level (SL) and, optionally, replicated. The PowerMax advertises six SLs to vSphere shown in [Figure 97](#):

- Diamond
- Platinum
- Gold
- Silver
- Bronze
- Optimized

An SL has a pre-defined target response time that the PowerMax attempts to deliver for the vVol device. There is also a defined Compliance Range as the response time will fluctuate. Be aware that adding replication to the VM Storage Policy will generally increase the pre-defined target response time by .2 ms.




Figure 97. PowerMax storage capabilities in Unisphere for PowerMax

11.1 VM Storage Policy with SRDF Replication

The following provides a step-by-step process for creating a VM storage policy with replication. A storage policy is not required in order to use vVol storage, however any VMs created in a vVol datastore without using a storage policy will be assigned to the default storage policy of **vVol No Requirements Policy** and assigned the SL with the highest response time targets and will not be replicated. The default storage policy can be changed, however, by following the section Default Storage Policy.

Note: As the Optimized SL has no response time target, it is considered to be the lowest SL if it is available in the storage container; however, this does not hold true for the swap file. Unless Optimized is the only SL available, the next lowest SL will be used for swap.

11.1.1 VM Storage Policy wizard

Start by accessing the **VM Storage Policies** icon in the Home page of the vSphere Client as shown in [Figure 98](#).

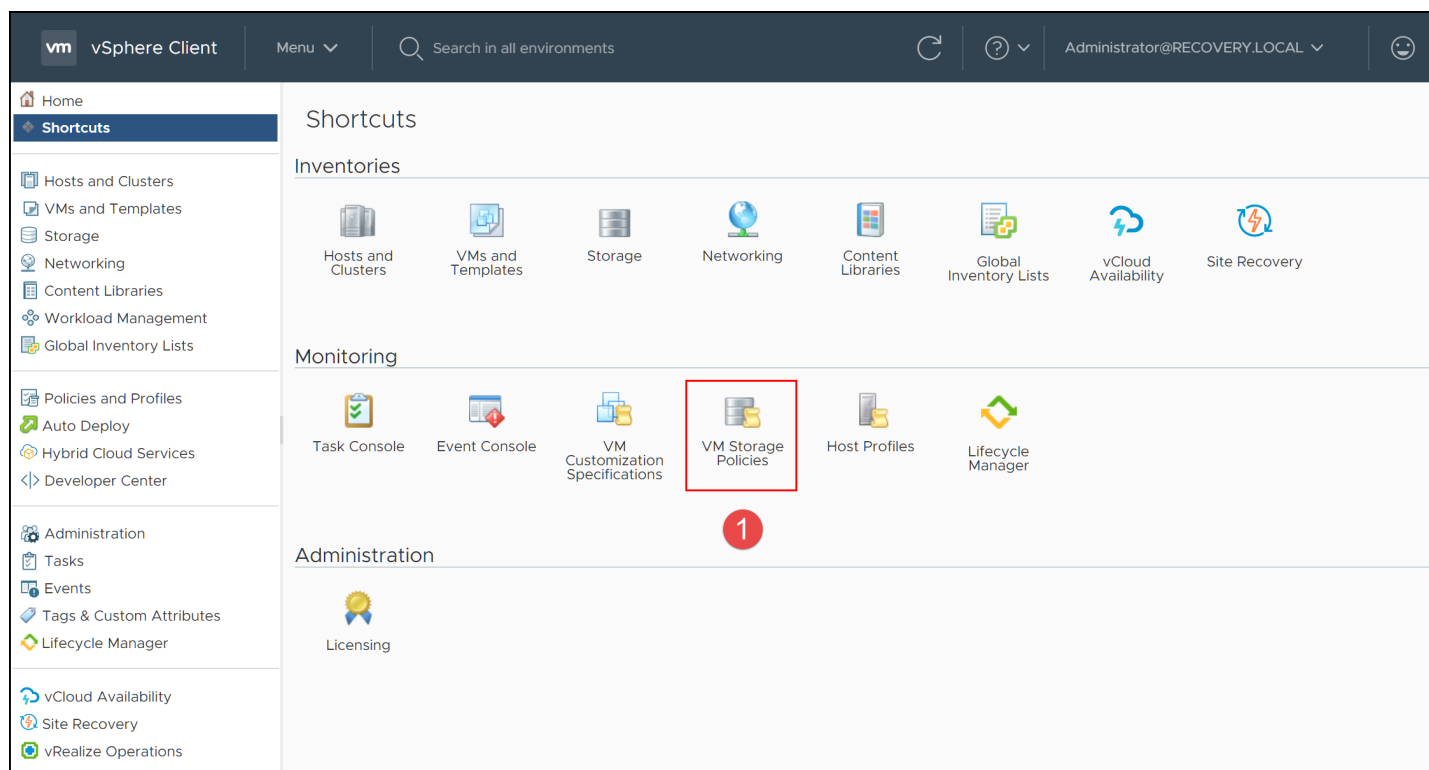


Figure 98. Creating a VM Storage Policy - step 1

Next in step 2 in [Figure 99](#) select the icon **Create VM Storage Policy** to create a new VM storage policy.

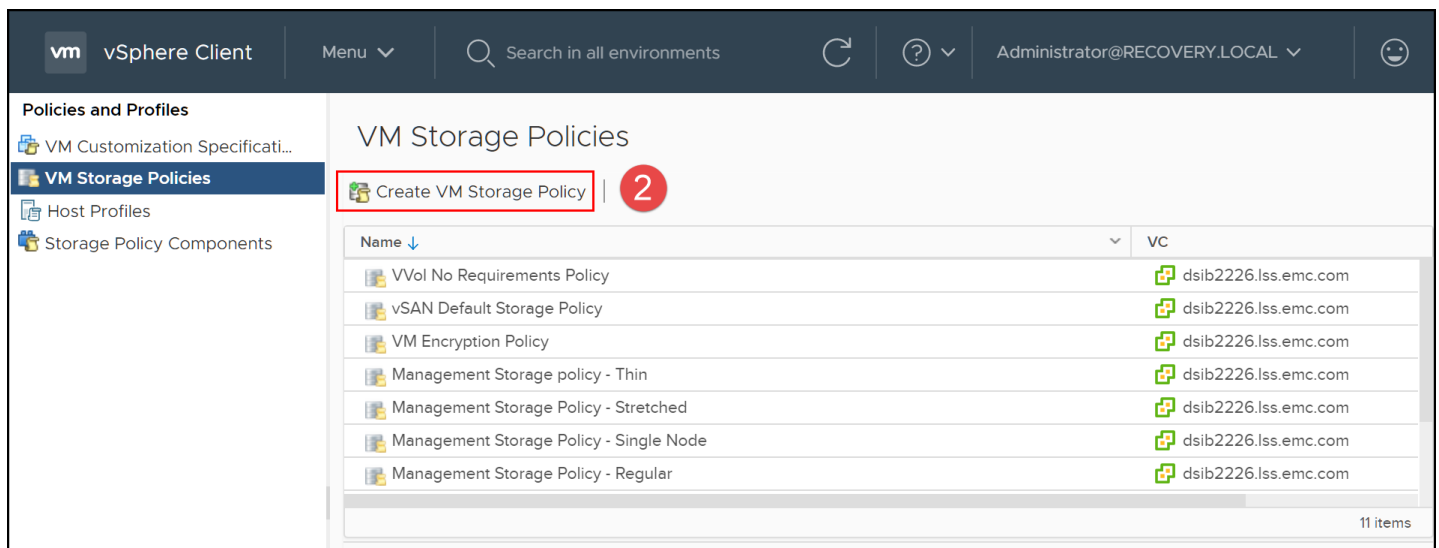


Figure 99. Creating a VM Storage Policy - step 2

Step 3 (step 1 of the wizard) formally starts the wizard. If the environment shares a Platform Service Controller, then begin by selecting the appropriate vCenter. Enter a name for the policy, preferably one that reflects the capabilities that will be associated with the policy as this is the name the VMware user will see. Finally, if desired enter a description. An example is shown in Figure 100.

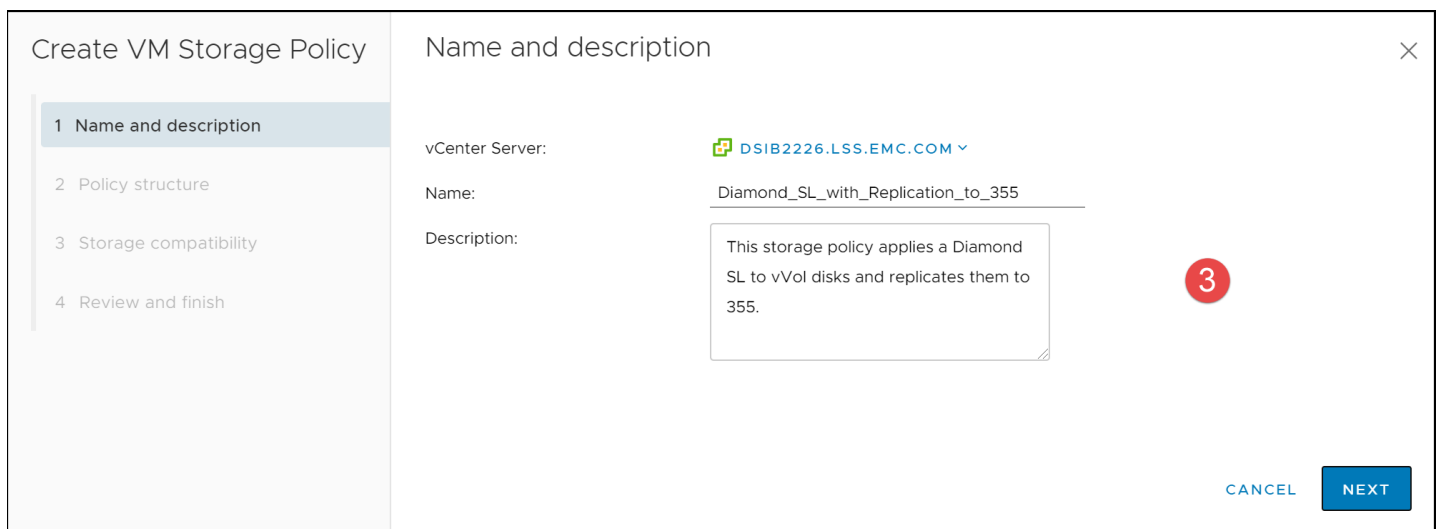


Figure 100. Creating a VM Storage Policy - step 3

Step 4 is the Policy structure definition. It covers adding host-based rules, like Storage IO Control, as well as assigning the Rule Set.

The Dell VASA Provider advertises a single set of rules for VP:

- VmaxVVolProvider

Select the **VmaxVVolProvider** as demonstrated in Figure 101. Note depending on the vSphere release, there may be additional rules listed.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 VmaxVVolProvider rules

4 Storage compatibility

5 Review and finish

Policy structure

Host based services

Create rules for data services provided by hosts. Available data services could include encryption, I/O control, caching, etc. Host based services will be applied in addition to any datastore specific rules.

☐ Enable host based rules

Datastore specific rules

Create rules for a specific storage type to configure data services provided by the datastores. The rules will be applied when VMs are placed on the specific storage type.

☐ Enable rules for "vSAN" storage

☒ Enable rules for "VmaxVVolProvider" storage

☐ Enable tag based placement rules

CANCEL BACK NEXT

Figure 101. Creating a VM Storage Policy - step 4

Once the data service is selected, the advertised capabilities may be added as rules for the policy. Through the VASA Provider the PowerMax presents the Service Level Objective as the capability. It is comprised of two rules, from a VMware perspective: Performance Index and Workload Hints. Start by selecting **Performance Index** as in Figure 102. The other rule, Workload Hints, is not used with PowerMax.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 VmaxVVolProvider rules

4 Storage compatibility

5 Review and finish

VmaxVVolProvider rules

Placement Replication Tags

ADD RULE

Performance Index

Workload Hints

CANCEL BACK NEXT

Figure 102. Creating a VM Storage Policy - step 5

In step 6, Figure 103, one can see that the Performance Index for PowerMax will be translated into **"Service Level"** and all the SLs available on the box will be shown in a drop-down list. Select an SL that maps to a storage resource available in the storage container on the PowerMax. Note that all service levels are advertised here. In this step VMware has no idea what each vVol datastore can support.

Figure 103. Creating a VM Storage Policy - step 6

Once the service level is chosen, select the **Replication** tab to view the options available for add this feature to the storage policy.

Figure 104. Creating a VM Storage Policy - step 7

In the Replication tab choose the “Custom” radio button. If PowerMax is the only VASA Provider registered, all fields will default. The only field which can be modified at this time is the “TargetFaultDomain”. If there is more than one remote array configured with a VASA Replication Group, it will be available in the drop-down of the TargetFaultDomain. [Figure 105](#) has the details for this target array.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 VmaxVVolProvider rules

4 Storage compatibility

5 Review and finish

VmaxVVolProvider rules

PlacementReplicationTags

Disabled

Use storage policy component<Select component>

Custom8

Provider:VmaxVVolProvider.RemoteReplication

DELLEMC PowerMax VVol Remote Replication Capabilities

Replication TypeAsynchronous

TargetFaultDomainPMAX_000197600355

Recovery Point Objective(RPO)300 seconds

CANCEL

BACK

NEXT

Figure 105. Creating a VM Storage Policy - step 8

VMware now takes the supplied parameters and compares it against the available vVol datastores to see if any are compatible. In Figure 106 the 358_vVol_Datastore is compatible with the Diamond SL and replication. Note that when replication is part of a storage policy, there must be a replication group going from the local storage container to the remote storage container or the vVol datastore associated with that storage container will not show as compatible. When creating the storage policy on the protection vCenter, the vVol datastore should always show as compatible; however, when creating the storage policy on the recovery vCenter, there may be no vVol datastores as compatible. This is not a concern because if a failover is run, the replication group will switch identity and the vVol datastore on the recovery vCenter will be compatible.

Create VM Storage Policy

1 Name and description

2 Policy structure

3 VmaxVVolProvider rules

4 Storage compatibility

5 Review and finish

Storage compatibility

Compatible storage 15.63 TB (15.5 TB free)Compatible

Expand datastore clusters

Name	Datacenter	Type	Free Space	Capacity	Warnings
358_vVol_Datastore	London	vVol	15.5 TB	15.63 TB	

9

CANCEL

BACK

NEXT

Figure 106. Creating a VM Storage Policy - step 9

A summary page in Figure 107 complete the VM Storage Policy.

95

Using VMware vSphere Virtual Volumes and VASA with Dell PowerMax | h19812

DELLTechnologies

Create VM Storage Policy

- 1 Name and description
- 2 Policy structure
- 3 VmaxVVolProvider rules
- 4 Storage compatibility
- 5 Review and finish

Review and finish

General

Name: Diamond_SL_with_Replication_to_355

Description: This storage policy applies a Diamond SL to vVol disks and replicates them to 355.

vCenter Server: dsib2226.lss.emc.com

VmaxVVolProvider rules

Placement: "Service Level" Diamond

Replication: Provider VmaxVVolProvider.RemoteReplication

DELLEMC PowerMax VVol Remote Replication Capabilities

Replication Type: Asynchronous

TargetFaultDomain: PMAX_000197600355

Recovery Point Objective(RPO): 300 seconds

CANCEL BACK FINISH

Figure 107. Creating a VM Storage Policy - step 10

11.2 Storage policy component

Another option available in [Figure 108](#), is to use a storage policy component for replication, rather than doing a custom application each time. The benefit of the storage policy component is that a single component for replication to a fault domain can be created once and then applied each time to a storage policy. To create the storage policy component, start by navigating to the VM Storage Policies as in [Figure 98](#). Creating a VM Storage Policy - step 1. Next, select from the left-hand menu, Storage Policy Components and then Create Storage Policy Component in [Figure 108](#).

vm vSphere Client

Menu

Search in all environments

Administrator@PROTECTION.LOCAL

Policies and Profiles

VM Customization Specifications

VM Storage Policies

Host Profiles

Storage Policy Components

Storage Policy Components

+ Create Storage Policy Component

Name	Description	Category	VC
Default encryption properties	Storage policy component for VM and...	Encryption	dsib2224.lss.emc.c
High IO shares allocation	Storage policy component for High SI...	Storage I/O Control	dsib2224.lss.emc.c
Low IO shares allocation	Storage policy component for Low SI...	Storage I/O Control	dsib2224.lss.emc.c
Normal IO shares allocation	Storage policy component for Mediu...	Storage I/O Control	dsib2224.lss.emc.c
Replication_to_358_SP_Component	This storage policy component will co...	Replication	dsib2224.lss.emc.c

Figure 108. Creating a Storage Policy Component - step 2

In step 3, provide a Name and set the Category to Replication which will automatically fill in the other fields. If there is more than one fault domain, use the drop-down to select the desired array. The final result is seen in [Figure 109](#).

New Storage Policy Component

3

vCenter Server:

DSIB2224.LSS.EMC.COM

Name:

Replication_to_450

Description:

This storage policy component will add replication to array 450 to a storage policy.

Category:

Replication

Provider:

VmaxVVolProvider.RemoteReplication

DELLEMC PowerMax VVol Remote Replication Capabilities

Replication Type

Asynchronous

TargetFaultDomain

PMAX_000197600450

Recovery Point Objective(RPO)

300 seconds

CANCEL

OK

Figure 109. Creating a Storage Policy Component - step 3

Once created, next time a storage policy is created, the component will be available for selection. An example is in [Figure 110](#).

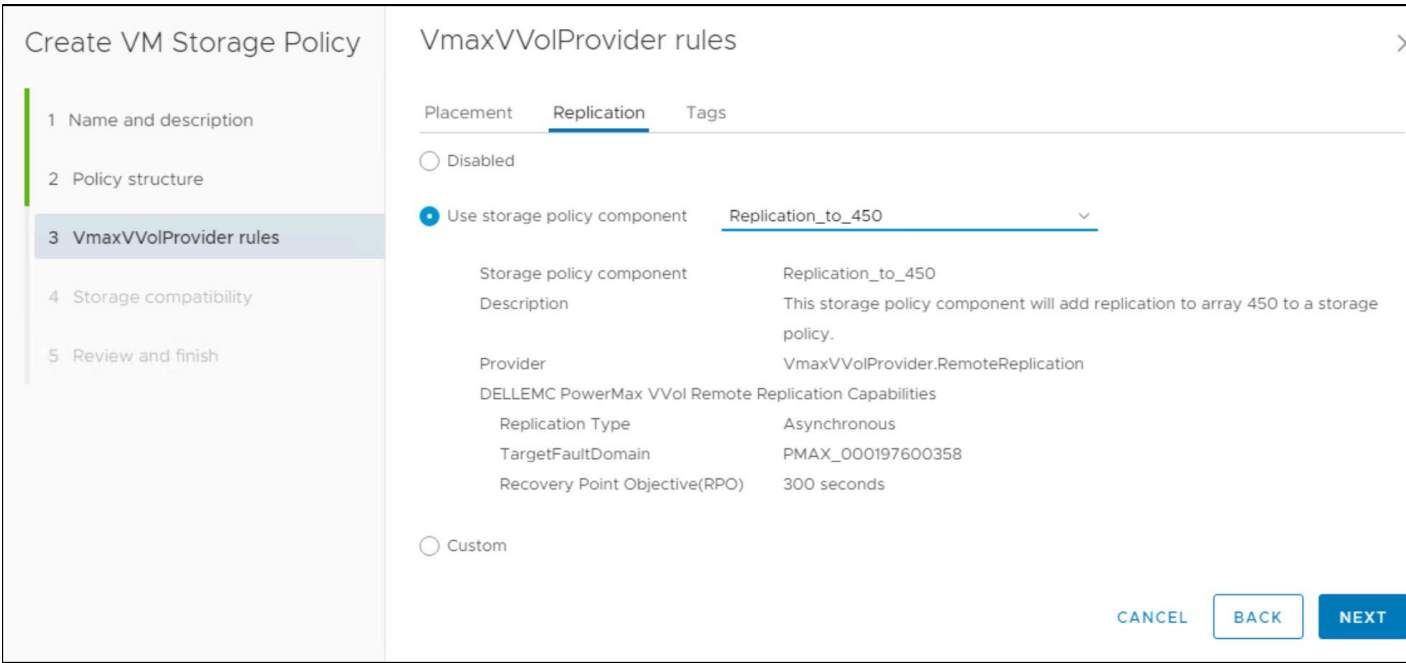


Figure 110. Using a Storage Policy Component

11.3 Creating a replicated VM with vVol storage

Creating a VM using vVol storage in the vSphere Client is no different than creating a VM with VMFS storage. As both types of storage are represented by datastores, it is simply a matter of selecting which datastore type the user wishes to use for the VM. It is essential for vVols, however, to utilize the VM storage policy to not only select the correct vVol datastore, but ensure the correct SL is sent with the creation command, along with a requirement to replicate it, if desired. Rather than walk through the entire creation of the VM, only step 4 from the wizard is included below to illustrate assigning the proper storage policy.

Note: The example below of creating a VM using vVol storage utilizes a single storage policy as it is using replication and only a single storage policy per VM is supported with VMware SRM. If a non-replicated VM has more than one vmdk, however, each vmdk could be assigned a different VM storage policy, and thus SL. For example, a user creating an application VM with two vmdks might select the VM storage policy for the Bronze SL for the OS vmdk, while assigning the Diamond SL to the application vmdk. If using replication, however, only a single storage policy is supported with VMware SRM, no matter how many vmdks the VM has.

11.3.1 VM Creation

In step 4 of the VM wizard, the storage is selected for the VM. By selecting a VM storage policy as in [Figure 111](#), the user can ensure that the vmdks that make up the VM will be assigned the desired SL in the vVol datastore, and if required replicated.

New Virtual Machine

- ✓ 1 Select a creation type
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- 4 Select storage**
- 5 Select compatibility
- 6 Select a guest OS
- 7 Customize hardware
- 8 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy:

Name	Capacity	Provisioned
Storage Compatibility: Compatible		
358_vVol_Datastore	15.63 TB	132.01 GB
Storage Compatibility: Incompatible		
dsib0184_local	150.75 GB	1.41 GB
dsib0186_local	150.75 GB	1.41 GB
HA_HEART_1	9.75 GB	13.33 GB
HA_HEART_2	9.75 GB	1.41 GB
vCENTER_358	2 TB	139.59 GB

Diamond_SL_with_Replication_to_355
Datastore Default
Bronze_358_with_Replication
Diamond_SL_with_Replication_to_355
Management Storage policy - Encryption
Management Storage Policy - Large
Management Storage Policy - Regular
Management Storage Policy - Single Node
Management Storage Policy - Stretched
Management Storage policy - Thin
VM Encryption Policy
vSAN Default Storage Policy
VVVol No Requirements Policy

Compatibility

CANCEL
BACK
NEXT

Figure 111. Creating VM (wizard step 4) – Select VM Storage Policy

For instance, in the example in Figure 112, the VM storage policy previously created “Diamond_SL_with_Replication_to_355” is selected. This storage policy has a diamond SL and will replicate any VM (vmdks) assigned to it. Once selected, VMware determines which vVol datastores are compatible with the policy. In this example, only one datastore, 358_vVol_Datastore, shows as compatible. When the user selects that datastore, another option immediately appears which indicates a source Replication Group (VASA) must be selected. In this example, there is only one VRG available, VP6.

New Virtual Machine

✓ 1 Select a creation type

✓ 2 Select a name and folder

✓ 3 Select a compute resource

4 Select storage

5 Select compatibility

6 Select a guest OS

7 Customize hardware

8 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy:

Diamond_SL_with_Replication_to_355

Name	Capacity	Provisioned	Free	Type	Cluster
Storage Compatibility: Compatible					
358_vVol_Datastore	15.63 TB	132.01 GB	15.5 TB	vVol	
Storage Compatibility: Incompatible					
dsib0184_local	150.75 GB	1.41 GB	149.34 GB	VMFS 6	
dsib0186_local	150.75 GB	1.41 GB	149.34 GB	VMFS 6	
HA_HEART_1	9.75 GB	13.33 GB	8.34 GB	VMFS 6	
HA_HEART_2	9.75 GB	1.41 GB	8.34 GB	VMFS 6	
vCENTER_358	2 TB	139.59 GB	1.86 TB	VMFS 6	

Replication Group:

<Select replication group>

Compatibility

VP6

ⓘ

 Select a valid replication group.

CANCEL

BACK

NEXT

Figure 112. Creating VM – Select datastore and assign Replication Group

The final result is shown in [Figure 113](#).

New Virtual Machine

✓ 1 Select a creation type

✓ 2 Select a name and folder

✓ 3 Select a compute resource

4 Select storage

5 Select compatibility

6 Select a guest OS

7 Customize hardware

8 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

VM Storage Policy:

Diamond_SL_with_Replication_to_355

Name	Capacity	Provisioned	Free	Type	Cluster
Storage Compatibility: Compatible					
358_vVol_Datastore	15.63 TB	132.01 GB	15.5 TB	vVol	
Storage Compatibility: Incompatible					
dsib0184_local	150.75 GB	1.41 GB	149.34 GB	VMFS 6	
dsib0186_local	150.75 GB	1.41 GB	149.34 GB	VMFS 6	
HA_HEART_1	9.75 GB	13.33 GB	8.34 GB	VMFS 6	
HA_HEART_2	9.75 GB	1.41 GB	8.34 GB	VMFS 6	
vCENTER_358	2 TB	139.59 GB	1.86 TB	VMFS 6	

Replication Group:

VP6

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

Figure 113. Creating VM – Complete

Note: When a Replication Group is in test mode, the name will be prefixed with “spbmReplicationGroups.intestState”. Do not choose this group until the test is complete.

Note: VMware does not support deploying an OVA/OVF template in a storage policy with replication enabled because the Replication Group cannot be selected in the wizard. The deployment will fail if the user proceeds. As a workaround, the VM can be deployed with a storage policy that does not include replication and adjusted after VM creation to a storage policy with replication.

Upon creation, a VM that is created in a vVol datastore will be comprised of two initial vVols – a *data* vVol for the Hard disk which is the size chosen (16 GB in this case), and a 4 GB *config* vVol for the configuration files (e.g., vmx, vmsd), represented as a directory. A third *swap* vVol is generated when the VM is powered on in the size of the VM memory (1 GB in this case). All vVols are present in [Figure 114](#).

```

10.228.246.17 - PuTTY
dsib2017:~ # symdev list -sid 358 -vvol

Symmetrix ID: 000197600358

-----
Device Name          Dir          Device
-----
Sym  Physical        SA :P   Config      Attribute    Sts    Cap
(MB)
-----
001D0 Not Visible    ??:??? VVOL        N/Grp'd     RW     4097
001D1 Not Visible    ??:??? VVOL        N/Grp'd     RW    16386
001D2 Not Visible    ??:??? VVOL        N/Grp'd     RW     1026

dsib2017:~ #

```

Figure 114. vVols for VM Diamond_Replication_VM

Figure 115 displays the contents of the VM “Diamond_Replication_VM” in the vVol datastore 358_vVol_Datastore. Note that despite the many files listed, there are only three vVols that are created on the PowerMax array. The config vVol is comprised of multiple files, namely the VM metadata, which accounts for the other files in the directory. Recall that a config vVol is represented by the Diamond_Replication_VM directory and is formatted with VMFS so is able to store these many files.

Name	Size	Modified	Type
.sdd.sf		06/01/2020, 2:46:21 PM	Folder
Diamond_Replication_VM-09f70577.vswp	0.29 KB	06/01/2020, 2:55:59 PM	File
Diamond_Replication_VM-09f70577.vswp.lck	0 KB	06/01/2020, 2:55:57 PM	File
Diamond_Replication_VM-5a8cfc2b.hlog	0.29 KB	06/01/2020, 2:46:25 PM	File
Diamond_Replication_VM.nvram	8.48 KB	06/01/2020, 2:56:32 PM	Non-volatile Memory Fi...
Diamond_Replication_VM.vmdk	16,777,216 KB	06/01/2020, 2:46:24 PM	Virtual Disk
Diamond_Replication_VM.vmsd	0 KB	06/01/2020, 2:46:25 PM	File
Diamond_Replication_VM.vmx	2.91 KB	06/01/2020, 2:56:01 PM	Virtual Machine
Diamond_Replication_VM.vmx.lck	0 KB	06/01/2020, 2:55:57 PM	File
Diamond_Replication_VM.vmx~	2.86 KB	06/01/2020, 2:56:01 PM	File
vmware.log	112.19 KB	06/01/2020, 2:57:41 PM	VM Log File
vmx-Diamond_Replication_-1cb92de776268d90dd...	81,920 KB	06/01/2020, 2:55:56 PM	File

Figure 115. VM contents in a vVol datastore

11.3.1.1 Replication

As replication was part of the storage policy, the vVols that make up the VM are now in SRDF/A pairs with the remote array. Note that the swap vVol is not replicated since a new vVol will be generated at the remote site if a failover is run. The pairs cannot be seen in Unisphere for PowerMax; however, they are available in Solutions Enabler in Figure 116.

```

10.228.246.17 - PuTTY
dsib2017:~ # symrdf list -sid 358 -all

Symmetrix ID: 000197600358

-----
Local Device View
-----
Sym   Sym   RDF   STATUS   FLAGS   R1 Inv   R2 Inv   RDF   S T A T E S
Dev   RDev  Typ:G  SA RA LNK MTES  Tracks  Tracks  Dev RDev Pair
-----
001D0 000E0  NA:6   ?? RW RW  A1.E      0      0 RW  WD  Consistent
001D1 000E3  NA:6   ?? RW RW  A1.E      0      0 RW  WD  Consistent
-----

Total
Track(s)      0      0
MB(s)        0.0    0.0

Legend for FLAGS:

(M)ode of Operation : A = Async, S = Sync, E = Semi-sync, D = Adaptive Copy Disk Mode
                    : W = Adaptive Copy WP Mode, M = Mixed, T = Active
Mirror (T)ype       : 1 = R1, 2 = R2
(E)xempt            : X = Enabled, . = Disabled, M = Mixed, - = N/A
R1/R2 Device (S)ize : E = R1 EQ R2, 1 = R1 GT R2, 2 = R2 GT R1, - = N/A

dsib2017:~ #

```

Figure 116. SRDF/A pairs for replication VM

11.3.1.2 Caveats

vVol SRDF pairs do not have an R1/R2 designation under the “RDF Typ:G” column, yet the mirror type flag will still identify the R1 from the R2. There are some other important points about vVol replication to keep in mind:

- A VM can be replicated even if there is no mounted vVol datastore at the remote array. In fact, a remote VMware infrastructure does not even have to exist at the time of creation. It can be built at a later time to facilitate DR.
- If a remote vVol datastore is associated with the container in the VASA Replication Group, it will not contain any visible files for the remote vVols. Until such time as the remote vVols are used in a testfailover or failover operation with SRM, they will not be visible.
- During the initial creation of a replicated VM, vCenter will report that the VM is not in compliance until synchronization of the SRDF pairs completes as seen in [Figure 117](#). When the pairs are in a “Consistent” state, the VM will move to “Compliant” in [Figure 118](#). The time it takes for VM to change compliancy, is directly related to how much data must be synchronized. Using the previously supplied Solutions Enabler command in [Figure 116](#), that progress can be monitored.
 - When new or migrated VMs are added to a replication group, all VMs in the VRG become Noncompliant. This is because the group is inconsistent while the new SRDF pairs are synchronizing.

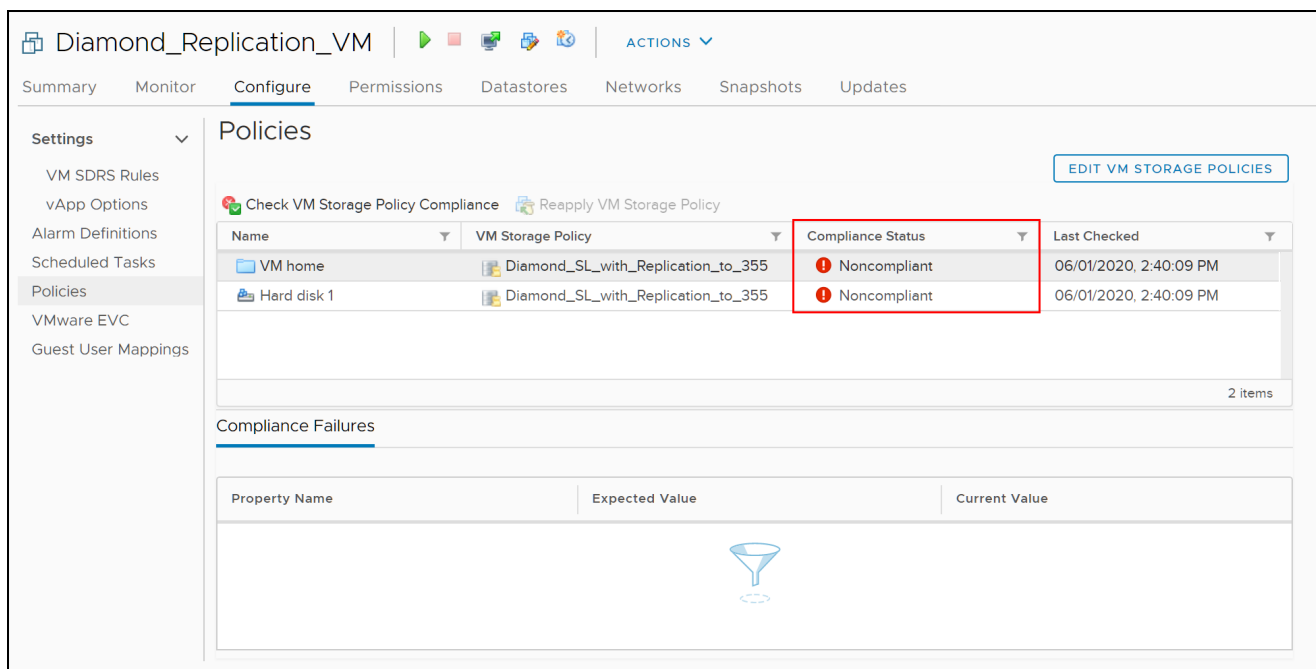


Figure 117. A replicated VM in Noncompliant status

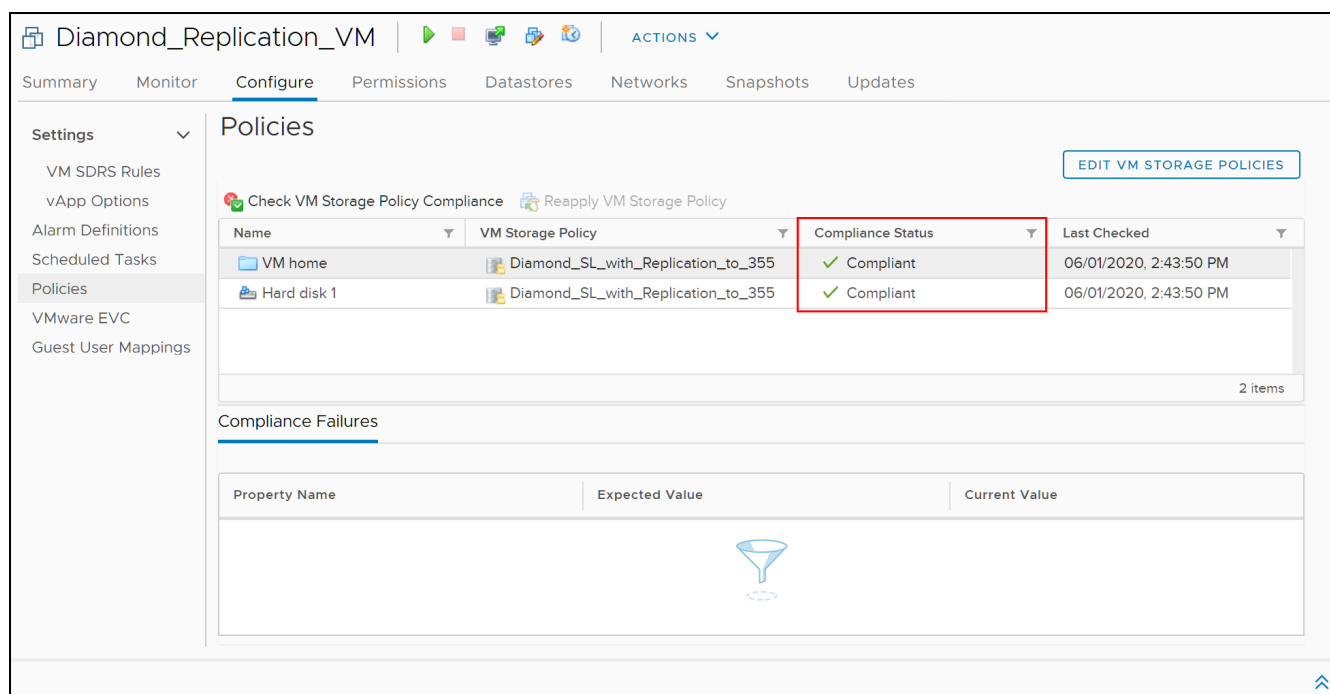


Figure 118. A replicated VM in Compliant status

- Because the underlying replication mode with vVols is SRDF/A, all vVols in a VASA Replication Group are failed over together. Partial pair failover is not possible. The only way to enable a VM to failover independently of other VMs is to put it in a separate VRG. Therefore, Dell recommends using a single VRG for each application, much like they recommend one application per storage group for traditional devices.
- Replicated vVols cannot be resized.

11.3.1.3 VMware disk type

By default, all vVols are created with a thin disk type. The vmdks will not grow unless data is actually written to them. In fact, even if the disk type is changed to “thick”, or an existing vmdk is inflated, the size will not change.

In a traditional VMFS environment, the default disk type is zeroedthick (aka lazy). A zeroedthick disk reserves the space in the VMFS datastore, though it does not actually allocate the space on the array, an important distinction. By contrast, a thin disk type in VMFS does not reserve any space on the datastore and only grows as data is written to it. This can be dangerous in a VMFS environment because a datastore could fill before a thin vmdk is full (due to other vmdks in the datastore).

In the PowerMax implementation of vVols, a vVol datastore behaves differently when it comes to thin vmdks. A thin vVol vmdk reserves space in the vVol datastore and storage resource/storage container on the array. This ensures the VM will not run out of space as it is used. The reason it works this way is because a vVol is not just a vmdk stored in a large datastore represented by a single device on the array. Every vVol is its own device (TDEV). If that thin vVol vmdk was treated as truly thin, the backing device on the array would need constant resizing. That would be an incredibly expensive process in terms of performance. Instead, the device is created with the requested size right away. It can be resized if non-replicated, of course, but the entire space of the vmdk is ready to be written to immediately. It is critical to remember, however, that storage resources do not reserve space in the SRP. So, while the space is reserved in the vVol datastore, if there is no available space in the SRP, the vVol will fail to allocate a new extent. In that case your entire array is out of space since most boxes have a single SRP.

Note: Since vVols do not support eagerzeroedthick disks, if the multi-writer flag is required for shared vmdks (e.g., Oracle RAC), it is set on the default thin devices.

11.4 Changing VM Storage Policy for a VM

During the course of regular business operations, the performance requirements of an application may change. Before virtual volumes, if a change in SL was needed it could only be done at the VMFS datastore level which impacted all VMs on that datastore. Alternatively, the VM could also be migrated off the source datastore to one with the required SL, but that meant a potential decrease in performance of the VM during the move. With virtual volumes a single VM can now be assigned a new SL without impacting any other VMs in the environment. In fact, a single vmdk of that VM could be assigned a new SL. There is a very simple process to do this.

Note: Changing the storage policy of a VM or VM disk within the same vVol datastore does not initiate a Storage vMotion. The array reassigns the vVols on the backend as required; however, if a VM is moved between vVol datastores, regardless of if those datastores are on the same array, a Storage vMotion is required.

The following walks through changing the policy for a vmdk (Hard disk).

11.4.1 Change Storage Policy wizard

Access the VM from the left-hand menu. Select the **Configure** tab on the right and the **Policies** menu on the left. Each Hard disk is listed in this panel. In the right-hand corner start by selecting the **EDIT VM STORAGE POLICIES** button as in [Figure 119](#).

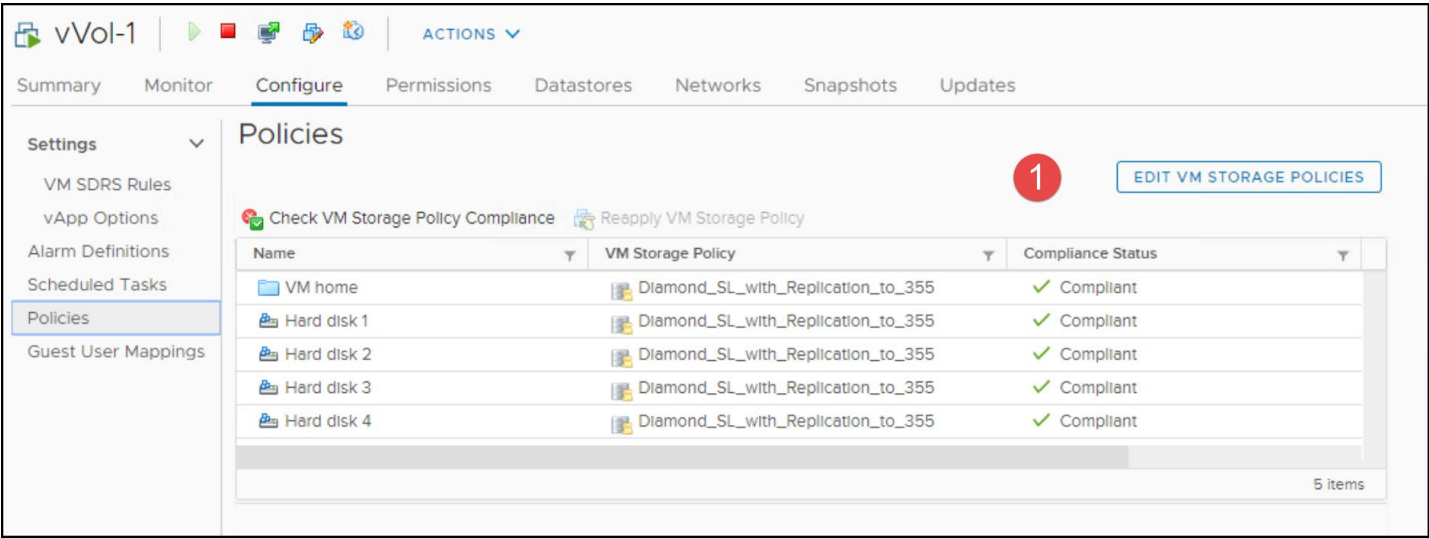


Figure 119. Changing a Storage Policy for VM - step 1

When changing the storage policy, there are two options available. By default, changing the VM storage policy will apply to all disks. Simply use the drop-down box highlighted in the red box in step 2 in Figure 120, select the new policy, and select OK. The second option is available in step 3 by toggling the option in the right-hand corner, **Configure per disk**. If using VMware Site Recovery Manager (SRM), do not select this option since SRM does not support using multiple storage policies in the same VM.

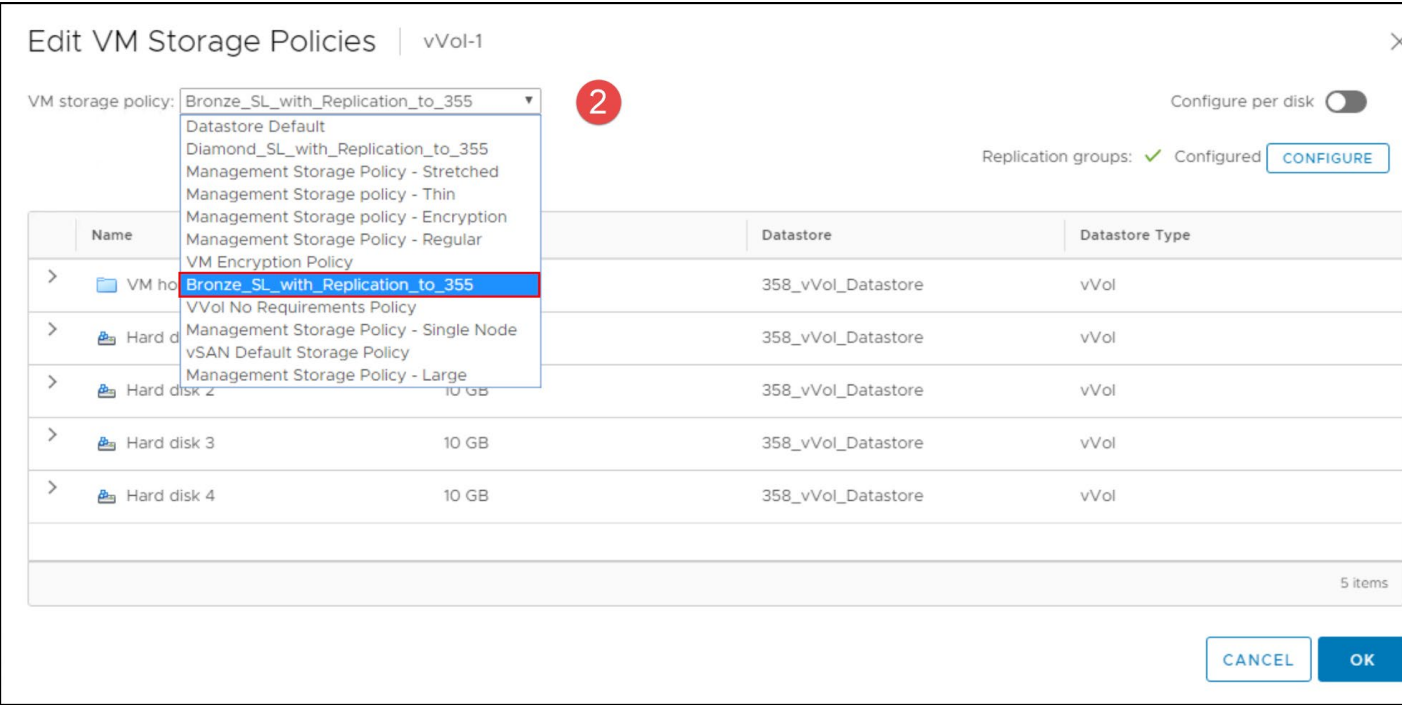


Figure 120. Changing a Storage Policy for VM - step 2

VMware makes the appropriate calls to the VASA Provider which calls to the array to move the vVol to the appropriate SL. The new storage policy for the VM is seen in Figure 121.

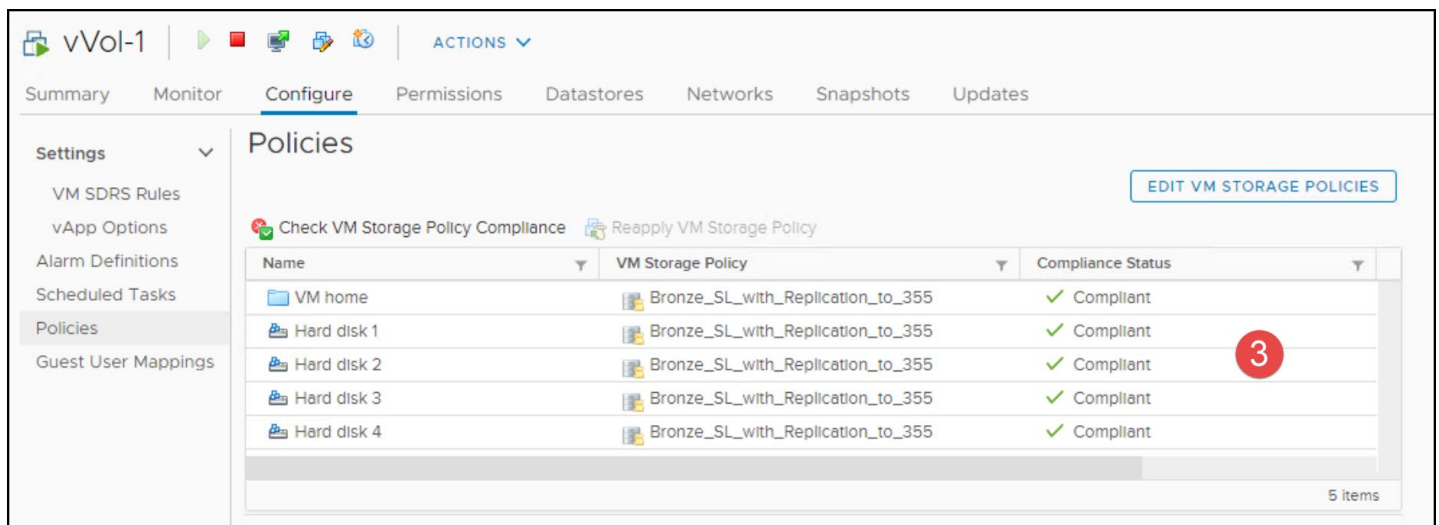


Figure 121. Changing a Storage Policy for VM - step 5

11.5 VMware CLI for vVols

VMware offers some ability to view the setup of the PowerMax vVols on a host. The command, `esxcli storage vvol`, pictured in Figure 122, has five available namespace commands prior to 8.0.1 but has expanded in 8.0.2+ in Figure 123.

```

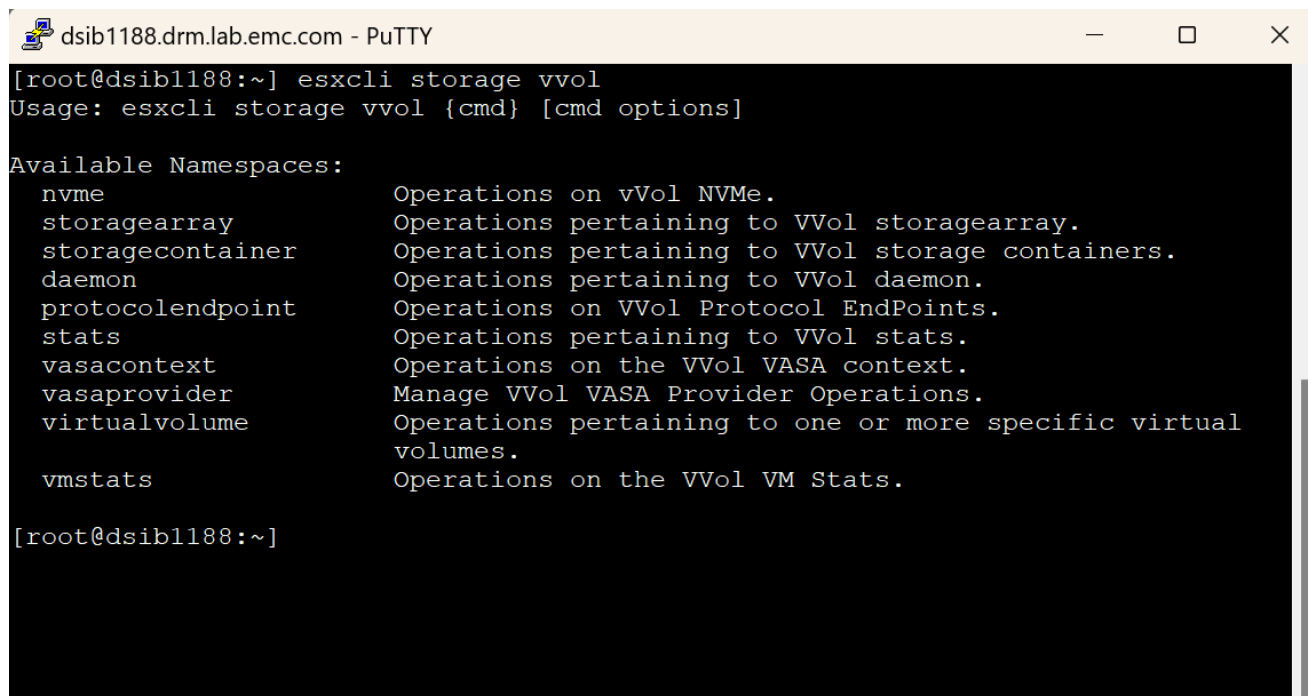
10.228.244.184 - PuTTY
[root@dsib0184:~] esxcli storage vvol
Usage: esxcli storage vvol {cmd} [cmd options]

Available Namespaces:
  storagecontainer  Operations pertaining to VVol storage containers.
  daemon           Operations pertaining to VVol daemon.
  protocolendpoint  Operations on VVol Protocol EndPoints.
  vasacontext       Operations on the VVol VASA context.
  vasaprovider      Manage VVol VASA Provider Operations.

[root@dsib0184:~]

```

Figure 122. VMware esxcli options for displaying vVol information up to 8.0.1



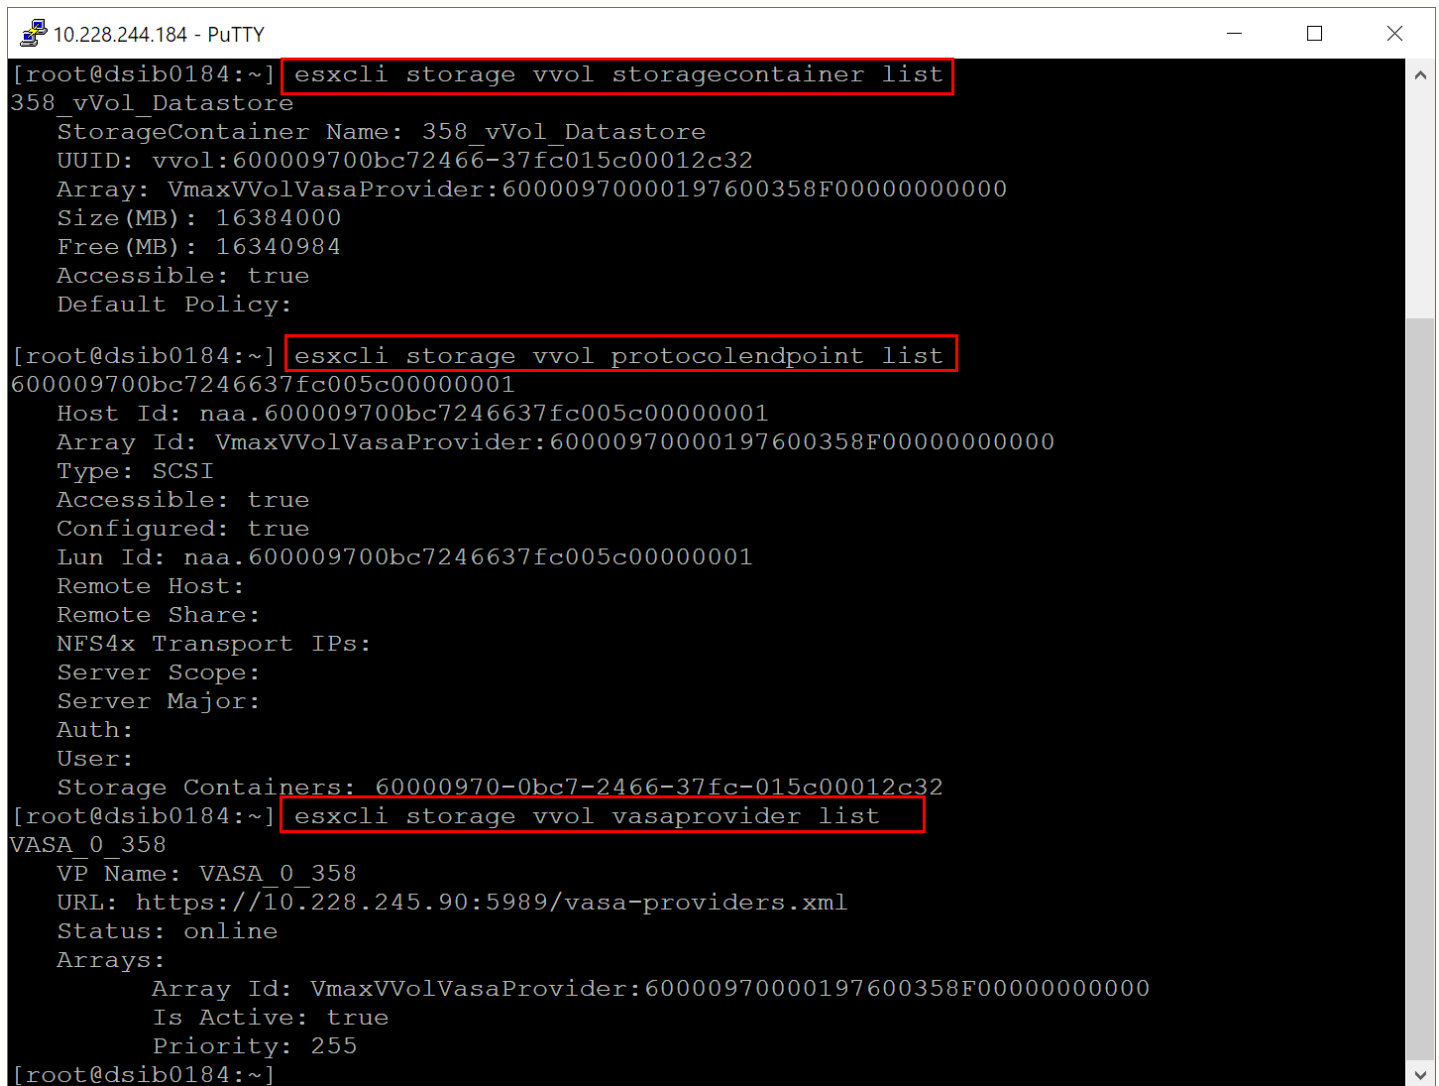
```
dsib1188.drm.lab.emc.com - PuTTY
[root@dsib1188:~] esxcli storage vvol
Usage: esxcli storage vvol {cmd} [cmd options]

Available Namespaces:
  nvme                Operations on vVol NVMe.
  storagearray        Operations pertaining to VVol storagearray.
  storagecontainer    Operations pertaining to VVol storage containers.
  daemon              Operations pertaining to VVol daemon.
  protocolendpoint    Operations on VVol Protocol EndPoints.
  stats               Operations pertaining to VVol stats.
  vasacontext          Operations on the VVol VASA context.
  vasaprovider         Manage VVol VASA Provider Operations.
  virtualvolume        Operations pertaining to one or more specific virtual
                       volumes.
  vmstats              Operations on the VVol VM Stats.

[root@dsib1188:~]
```

Figure 123. VMware esxcli options for displaying vVol information in 8.0.2+

Most of these commands are simply to list objects such as containers, protocol endpoints, or even the VASA Provider; however, the *vasacontext* command gets the vCenter UUID, while the *daemon* command can cause disruption as it unbinds all virtual volumes from the known VASA Provider. Three of the most useful commands and their output are shown in [Figure 124](#).



```

10.228.244.184 - PuTTY
[root@dsib0184:~] esxcli storage vvol storagecontainer list
358_vVol_Datastore
StorageContainer Name: 358_vVol_Datastore
UUID: vvol:600009700bc72466-37fc015c00012c32
Array: VmaxVVolVasaProvider:60000970000197600358F00000000000
Size(MB): 16384000
Free(MB): 16340984
Accessible: true
Default Policy:

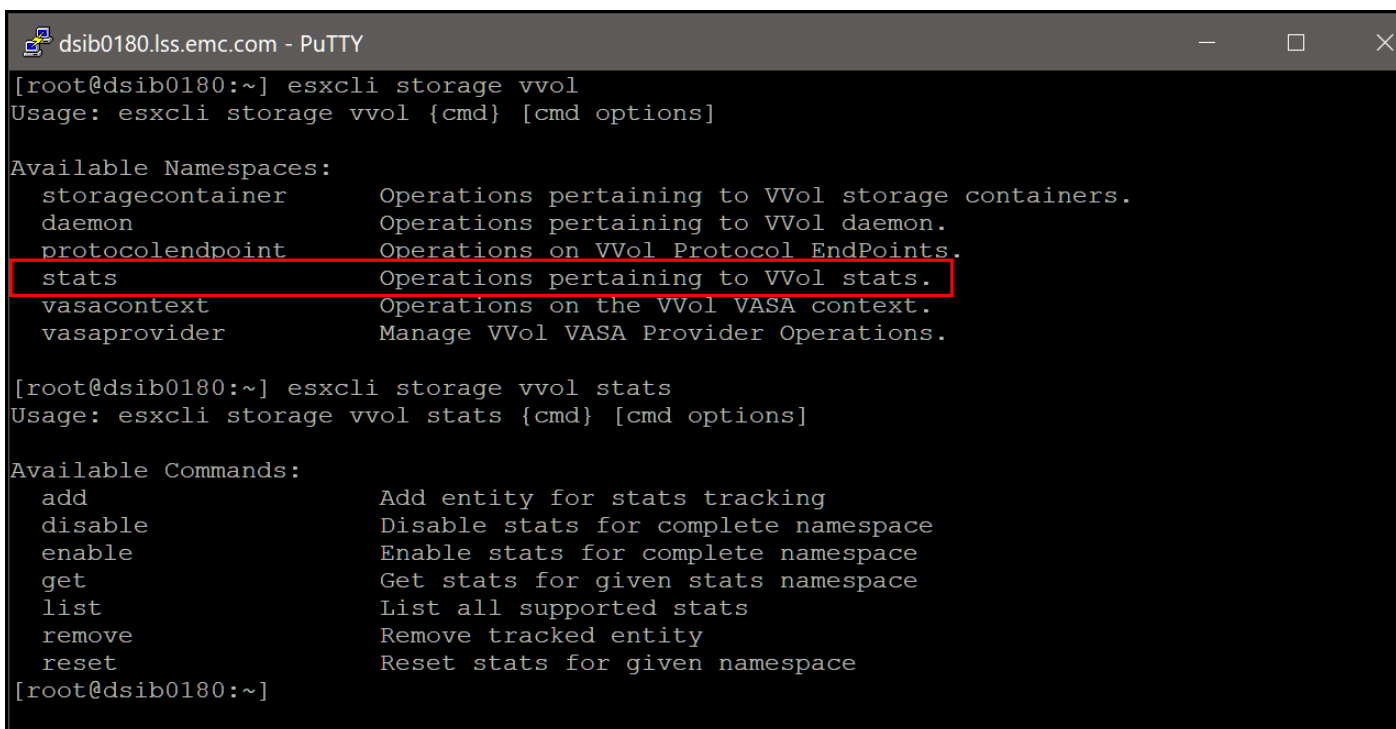
[root@dsib0184:~] esxcli storage vvol protocolendpoint list
600009700bc7246637fc005c00000001
Host Id: naa.600009700bc7246637fc005c00000001
Array Id: VmaxVVolVasaProvider:60000970000197600358F00000000000
Type: SCSI
Accessible: true
Configured: true
Lun Id: naa.600009700bc7246637fc005c00000001
Remote Host:
Remote Share:
NFS4x Transport IPs:
Server Scope:
Server Major:
Auth:
User:
Storage Containers: 60000970-0bc7-2466-37fc-015c00012c32
[root@dsib0184:~] esxcli storage vvol vasaprovider list
VASA_0_358
VP Name: VASA_0_358
URL: https://10.228.245.90:5989/vasa-providers.xml
Status: online
Arrays:
Array Id: VmaxVVolVasaProvider:60000970000197600358F00000000000
Is Active: true
Priority: 255
[root@dsib0184:~]

```

Figure 124. Namespace commands for vVol object listing

11.5.1 Stats

The namespace command, stats, is shown in [Figure 125](#).



```

dsib0180.lss.emc.com - PuTTY
[root@dsib0180:~] esxcli storage vvol
Usage: esxcli storage vvol {cmd} [cmd options]

Available Namespaces:
  storagecontainer  Operations pertaining to VVol storage containers.
  daemon            Operations pertaining to VVol daemon.
  protocolendpoint  Operations on VVol Protocol EndPoints.
  stats             Operations pertaining to VVol stats.
  vasacontext       Operations on the VVol VASA context.
  vasaprovider      Manage VVol VASA Provider Operations.

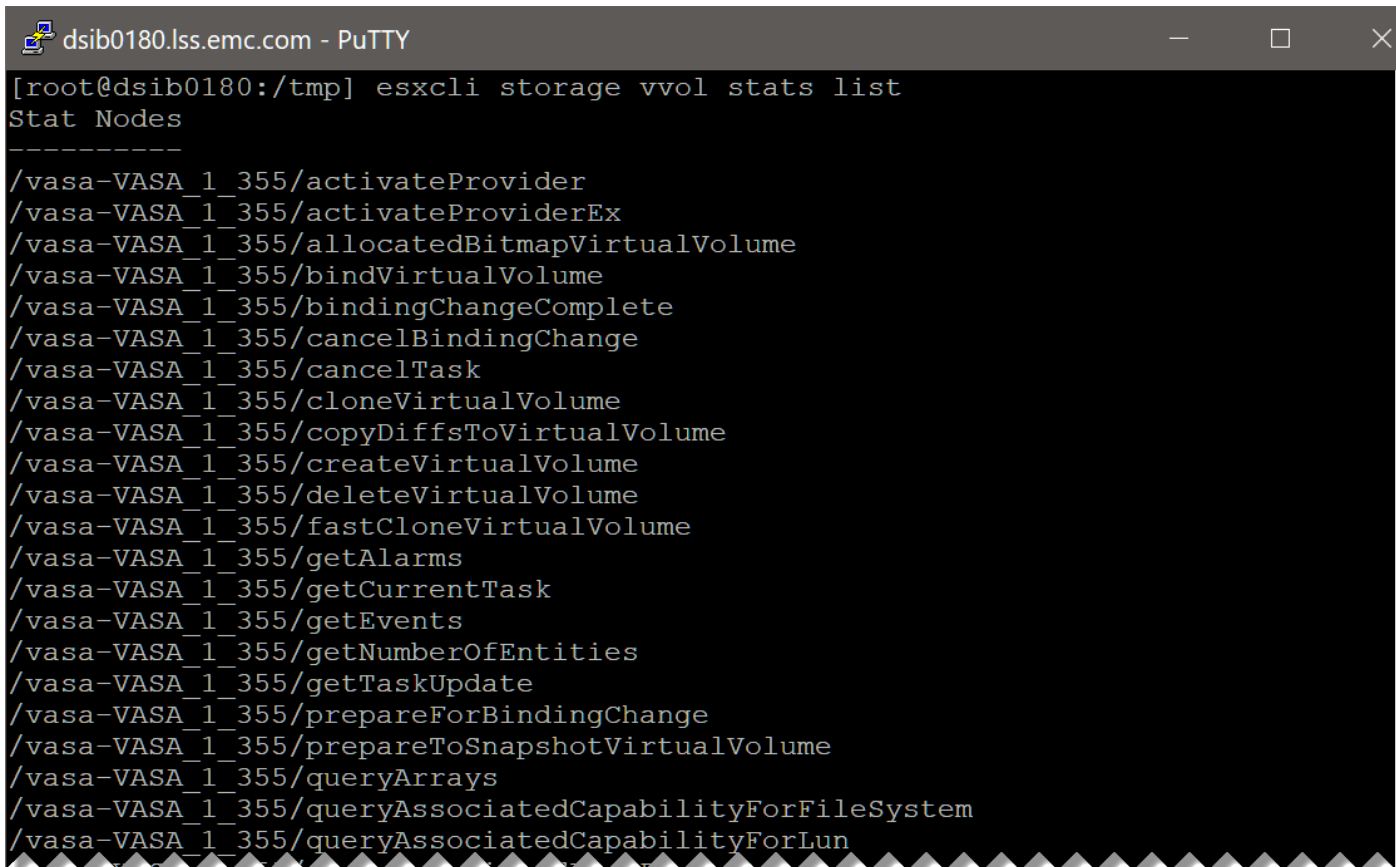
[root@dsib0180:~] esxcli storage vvol stats
Usage: esxcli storage vvol stats {cmd} [cmd options]

Available Commands:
  add      Add entity for stats tracking
  disable  Disable stats for complete namespace
  enable   Enable stats for complete namespace
  get      Get stats for given stats namespace
  list     List all supported stats
  remove   Remove tracked entity
  reset    Reset stats for given namespace
[root@dsib0180:~]

```

Figure 125. Stats namespace

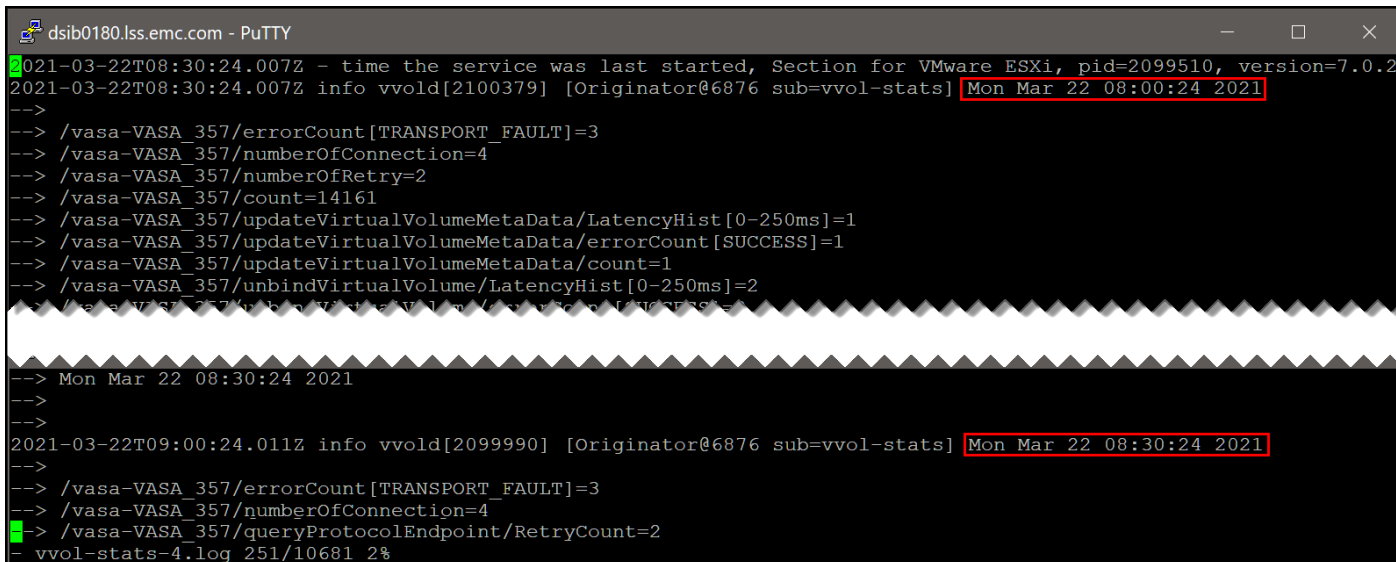
The command allows the user to monitor some performance stats, and, VMware suggests, identify issues with a slow VASA Provider. One can enable or disable all stats, or just a collection of them through the add/remove commands. One can get a list by running **esxcli storage vvol list**, in Figure 126, whether or not stats are enabled.



```
dsib0180.lss.emc.com - PuTTY
[root@dsib0180:/tmp] esxcli storage vvol stats list
Stat Nodes
-----
/vasa-VASA_1_355/activateProvider
/vasa-VASA_1_355/activateProviderEx
/vasa-VASA_1_355/allocatedBitmapVirtualVolume
/vasa-VASA_1_355/bindVirtualVolume
/vasa-VASA_1_355/bindingChangeComplete
/vasa-VASA_1_355/cancelBindingChange
/vasa-VASA_1_355/cancelTask
/vasa-VASA_1_355/cloneVirtualVolume
/vasa-VASA_1_355/copyDiffsToVirtualVolume
/vasa-VASA_1_355/createVirtualVolume
/vasa-VASA_1_355/deleteVirtualVolume
/vasa-VASA_1_355/fastCloneVirtualVolume
/vasa-VASA_1_355/getAlarms
/vasa-VASA_1_355/getCurrentTask
/vasa-VASA_1_355/getEvents
/vasa-VASA_1_355/getNumberOfEntities
/vasa-VASA_1_355/getTaskUpdate
/vasa-VASA_1_355/prepareForBindingChange
/vasa-VASA_1_355/prepareToSnapshotVirtualVolume
/vasa-VASA_1_355/queryArrays
/vasa-VASA_1_355/queryAssociatedCapabilityForFileSystem
/vasa-VASA_1_355/queryAssociatedCapabilityForLun
```

Figure 126. Listing vVol statistics with esxcli

If stats are enabled, VMware will create logs in the /var/log directory called **vvol-stats-*<number>*.log**. In these logs the stats are dumped every 30 minutes. An example is shown in Figure 127.

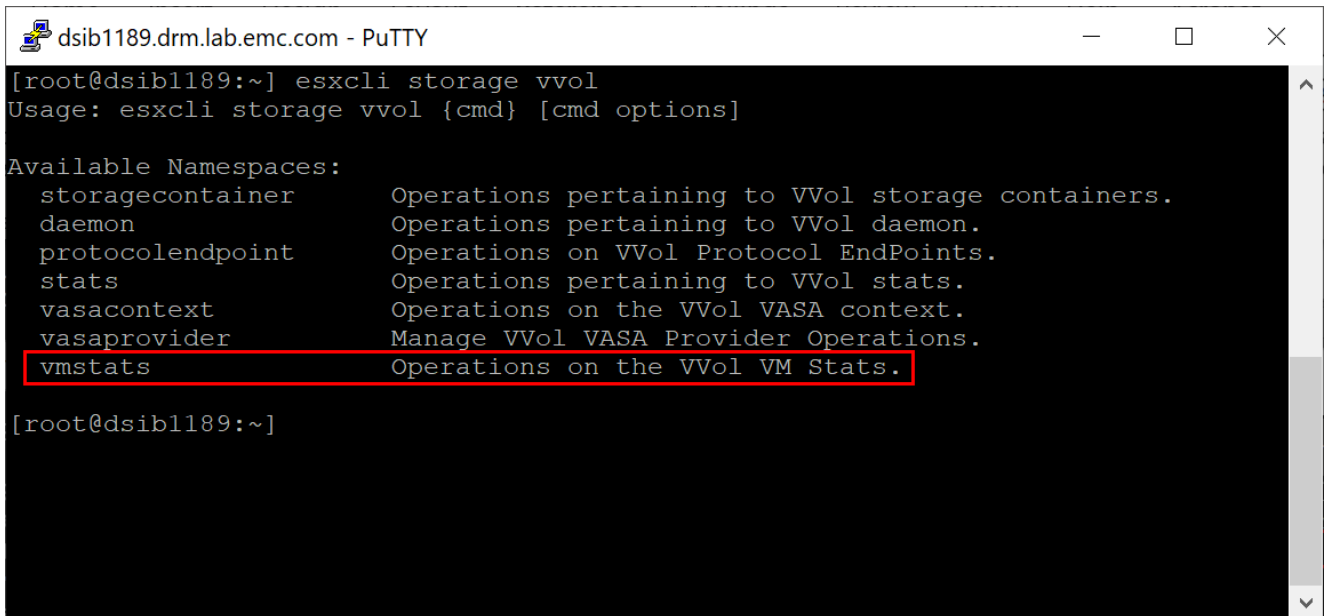


```
dsib0180.lss.emc.com - PuTTY
2021-03-22T08:30:24.007Z - time the service was last started, Section for VMware ESXi, pid=2099510, version=7.0.2
2021-03-22T08:30:24.007Z info vvol[d2100379] [Originator@6876 sub=vvol-stats] Mon Mar 22 08:00:24 2021
-->
--> /vasa-VASA_357/errorCount[TRANSPORT_FAULT]=3
--> /vasa-VASA_357/numberOfConnection=4
--> /vasa-VASA_357/numberOfRetry=2
--> /vasa-VASA_357/count=14161
--> /vasa-VASA_357/updateVirtualVolumeMetaData/LatencyHist[0-250ms]=1
--> /vasa-VASA_357/updateVirtualVolumeMetaData/errorCount[SUCCESS]=1
--> /vasa-VASA_357/updateVirtualVolumeMetaData/count=1
--> /vasa-VASA_357/unbindVirtualVolume/LatencyHist[0-250ms]=2
--> /vasa-VASA_357/unbindVirtualVolume/LatencyHist[0-250ms]=2
--> /vasa-VASA_357/unbindVirtualVolume/LatencyHist[0-250ms]=2
--> Mon Mar 22 08:30:24 2021
-->
-->
2021-03-22T09:00:24.011Z info vvol[d2099990] [Originator@6876 sub=vvol-stats] Mon Mar 22 08:30:24 2021
-->
--> /vasa-VASA_357/errorCount[TRANSPORT_FAULT]=3
--> /vasa-VASA_357/numberOfConnection=4
--> /vasa-VASA_357/queryProtocolEndpoint/RetryCount=2
- vvol-stats-4.log 251/10681 2%
```

Figure 127. vVol stats generated in the /var/log directory

11.5.2 Vmstats

The namespace command, **vmstats**, is shown in [Figure 128](#).



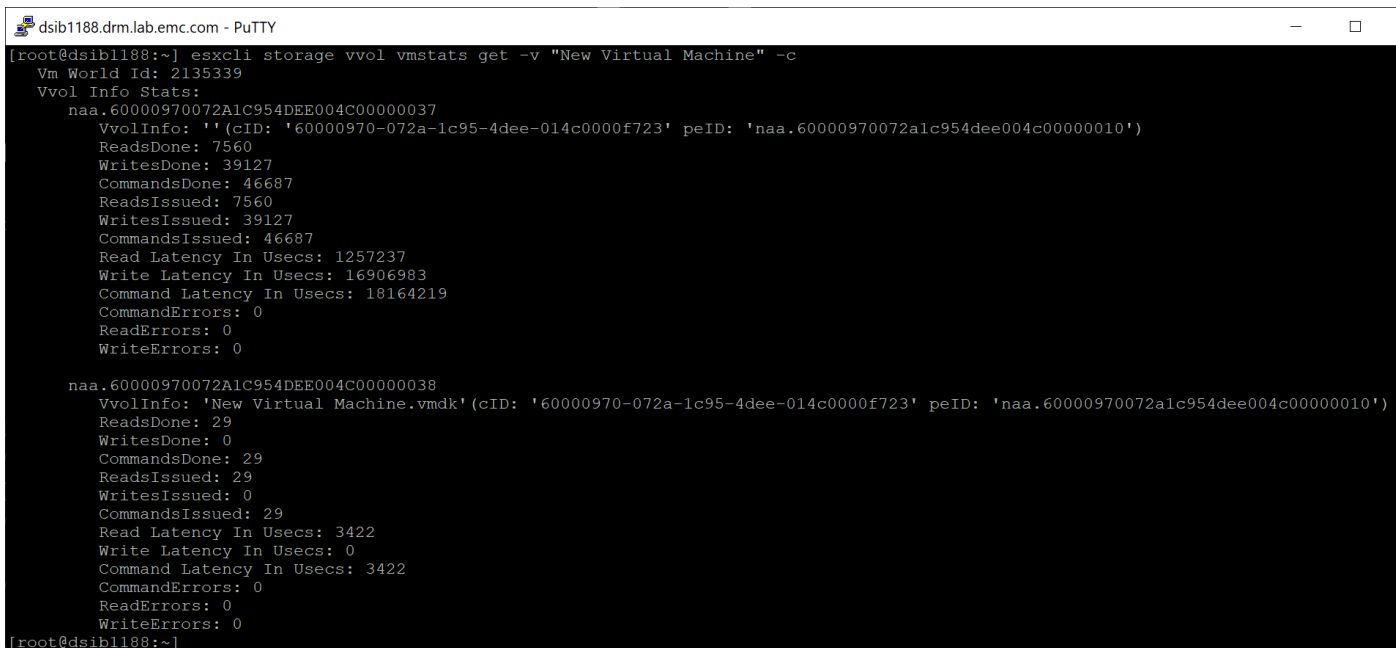
```
dsib1189.drm.lab.emc.com - PuTTY
[root@dsib1189:~] esxcli storage vvol
Usage: esxcli storage vvol {cmd} [cmd options]

Available Namespaces:
  storagecontainer  Operations pertaining to VVol storage containers.
  daemon           Operations pertaining to VVol daemon.
  protocolendpoint  Operations on VVol Protocol EndPoints.
  stats            Operations pertaining to VVol stats.
  vasacontext       Operations on the VVol VASA context.
  vasaprovider      Manage VVol VASA Provider Operations.
  vmstats          Operations on the VVol VM Stats.

[root@dsib1189:~]
```

Figure 128. Vmstats namespace

The command can be used to pull statistics of an individual vVol VM as seen in [Figure 129](#).



```
dsib1188.drm.lab.emc.com - PuTTY
[root@dsib1188:~] esxcli storage vvol vmstats get -v "New Virtual Machine" -c
Vm World Id: 2135339
Vvol Info Stats:
  naa.60000970072A1C954DEE004C00000037
    VvolInfo: '(cID: '60000970-072a-1c95-4dee-014c0000f723' peID: 'naa.60000970072a1c954dee004c00000010')
    ReadsDone: 7560
    WritesDone: 39127
    CommandsDone: 46687
    ReadsIssued: 7560
    WritesIssued: 39127
    CommandsIssued: 46687
    Read Latency In Usecs: 1257237
    Write Latency In Usecs: 16906983
    Command Latency In Usecs: 18164219
    CommandErrors: 0
    ReadErrors: 0
    WriteErrors: 0

  naa.60000970072A1C954DEE004C00000038
    VvolInfo: 'New Virtual Machine.vmdk'(cID: '60000970-072a-1c95-4dee-014c0000f723' peID: 'naa.60000970072a1c954dee004c00000010')
    ReadsDone: 29
    WritesDone: 0
    CommandsDone: 29
    ReadsIssued: 29
    WritesIssued: 0
    CommandsIssued: 29
    Read Latency In Usecs: 3422
    Write Latency In Usecs: 0
    Command Latency In Usecs: 3422
    CommandErrors: 0
    ReadErrors: 0
    WriteErrors: 0
[root@dsib1188:~]
```

Figure 129. Listing vVol VM statistics with esxcli

11.5.3 vVol datastores in a cluster

In a traditional VMFS environment, devices are presented in a single storage group to a VMware cluster. Each host in the cluster sees the same device and therefore when a datastore is created on

that device, upon rescan each host will recognize the new datastore. With vVols, there is no storage group. Each host in a cluster is presented a unique PE to which vVols are bound. When creating a datastore, therefore, the wizard behaves similarly to NFS and will present the hosts in the cluster as available for mounting the datastore. One simply checks the boxes of the hosts which have a PE as in [Figure 130](#).

New Datastore

✓ 1 Type
✓ 2 Name and container selection
3 Select hosts accessibility
4 Ready to complete

Select hosts accessibility
Specify which hosts will have access to the datastore.

Host	Cluster
<input checked="" type="checkbox"/> dsib0186.lss.emc.com	<input checked="" type="checkbox"/> London_Cluster
<input checked="" type="checkbox"/> dsib0184.lss.emc.com	<input checked="" type="checkbox"/> London_Cluster

2 items

CANCEL BACK NEXT

Figure 130. Creating vVol datastores in a cluster

In addition, if hosts are added to the cluster in the future (or if the datastore wizard was originally started at the host and not the cluster), the vVol datastore(s) can be mounted to them as in [Figure 131](#).

Mount Datastore to Additional Hosts... 355_vVol_Datastore

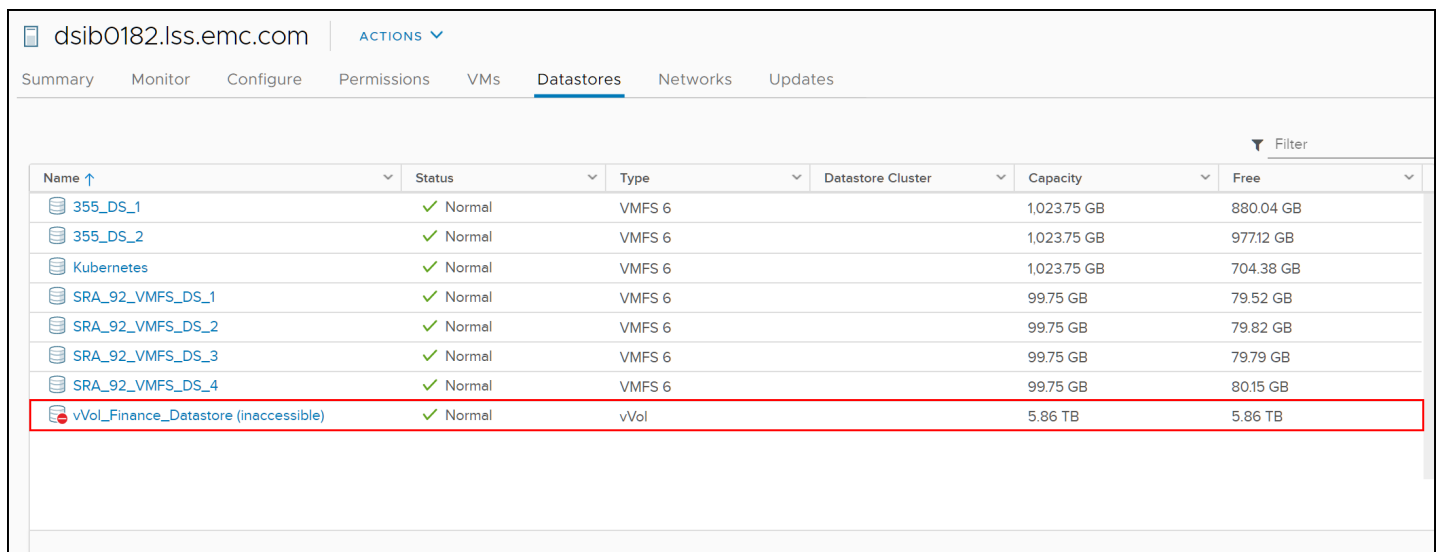
Host	Cluster
<input checked="" type="checkbox"/> dsib0182.lss.emc.com	<input checked="" type="checkbox"/> Boston_Cluster

1 items

CANCEL OK

Figure 131. Mounting vVol datastores to additional hosts in a cluster

Note that VMware does not validate that each host has a presented PE, and therefore if an attempt is made to mount the datastore to a host without a PE, it will show as inaccessible ([Figure 132](#)).



dsib0182.lss.emc.com ACTIONS

Summary Monitor Configure Permissions VMs **Datastores** Networks Updates

Filter

Name	Status	Type	Datastore Cluster	Capacity	Free
355_DS_1	✓ Normal	VMFS 6		1,023.75 GB	880.04 GB
355_DS_2	✓ Normal	VMFS 6		1,023.75 GB	977.12 GB
Kubernetes	✓ Normal	VMFS 6		1,023.75 GB	704.38 GB
SRA_92_VMFS_DS_1	✓ Normal	VMFS 6		99.75 GB	79.52 GB
SRA_92_VMFS_DS_2	✓ Normal	VMFS 6		99.75 GB	79.82 GB
SRA_92_VMFS_DS_3	✓ Normal	VMFS 6		99.75 GB	79.79 GB
SRA_92_VMFS_DS_4	✓ Normal	VMFS 6		99.75 GB	80.15 GB
vVol_Finance_Datastore (inaccessible)	✓ Normal	vVol		5.86 TB	5.86 TB

Figure 132. ESXi host with no presented PE

11.5.4 VMware High Availability (HA)

VMware HA is supported with vVols, though there are some important things to keep in mind. Firstly, as has been made clear, each host must see a unique PE. ESXi hosts in any cluster, including HA, may not share a PE. Secondly, VM Component Protection (VMCP) is not supported with vVols and hence if APD or PDL situations arise, the way the host reacts may not be the same as it will be with VMFS. This includes, but not limited to, vVol VMs not failing over when an APD event is experienced.

If HA is enabled on the cluster where the vVol datastore is created, or enabled after, VMware will request a configuration vVol to serve the purposes of HA in that datastore. A new directory will be created in the datastore named “.vSphere-HA”.

11.5.5 Default profile/capability sets and default Storage Policy for vVol datastores

There are two different default capabilities related to vVol datastores – a default profile based on capability sets and a default Storage Policy.

11.5.5.1 Default profile/capability sets

The VASA 3 and 4 specifications do not support capability sets or default profiles. This is a change from the VASA 2 specification which did support them. Within vCenter, therefore, the only way to determine what capabilities a storage container has from the array, is by testing storage policies and vVol datastore compatibility. In the vCenter, both fields will show empty, with no ability to modify as in Figure 133.

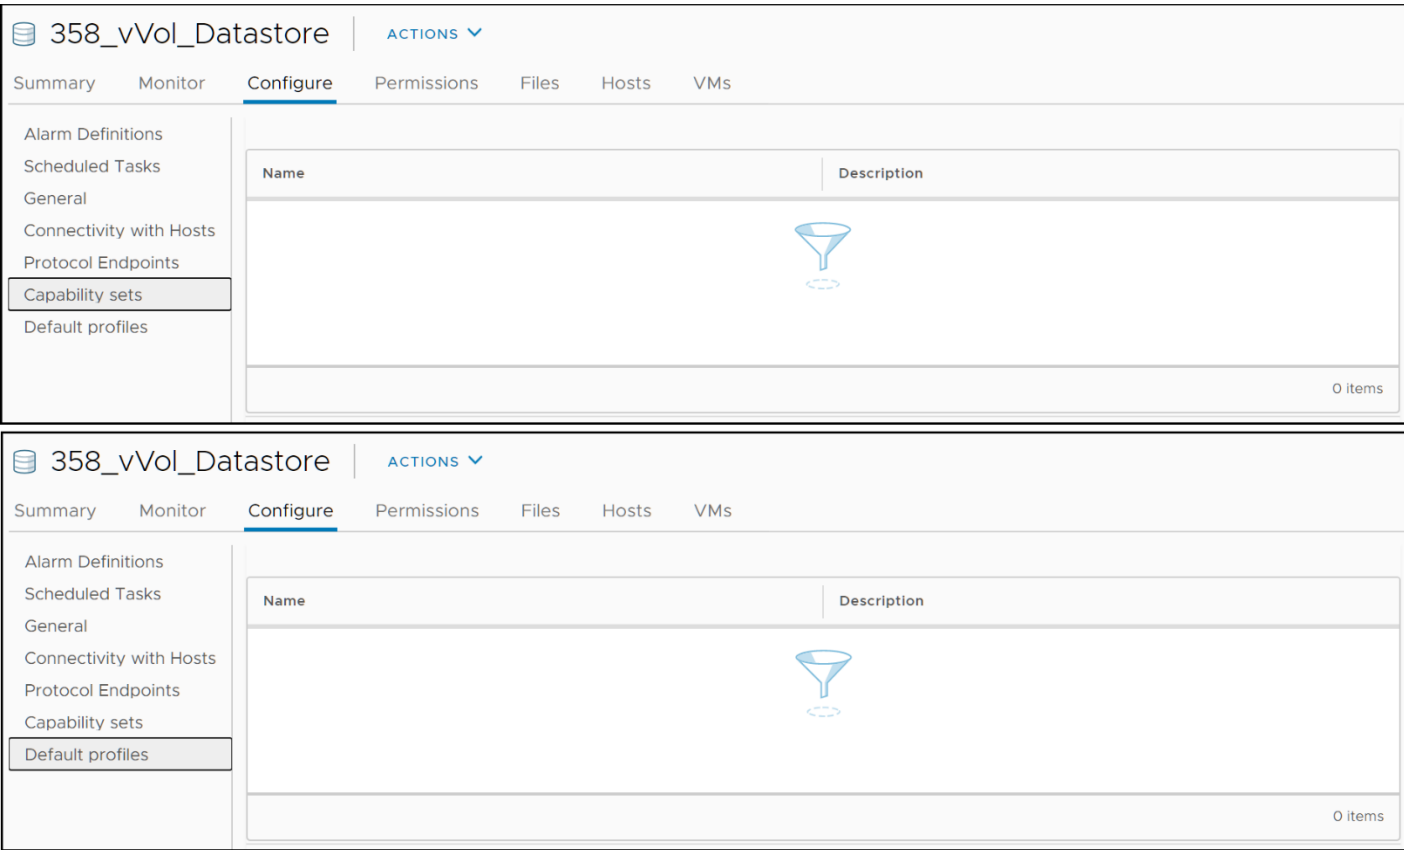


Figure 133. vVol Capability sets and Default profiles

11.5.5.2 Default Storage Policy

It is possible to set a default storage policy on a vVol datastore. The initial policy for a vVol datastore is “VVol No Requirements Policy” present in Figure 134. As set, this policy means the storage resource with the least performance service level in the storage container will be used for the VM files.

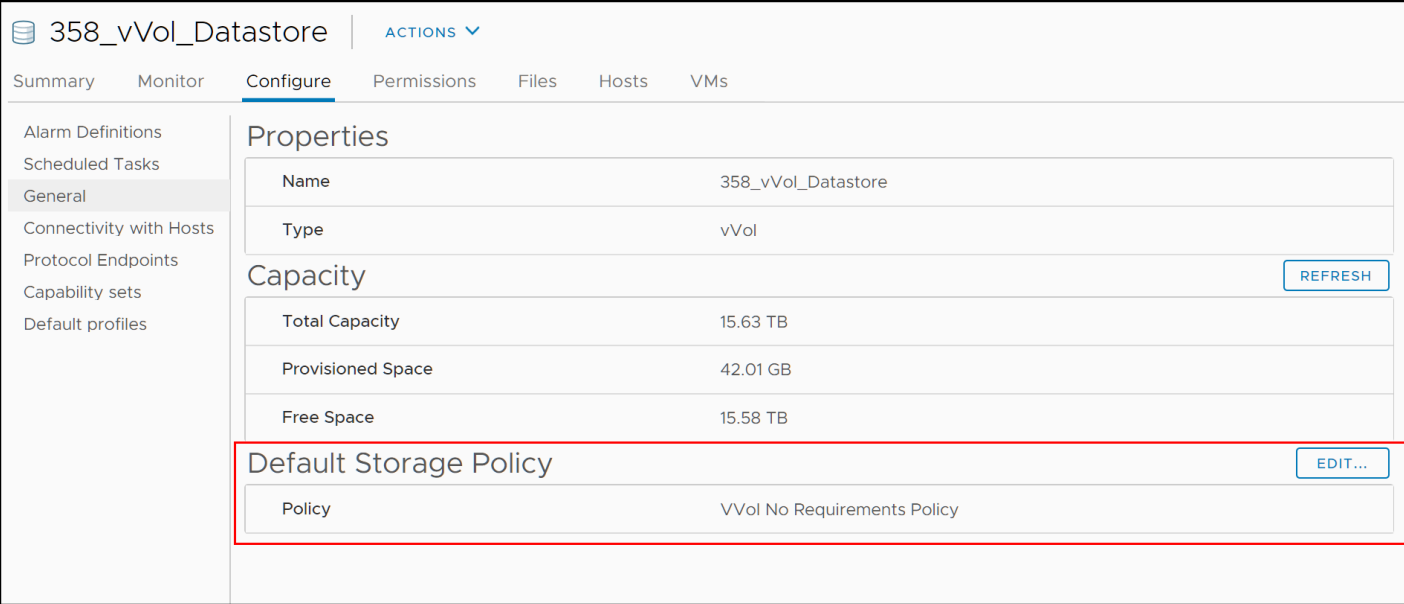


Figure 134. Default Storage Policy

To change the policy to one of the storage policies previously created, start by selecting the **EDIT** button in Figure 134.

In the dialog box that appears in Figure 135, select the desired default storage policy from the available policies and select OK.

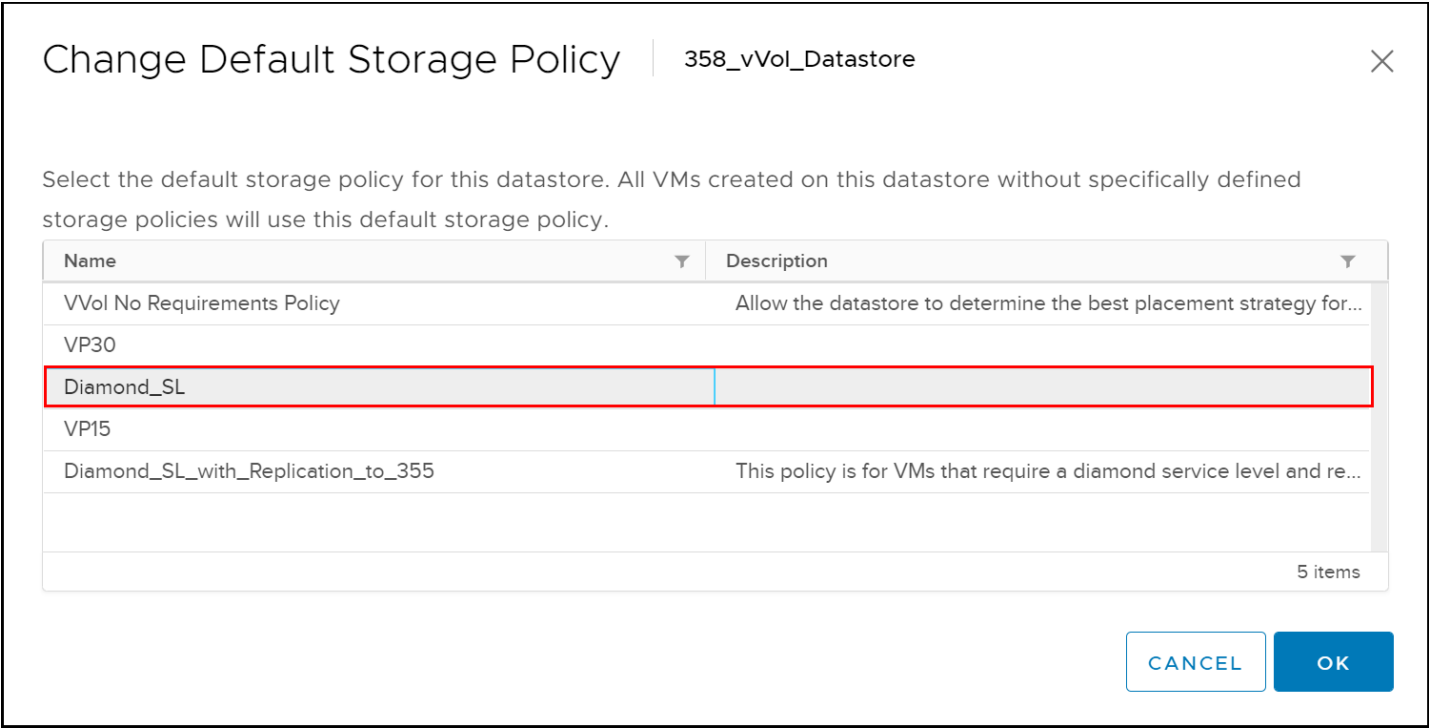


Figure 135. Select new default storage policy

The new policy is now set and can be seen in Figure 136. This policy will now be used when no policy is selected during VM or vmdk creation. Note that if a default policy includes replication, the VM creation will fail since VMware does not offer the ability to set the Replication Group during creation.

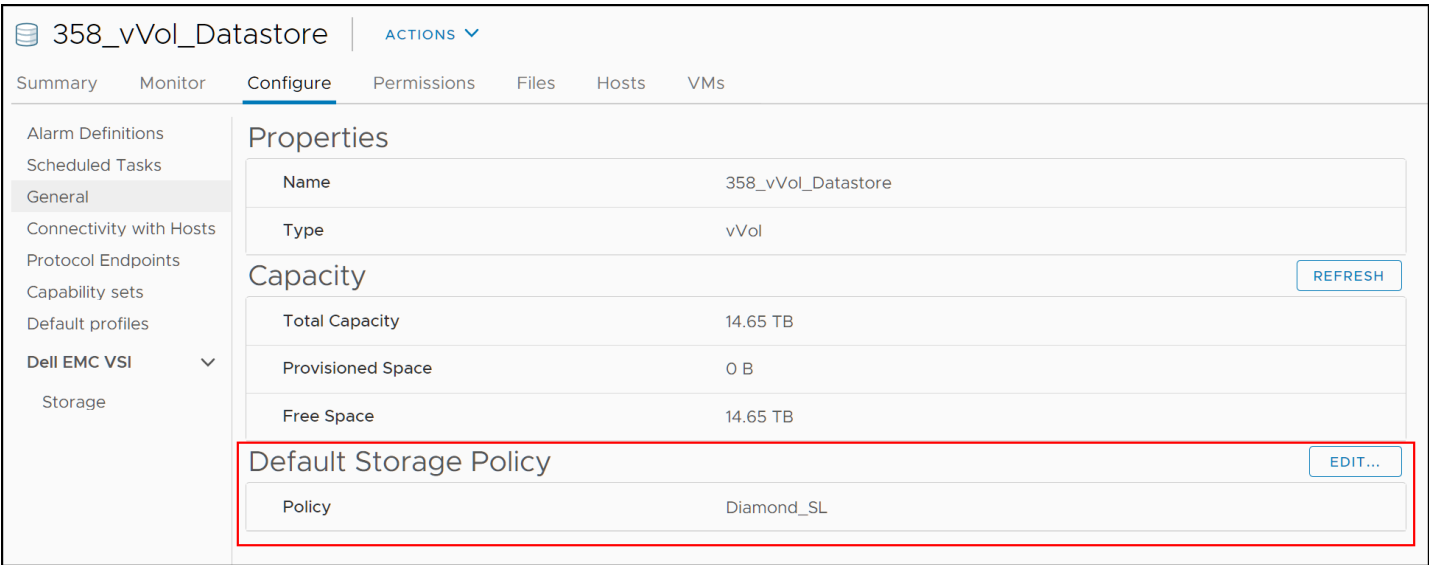


Figure 136. New Default Storage Policy

12 VMFS and vVol Cloning/Migrations

Cloning and migration of VMs between vVol and VMFS datastores is fully supported through the normal Storage vMotion function. Table 1 contains the most common tasks related to cloning and migrating virtual machines and what APIs (simplified to VAAI or VASA) are used to accomplish those tasks. The table is generic, and simply indicates the default and failback (if applicable) functions for these vSphere tasks. The tasks that list both VASA and VAAI mean that VASA must be used for part of the process, but that VAAI can be replaced by host-based copy if the array does not support XCOPY, as is the case with vVols on the PowerMax.

Table 1. Cloning/Migration tasks and functions

Task	Default Function	Failback Function
Clone from VMFS to vVol	VAAI	host-based copy
Clone vVol to VMFS	host-based copy	
Migrate vVol to VMFS	host-based copy	
Clone vVol to vVol in same container	VASA	
Clone vVol to vVol in different container, same array	VASA, VAAI	host-based copy
SvMotion (powered on) without snapshots	VASA, VAAI	host-based copy
SvMotion (powered on) with snapshots	VASA, VAAI	host-based copy
SvMotion (powered off) without snapshots	VASA, VAAI	host-based copy
SvMotion (powered off) with snapshots	VASA, VAAI	host-based copy
Clone vVol to vVol in different container, different array	host-based copy	
Migrate vVol within the different container, different array	host-based copy	

13 vVol identification and monitoring in Unisphere

The following sections describe how a user can utilize Unisphere for PowerMax to obtain general information about vVols as well as performance data.

13.1 Identifying vVol WWN in Unisphere

Although vCenter provides no means to map a vVol to the underlying array device, Unisphere for PowerMax does offer this capability. In order to take advantage of the feature, the vCenter involved with vVols needs to be added to Unisphere. This can be done through the **VMWARE -> vCenters and ESXi** menu in Unisphere seen in [Figure 137](#).

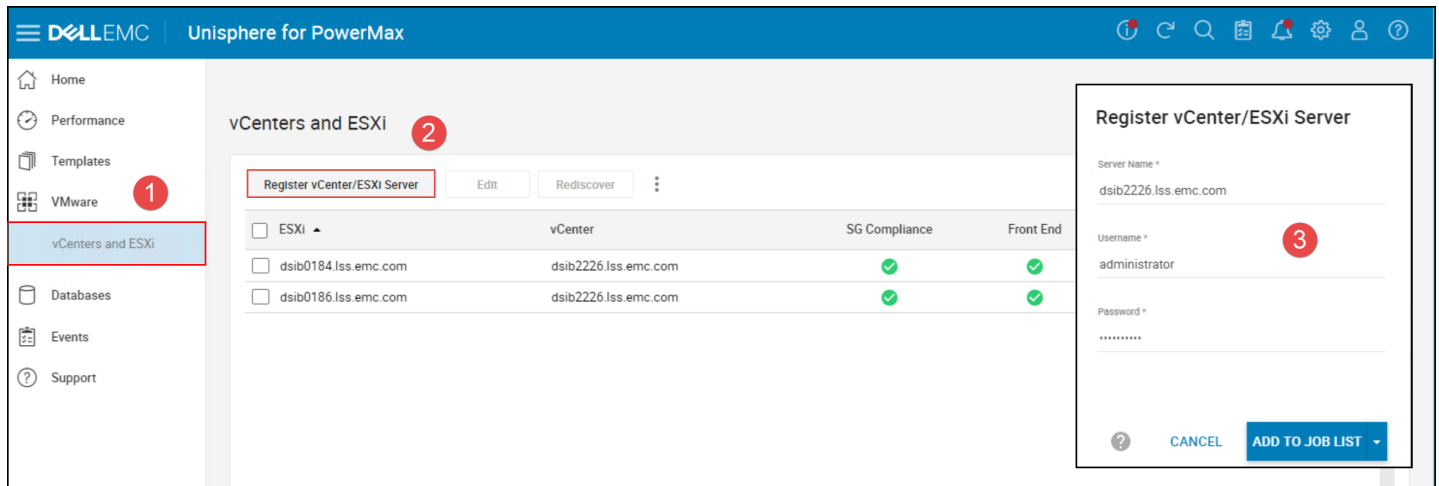


Figure 137. Adding vCenter to Unisphere

Once the vCenter is added, start in [Figure 138](#) by selecting an ESXi host in that vCenter and double-clicking.

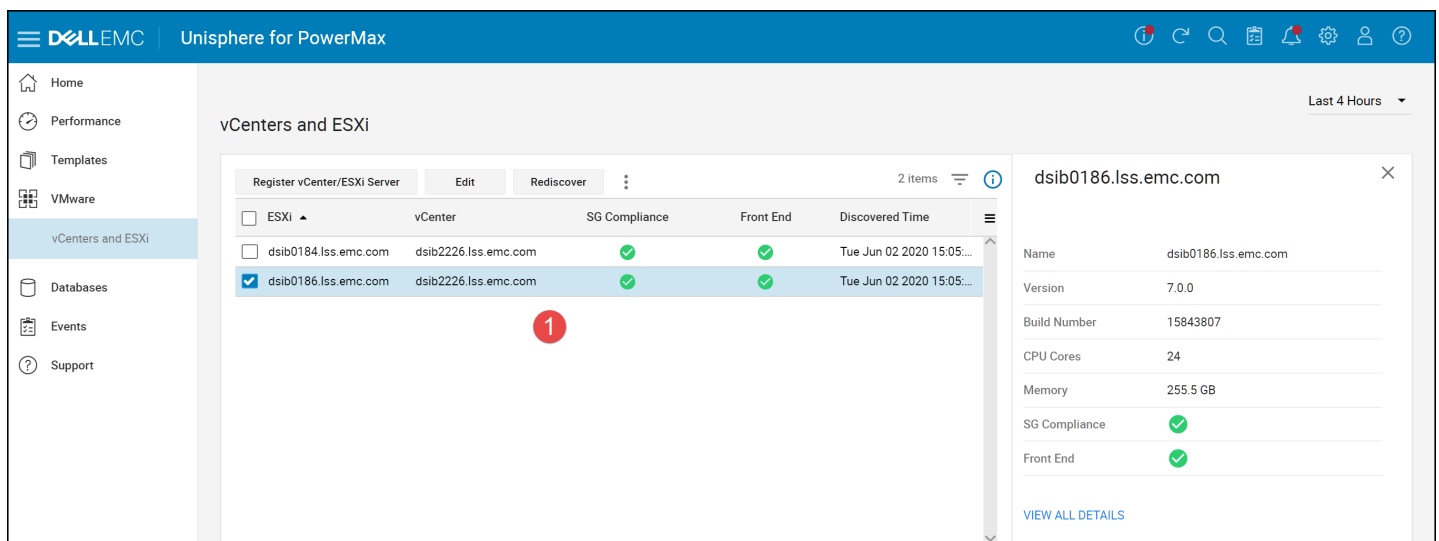


Figure 138. VM vmdk mapping to vVol device - step 1

In steps 2-4 in [Figure 139](#), begin by selecting the Virtual Machines tab, double-click a vVol VM, and in the side panel on the right that will appear, select the hyperlink next to the Virtual Disks row.

Unisphere for PowerMax

VMware > dsib0186.lss.emc.com

DETAILS MASKING VIEWS VIRTUAL MACHINES PERFORMANCE

4 items

Name	Powered	CPU Count	Total Memory (MB)	OS Name
<input checked="" type="checkbox"/> vVol-1	On	4	8192	Microsoft Windows...
<input type="checkbox"/> vVol-8	On	2	4096	Microsoft Windows...
<input type="checkbox"/> vVol-3	On	2	4096	Microsoft Windows...
<input type="checkbox"/> vVol-10	On	2	4096	Microsoft Windows...

Details for vVol-1:

- Name: vVol-1
- Powered: On
- CPU Count: 4
- Total Memory (MB): 8192
- OS Name: Microsoft Windows Server 2012 (64-bit)
- Address: fe80::295d:ea98:8d57:4c6d
- Host Name: dsib2229.lss.com
- Virtual Disks: 4

Figure 139. VM vmdk mapping to vVol device - steps 2-4

In step 5 in Figure 140, double-click on one of the Hard disks (vmdk) and in the right-hand panel that appears, all the information about the vVol is displayed. The red box highlighted in the figure contains the vVol WWN. Note that in Unisphere for PowerMax 10.0, the WWN will have a prefix of "naa.". When searching for the vVol with Solutions Enabler, drop the prefix.

Note: The VVOL WWN field was removed in Unisphere 10.1.0.0 but returned in 10.2.0.0.

Unisphere for PowerMax

VMware > dsib0186.lss.emc.com > Virtual Machines > vVol-1

4 items

Label	Bus	Disk Mode	UUID	Node	Filename	Capacity
<input type="checkbox"/> Hard disk 1	0	Persistent	6000C293-d73a-7540-d4e1-059...	0	[358_vVol_Datastore] 600009700BC724663...	100
<input type="checkbox"/> Hard disk 2	0	Persistent	6000C29a-a9b7-28f0-ab3a-bbce...	1	[358_vVol_Datastore] 600009700BC724663...	10
<input checked="" type="checkbox"/> Hard disk 3	0	Persistent	6000C290-0273-b41f-e7c4-ecc8...	2	[358_vVol_Datastore] 600009700BC724663...	10
<input type="checkbox"/> Hard disk 4	0	Persistent	6000C292-407d-66b6-0c48-3b0f...	3	[358_vVol_Datastore] 600009700BC724663...	10

Details for Hard disk 3:

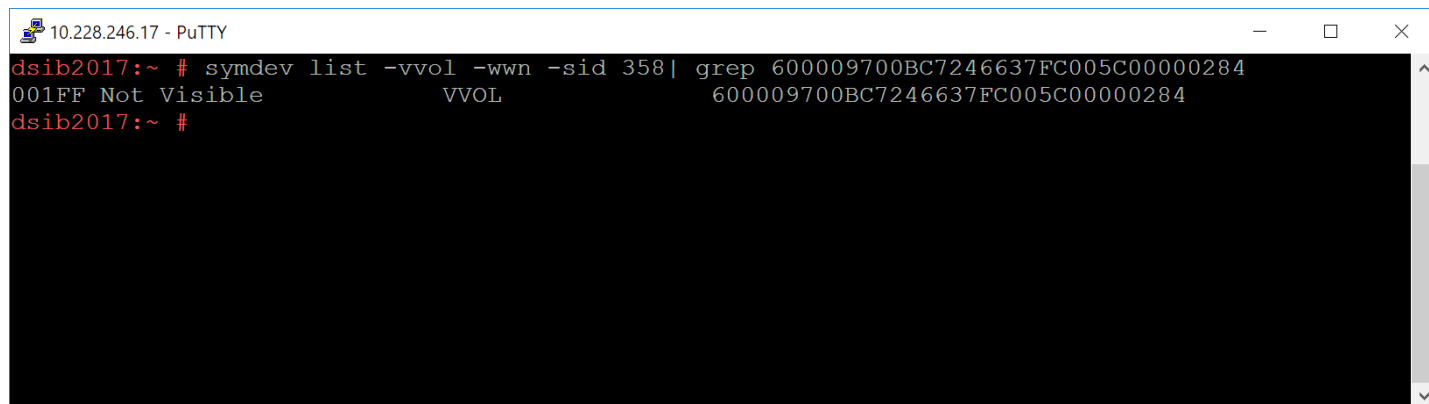
- Label: Hard disk 3
- Bus: 0
- Disk Mode: Persistent
- UUID: 6000C290-0273-b41f-e7c4-ecc8a7e3e543
- Node: 2
- Filename: [358_vVol_Datastore] 600009700BC7246637FC005C0000024D/vVol-1_2.vmdk
- Capacity (GB): 10
- VVOL WWN: 600009700BC7246637FC005C00000284
- Type: VVOL

Figure 140. VM vmdk mapping to vVol device - step 5

To determine the device ID, it is necessary to use Solutions Enabler as the ID is not included in the Unisphere output. The command to list the device IDs with the WWN is:

```
symdev list -vvol -wwn -sid 358 |grep <WWN>
```

Using this command, one can see in [Figure 141](#) that the device ID for the vVol in [Figure 140](#) is 001FF.



```
10.228.246.17 - PuTTY
dsib2017:~ # symdev list -vvol -wwn -sid 358 | grep 600009700BC7246637FC005C00000284
001FF Not Visible          VVOL          600009700BC7246637FC005C00000284
dsib2017:~ #
```

Figure 141. vVol device ID and WWN

In the next section, the device ID can be used to monitor the performance of the vVol.

13.2 vVol Performance Monitoring in Unisphere

As vVols on the PowerMax are not visible at the storage group level, performance monitoring needs to be conducted directly at the individual vVol. All metrics that are gathered for regular thin devices are also gathered for vVols. To view performance in Unisphere for PowerMax for vVols, first navigate to the **PERFORMANCE -> Charts** menu in steps 1 and 2 in [Figure 142](#).

Note: The vVol menu is dynamic and will only appear if vVols are active on the array.

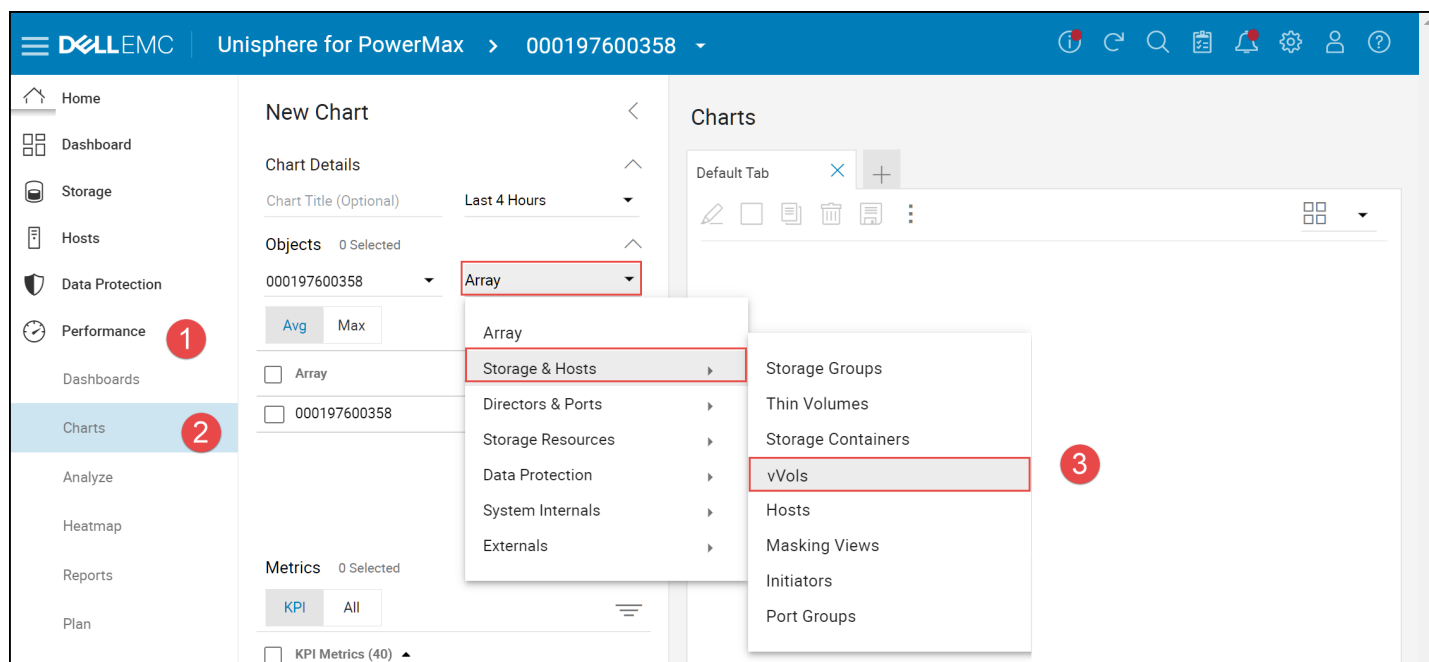


Figure 142. vVol performance monitoring

Next, using the drop-down menu in step 3, traverse to **Storage & Hosts -> vVols**. In this screen the user can select all, or some vVols, and then a set of metrics and generate one or more dashboards as in [Figure 143](#).




Figure 143. Creating vVol dashboards in Unisphere for PowerMax

In addition to vVol metrics, Unisphere for PowerMax also offers metrics for the Storage Containers and the VASA Replication Groups (displayed as RDFA Groups). See the Unisphere for PowerMax online help for more detail on creating dashboards.

14 Scalability

14.1 ESXi

With vVols, there is no change to the ESXi device host limit. As each VM is made up of many vVols, initially this might seem problematic and prevent scalability; however only the PE device counts against the limit. vVols are not mapped and masked directly to the hosts. This permits a vVol environment to scale.

In addition, a single PE can have a maximum of 16,383 vVols bound to it.

Note: The maximum size of an individual vVol on the PowerMax is 16 TB. This is a PowerMax restriction, not VMware's.

14.2 Storage Resources

Dell limits the number of vVols that a single storage resource can contain to a maximum of 8192. Attempts to create any vVols in that resource beyond that value will fail with a general error similar to the following: *"Error: "Cannot complete file creation operation. Operation failed."* It is important to remember that even if the current count of vVols in the storage resource is below the maximum, adding disks or VM snapshots that include memory to VMs in the storage resource, will generate additional vVols. While it is possible to place additional disks to those VMs in a different storage resource, the snapshot vVol will always be placed in the same storage resource. Therefore, in large environments it is important to keep track of the number of vVols in each storage resource. Unfortunately, there is no GUI interface to obtain this information, however the vVol count for a storage resource can be pulled with the CLI.

Utilizing the previously covered commands in the section Using Solutions Enabler with Virtual Volumes, the number of vVols can be pulled from both the storage container and the storage resources within that container. The key to pulling the detail of the number of vVols is to use the output command with the XML format. [Figure 144](#) shows the detail of storage container 358_Container using the following command:

```
symcfg show -sc -sc_name 358_Container -sid 358 -detail -out xml
```

The first red box in the figure highlights the name of the container while the second red box details the number of vVols, 37, and the number snapshots. The xml output does not distinguish between snapshots that include the swap and those that do not. Therefore, the total vVols of 37 may also include the snapshot vVols. As the max vVol limit does not apply to the storage container, each storage resource needs to be examined. The blue and green boxes in the figure contain the total vVols and the total number of snapshots in each resource.

```

10.228.246.17 - PuTTY
dsib2017:~ # symcfg show -sc -sc name 358_Container -sid 358 -detail -out xml
<?xml version="1.0" standalone="yes" ?>
<SymCLI ML>
  <Symmetrix>
    <Symm_Info>
      <syamid>000197600358</syamid>
    </Symm_Info>
    <Storage_Containers>
      <count>1</count>
      <Storage_Container_Info>
        <name>358_Container</name>
        <uuid>600009700bc7246637fc015c00012c32</uuid>
        <description>
        </description>
        <type>VVOL</type>
        <subscribed_max_gb>20000.0</subscribed_max_gb>
        <subscribed_max_tb>19.53</subscribed_max_tb>
        <subscribed_capacity_gigabytes>1256.0</subscribed_capacity_gigabytes>
        <subscribed_capacity_terabytes>1.23</subscribed_capacity_terabytes>
        <subscribed_capacity_pct>6</subscribed_capacity_pct>
        <num_devs>37</num_devs>
        <num_snapshots>3</num_snapshots>
      </Storage_Container_Info>
      <Storage_Resources>
        <count>2</count>
        <Storage_Resource_Info>
          <name>358_Container_resource_1</name>
          <uuid>600009700bc7246637fc015c00012c33</uuid>
          <type>StorageResource</type>
          <SLO_name>Bronze</SLO_name>
          <Workload>none</Workload>
          <SRP_name>SRP_1</SRP_name>
          <compression>Yes</compression>
          <compression_ratio>2.2:1</compression_ratio>
          <subscribed_max_gb>15000.0</subscribed_max_gb>
          <subscribed_max_tb>14.65</subscribed_max_tb>
          <subscribed_capacity_gigabytes>85.0</subscribed_capacity_gigabytes>
          <subscribed_capacity_terabytes>0.08</subscribed_capacity_terabytes>
          <subscribed_capacity_pct>0</subscribed_capacity_pct>
          <num_devs>14</num_devs>
          <num_snapshots>1</num_snapshots>
        </Storage_Resource_Info>
        <Storage_Resource_Info>
          <name>358_Container_resource_2</name>
          <uuid>600009700bc7246637fc015c0001f048</uuid>
          <type>StorageResource</type>
          <SLO_name>Diamond</SLO_name>
          <Workload>none</Workload>
          <SRP_name>SRP_1</SRP_name>
          <compression>Yes</compression>
          <compression_ratio>1.6:1</compression_ratio>
          <subscribed_max_gb>5000.0</subscribed_max_gb>
          <subscribed_max_tb>4.88</subscribed_max_tb>
          <subscribed_capacity_gigabytes>1171.0</subscribed_capacity_gigabytes>
          <subscribed_capacity_terabytes>1.14</subscribed_capacity_terabytes>
          <subscribed_capacity_pct>23</subscribed_capacity_pct>
          <num_devs>23</num_devs>
          <num_snapshots>2</num_snapshots>
        </Storage_Resource_Info>
      </Storage_Resources>
    </Storage_Containers>
  </Symmetrix>
</SymCLI ML>
dsib2017:~ #

```

Figure 144. Obtaining vVol count in vVol storage objects on PowerMax 2000/8000

Using this methodology, the vVol count per storage resource can be managed.

Since each storage container can only have a single storage resource of a particular service level, if a resource approaches the 8192 count it may be necessary to create a new storage container with a new storage resource of that service level. VMs or individual vmdks can then be moved to the new container.

On the V4 platform the XML output also has “provisioned” values which are mirrors of the subscribed values. An example is in [Figure 145](#).


```

root@dsib2005:~
[root@dsib2005 ~]# symcfg show -sc -sc_name ORACLE_302_PCLI -sid 302 -detail -out xml
<?xml version="1.0" standalone="yes" ?>
<SymCLI_ML>
  <Symmetrix>
    <Symm_Info>
      <symid>000120200302</symid>
    </Symm_Info>
    <Storage_Containers>
      <count>1</count>
      <Storage_Container_Info>
        <name>ORACLE_302_PCLI</name>
        <uuid>60000970072alc6e4a0501b20000000a</uuid>
        <description>
        </description>
        <type>VVOLS</type>
        <protocol>SCSI</protocol>
        <provisioned_max_gb>21000.0</provisioned_max_gb>
        <provisioned_max_tb>20.51</provisioned_max_tb>
        <provisioned_capacity_gigabytes>5314.0</provisioned_capacity_gigabytes>
        <provisioned_capacity_terabytes>5.19</provisioned_capacity_terabytes>
        <provisioned_capacity_pct>25</provisioned_capacity_pct>
        <subscribed_max_gb>21000.0</subscribed_max_gb>
        <subscribed_max_tb>20.51</subscribed_max_tb>
        <subscribed_capacity_gigabytes>5314.0</subscribed_capacity_gigabytes>
        <subscribed_capacity_terabytes>5.19</subscribed_capacity_terabytes>
        <subscribed_capacity_pct>25</subscribed_capacity_pct>
        <num_devs>30</num_devs>
        <num_snapshots>1</num_snapshots>
        <Storage_Resources>
          <count>2</count>
          <Storage_Resource_Info>
            <name>ORACLE_302_PCLI_resource_1</name>
            <uuid>60000970072alc6e4a0501b20000000b</uuid>
            <type>StorageResource</type>
            <SLO_name>Diamond</SLO_name>
            <Workload>none</Workload>
            <SRP_name>SRP_1</SRP_name>
            <compression>Yes</compression>
            <compression_ratio>1.0:1</compression_ratio>
            <provisioned_max_gb>20000.0</provisioned_max_gb>
            <provisioned_max_tb>19.53</provisioned_max_tb>
            <provisioned_capacity_gigabytes>5035.0</provisioned_capacity_gigabytes>
            <provisioned_capacity_terabytes>4.92</provisioned_capacity_terabytes>
            <provisioned_capacity_pct>25</provisioned_capacity_pct>
            <subscribed_max_gb>20000.0</subscribed_max_gb>
            <subscribed_max_tb>19.53</subscribed_max_tb>
            <subscribed_capacity_gigabytes>5035.0</subscribed_capacity_gigabytes>
            <subscribed_capacity_terabytes>4.92</subscribed_capacity_terabytes>
            <subscribed_capacity_pct>25</subscribed_capacity_pct>
            <num_devs>28</num_devs>
            <num_snapshots>1</num_snapshots>
          </Storage_Resource_Info>
          <Storage_Resource_Info>
            <name>ORACLE_302_PCLI_resource_3</name>
            <uuid>60000970072alc6e4a0501b200001848</uuid>
            <type>StorageResource</type>
            <SLO_name>Silver</SLO_name>
            <Workload>none</Workload>
            <SRP_name>SRP_1</SRP_name>
            <compression>Yes</compression>
            <compression_ratio>1.4:1</compression_ratio>
            <provisioned_max_gb>1000.0</provisioned_max_gb>
            <provisioned_max_tb>0.98</provisioned_max_tb>
            <provisioned_capacity_gigabytes>279.0</provisioned_capacity_gigabytes>
            <provisioned_capacity_terabytes>0.27</provisioned_capacity_terabytes>
            <provisioned_capacity_pct>27</provisioned_capacity_pct>
            <subscribed_max_gb>1000.0</subscribed_max_gb>
            <subscribed_max_tb>0.98</subscribed_max_tb>
            <subscribed_capacity_gigabytes>279.0</subscribed_capacity_gigabytes>
            <subscribed_capacity_terabytes>0.27</subscribed_capacity_terabytes>
            <subscribed_capacity_pct>27</subscribed_capacity_pct>
            <num_devs>2</num_devs>
            <num_snapshots>0</num_snapshots>
          </Storage_Resource_Info>
        </Storage_Resources>
      </Storage_Container_Info>
    </Storage_Containers>
  </Symmetrix>
</SymCLI_ML>
[root@dsib2005 ~]#

```

Figure 145. Obtaining vVol count in vVol storage objects on PowerMax 2500/8500

14.3 VM snapshot sizing

One of the benefits of vVols is the ability to use the array data services directly. Most of these are applied through Storage Policy Based Management (SPBM), but the most direct integration is snapshots. When taking a snapshot of a vVol-based VM, VMware passes off control to the array to take the snapshot using, in this case, TimeFinder technology. Array snapshots are incredibly efficient because they are targetless. Unlike traditional VMFS, there is no requirement to keep delta files on the file system. All deltas are kept on the array. In addition to the targetless snapshot, if the user wishes to capture the memory state of the VM, VASA generates a static vVol with that information, equaling the size of the VM memory.

Since tracks changes for snapshots are stored on the array, there is some storage usage. The question then is how is this accounted for? Is the usage deducted from the storage container/resource? Yes and no. The following example will explain the process.

14.3.1 Setup

A single storage container is created, Demo, which contains two storage resources in [Figure 146](#):

- Demo_resource_1 is 1000 GB with an SL of Gold
- Demo_resource_2 is 1000 GB with an SL of Silver

Therefore, there are two storage policies, one for Gold and one for Silver.

There is a single 100 GB vVol VM, i.e., a single 100 GB vmdk, created with a Gold storage policy. In the vVol datastore (the Demo storage container) there are 3 vVols – 1 – 4 GB config vVol, 1 – 8 GB swap vVol, and 1 – 100 GB data vVol, or 112 GB of allocated vVols. These are the storage resources in [Figure 146](#) listing the amount subscribed.

Demo_resource_1	×	Demo_resource_2	×
Name	Demo_resource_1	Name	Demo_resource_2
Storage Container	Demo	Storage Container	Demo
SRP	SRP_1	SRP	SRP_1
Service Level	Gold	Service Level	Silver
Workload	—	Workload	—
Data Reduction	✓	Data Reduction	✓
Compression Ratio	1:1	Compression Ratio	—
Subscribed Limit(GB)	1000	Subscribed Limit(GB)	1000
Subscribed Used(GB)	104	Subscribed Used(GB)	8
Subscribed Free(GB)	896	Subscribed Free(GB)	992

Figure 146. Storage Resources

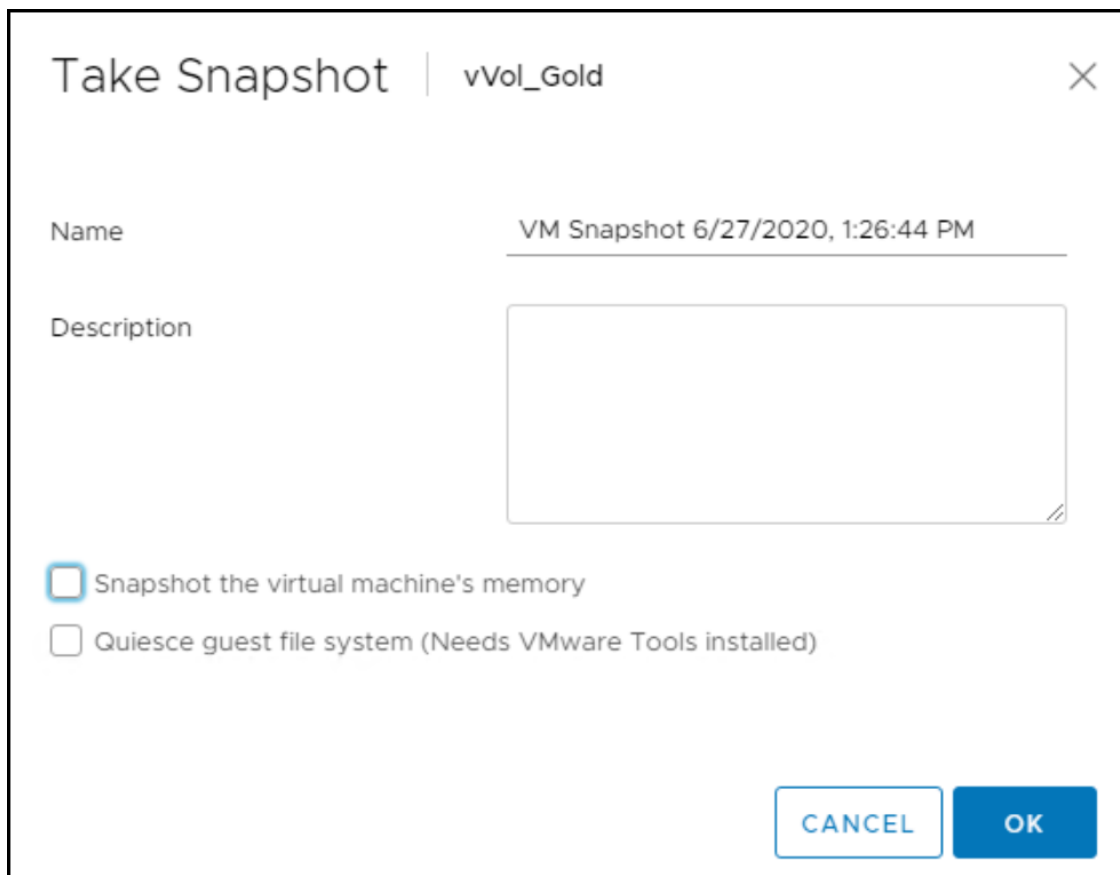
After reviewing one may wonder why there is only 104 GB subscribed in the Gold storage resource when the expected total would be 112 GB. The vVol implementation is designed such that the swap

vVol will always be placed in the lowest performing service level in the container, so in this case the 8 GB swap is in the Silver storage resource. As the data associated with that vVol is in memory on ESXi, the idea is to not waste space in the higher performing service levels.

14.3.2 Snapshot preserved space

14.3.2.1 Targetless snapshot

The first step is to take a targetless (no memory) snapshot of the VM in [Figure 147](#).



The screenshot shows a 'Take Snapshot' dialog box for a virtual machine named 'vVol_Gold'. The dialog has a title bar with the text 'Take Snapshot | vVol_Gold' and a close button (X) in the top right corner. Inside the dialog, there is a 'Name' field with the text 'VM Snapshot 6/27/2020, 1:26:44 PM' and a 'Description' field which is an empty text area. Below these fields, there are two checkboxes: 'Snapshot the virtual machine's memory' (which is checked) and 'Quiesce guest file system (Needs VMware Tools installed)' (which is unchecked). At the bottom right of the dialog, there are two buttons: 'CANCEL' and 'OK'.

Figure 147. Targetless snapshot

It finishes almost instantaneously since the array is doing it. Once it completes, the summary information for the VM in [Figure 148](#) will show that one snapshot is taken, utilizing 100 GB.

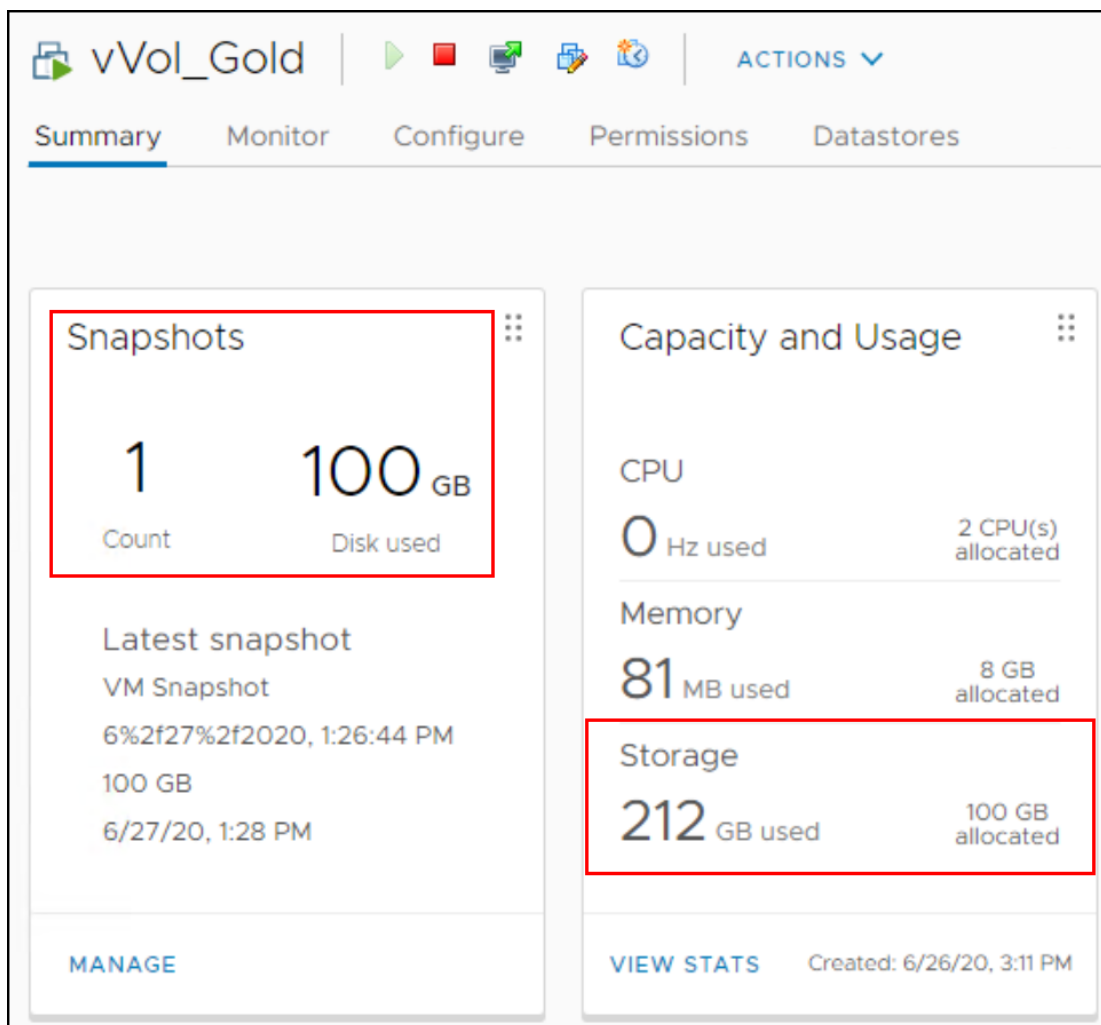


Figure 148. VM summary

VMware records the single snapshot of the data vVol (100 GB) and shows that the VM is using a total of 212 GB (note that VMware shows only 100 GB of that is allocated since the other 100 GB is the snapshot). Reviewing the storage resource on the array in [Figure 149](#), it now indicates there is 204 GB subscribed.

Demo_resource_1

</

Figure 149. Storage resource allocation without memory

The array has reserved that 100 GB of snapshot space that it is not associated with a vVol. The reservation is a preventive step for future restores. Since SnapVX technology is utilized on the array, the first thing the array will need to do for a restore is create a target device to link the snapshot to. If the space is not reserved up-front, there is always the potential that the storage resource could be out of space by the time the restore is run, and thus the restore would fail.

14.3.2.2 Memory snapshot

In step 2, take a memory snapshot. Note that the storage usage in [Figure 150](#) is now at 320 GB, not 312 GB, because of the 8 GB of memory.

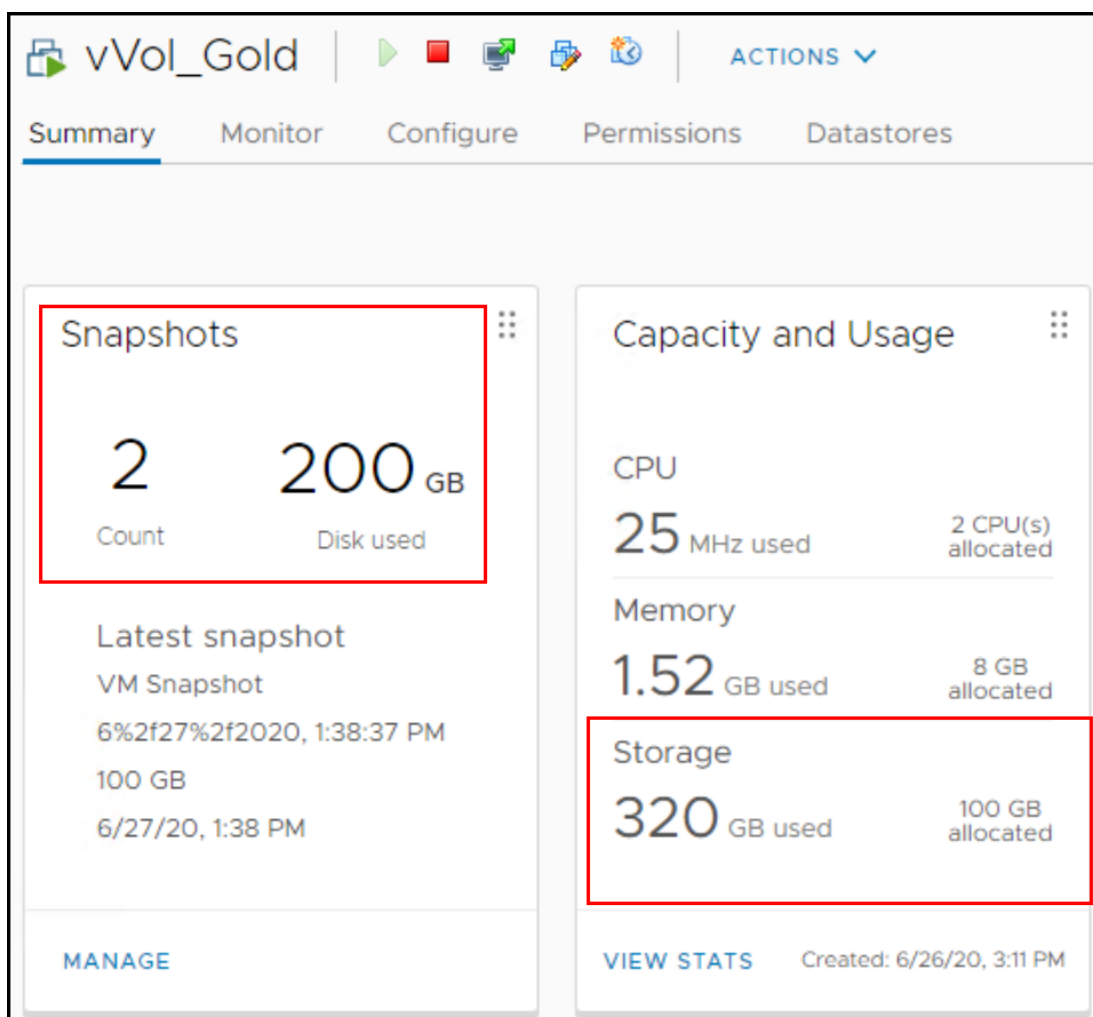


Figure 150. Memory snapshot

However, in [Figure 151](#) see how the memory snapshot vVol is stored in the Silver service level, not Gold, as Silver is the lower tier.

Demo_resource_1		Demo_resource_2	
Name	Demo_resource_1	Name	Demo_resource_2
Storage Container	Demo	Storage Container	Demo
SRP	SRP_1	SRP	SRP_1
Service Level	Gold	Service Level	Silver
Workload	—	Workload	—
Data Reduction	✓	Data Reduction	✓
Compression Ratio	1:1	Compression Ratio	1:1
Subscribed Limit(GB)	1000	Subscribed Limit(GB)	1000
Subscribed Used(GB)	304	Subscribed Used(GB)	16
Subscribed Free(GB)	696	Subscribed Free(GB)	984

Figure 151. Storage resource allocation with memory

14.3.2.3 Overallocation

The creation of the swap file is one of those tasks that will not fail due to lack of space. For example, if this VM is powered down, and then the capacity of the Silver resource adjusted down to 10 GB so that the 8 GB swap will no longer fit, VASA will over-allocate the resource to prevent failure as in [Figure 152](#).

Demo_resource_2

Name	Demo_resource_2
Storage Container	Demo
SRP	SRP_1
Service Level	Silver
Workload	—
Data Reduction	✓
Compression Ratio	1:1
Subscribed Limit(GB)	10.06
Subscribed Used(GB)	16
Subscribed Free(GB)	-5.94

Figure 152. Overallocation of the storage resource

Once a storage resource is oversubscribed, no additional files will be placed there.

14.4 Storage Demand Report

Although a Storage Container does not reserve space in the SRP, the active usage of the container is accounted for at the storage group level. Keeping track in the storage container would be resource intensive; however, it is still possible to determine the snapshot usage in the SRP for those storage resources. Within Unisphere for PowerMax there is a Storage Demand Report which shows how much each storage group is using, including snapshot space. Each storage resource in a storage container is represented by a storage group. There is a difference in navigation between the PowerMax platforms so screenshots will be shown for both, but only the first section will detail the report's usage.

14.4.1 PowerMax 2000/8000

First navigate to the main dashboard of the array and select **CAPACITY** in Figure 153.

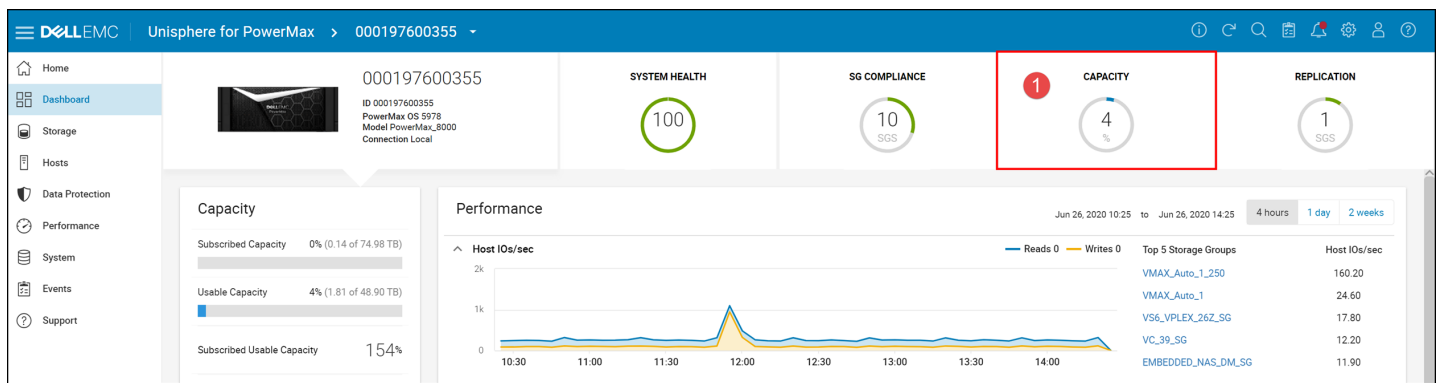


Figure 153. Storage Group Demand Report - step 1

In the second step, use the System drop-down to select the correct **SRP**. Then in step 3 select **STORAGE GROUP DEMAND** Action. These are outlined in Figure 154.

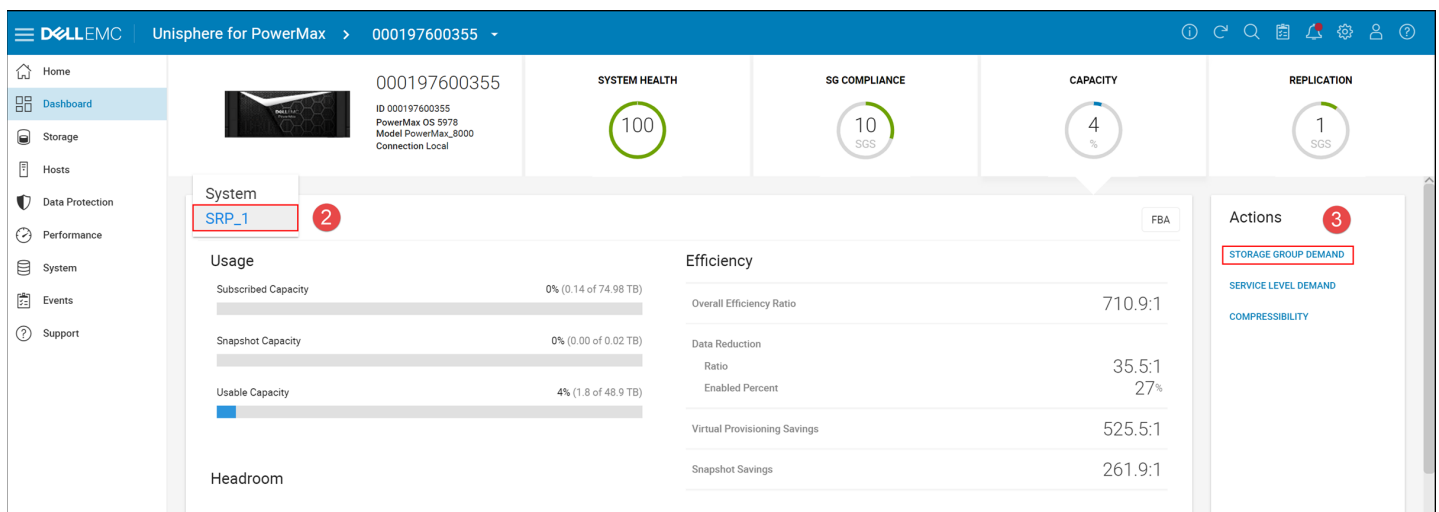


Figure 154. Storage Group Demand Report - step 2 - 3

The Storage Group Demand Report in Figure 155 will list all storage resources in the container as storage groups with the prefix “_VVOLS_”. Here one can see the snapshot information for two storage resources. Again, the snapshot data listed here will not be accounted for in the storage resource, just the SRP. Also note that this report does not include the snapshot reserved space, since that is only reserved in the storage resource and is not represented by vVols in a storage group.

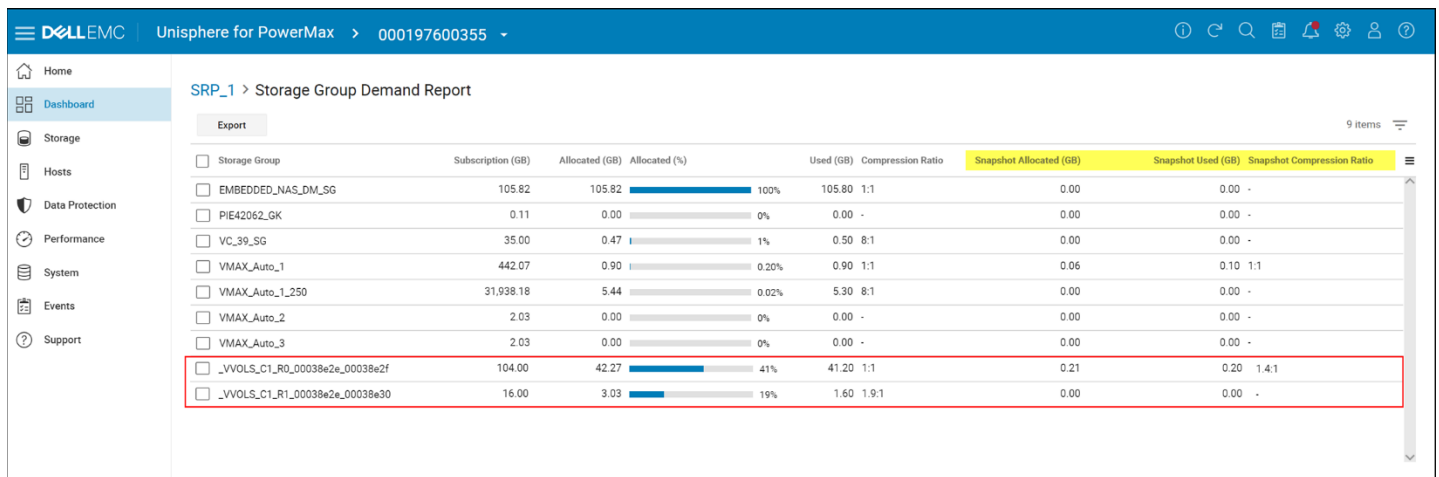


Figure 155. Storage Group Demand Report

Unfortunately, the naming of the storage groups for vVols is not straightforward. In fact, it looks downright cryptic:

- _VVOLS_C1_R0_00038e2e_00038e2f
- _VVOLS_C1_R1_00038e2e_00038e30

Fortunately, it can be decoded using the *symcfg* Solutions Enabler command by outputting the storage container in XML format. The XML format will give extra details about the container and its storage resources. Part of the detail is the UUID of the container itself and each storage resource. This is how the vVol storage group is constructed. The format is “_VVOLS_C” then the number for the order in which the container was created starting with 0. In this case the demo container is the second one, so it gets the number “1”. Next is “_R” followed by the resource number, again in the order created. The final two parts are comprised first of the last 8 characters of the container UUID then the last 8 characters of the resource UUID. The values are highlighted in Figure 156, but decoding:

- _VVOLS_C1_R0_00038e2e_00038e2f – Gold storage resource
- _VVOLS_C1_R1_00038e2e_00038e30 – Silver storage resource

```

dsib2017:~ # symcfg show -sc -sc_name demo -sid 355 -detail -out xml
<?xml version="1.0" standalone="yes" ?>
<SymCLI_ML>
  <Symmetrix>
    <Symm_Info>
      <symid>000197600355</symid>
    </Symm_Info>
    <Storage_Containers>
      <count>1</count>
      <Storage_Container_Info>
        <name>Demo</name>
        <uuid>600009700bc724633801013600038e2e</uuid>
        <description>
        </description>
        <type>VVOLS</type>
        <subscribed_max_gb>2000.0</subscribed_max_gb>
        <subscribed_max_tb>1.95</subscribed_max_tb>
        <subscribed_capacity_gigabytes>320.0</subscribed_capacity_gigabytes>
        <subscribed_capacity_terabytes>0.31</subscribed_capacity_terabytes>
        <subscribed_capacity_pct>16</subscribed_capacity_pct>
        <num_devs>4</num_devs>
        <num_snapshots>2</num_snapshots>
        <Storage_Resources>
          <count>2</count>
          <Storage_Resource_Info>
            <name>Demo_resource_1</name>
            <uuid>600009700bc724633801013600038e2f</uuid>
            <type>StorageResource</type>
            <SLO_name>Gold</SLO_name>
            <Workload>none</Workload>
            <SRP_name>SRP_1</SRP_name>
            <compression>Yes</compression>
            <compression_ratio>1.0:1</compression_ratio>
            <subscribed_max_gb>1000.0</subscribed_max_gb>
            <subscribed_max_tb>0.98</subscribed_max_tb>
            <subscribed_capacity_gigabytes>304.0</subscribed_capacity_gigabytes>
            <subscribed_capacity_terabytes>0.30</subscribed_capacity_terabytes>
            <subscribed_capacity_pct>30</subscribed_capacity_pct>
            <num_devs>2</num_devs>
            <num_snapshots>2</num_snapshots>
          </Storage_Resource_Info>
          <Storage_Resource_Info>
            <name>Demo_resource_2</name>
            <uuid>600009700bc724633801013600038e30</uuid>
            <type>StorageResource</type>
            <SLO_name>Silver</SLO_name>
            <Workload>none</Workload>
            <SRP_name>SRP_1</SRP_name>
            <compression>Yes</compression>
            <compression_ratio>1.9:1</compression_ratio>
            <subscribed_max_gb>1000.0</subscribed_max_gb>
            <subscribed_max_tb>0.98</subscribed_max_tb>
            <subscribed_capacity_gigabytes>16.0</subscribed_capacity_gigabytes>
            <subscribed_capacity_terabytes>0.02</subscribed_capacity_terabytes>
            <subscribed_capacity_pct>1</subscribed_capacity_pct>
            <num_devs>2</num_devs>
            <num_snapshots>0</num_snapshots>
          </Storage_Resource_Info>
        </Storage_Resources>
      </Storage_Container_Info>
    </Storage_Containers>
  </Symmetrix>
</SymCLI_ML>
dsib2017:~ #

```

Figure 156. XML output of storage resource

The XML output also has important information about the container. In particular it will show you how many vVols are in each storage resource. For this example, there are 2 vVols in the Gold resource and 2 snapshots, while in the Silver resource there is also 2 vVols, those being the swap (memory) vVols. Additional information about this command can be found in the section Storage Resources.

14.4.2 PowerMax 2500/8500

On the V4 platform it is no longer called the Storage Demand Report, rather it is the **Storage Group Data Reduction** screen. First navigate to the Capacity screen in step 1 and then in step 2 click on the area of **Data Reduction Ratio/Reducing Data**.

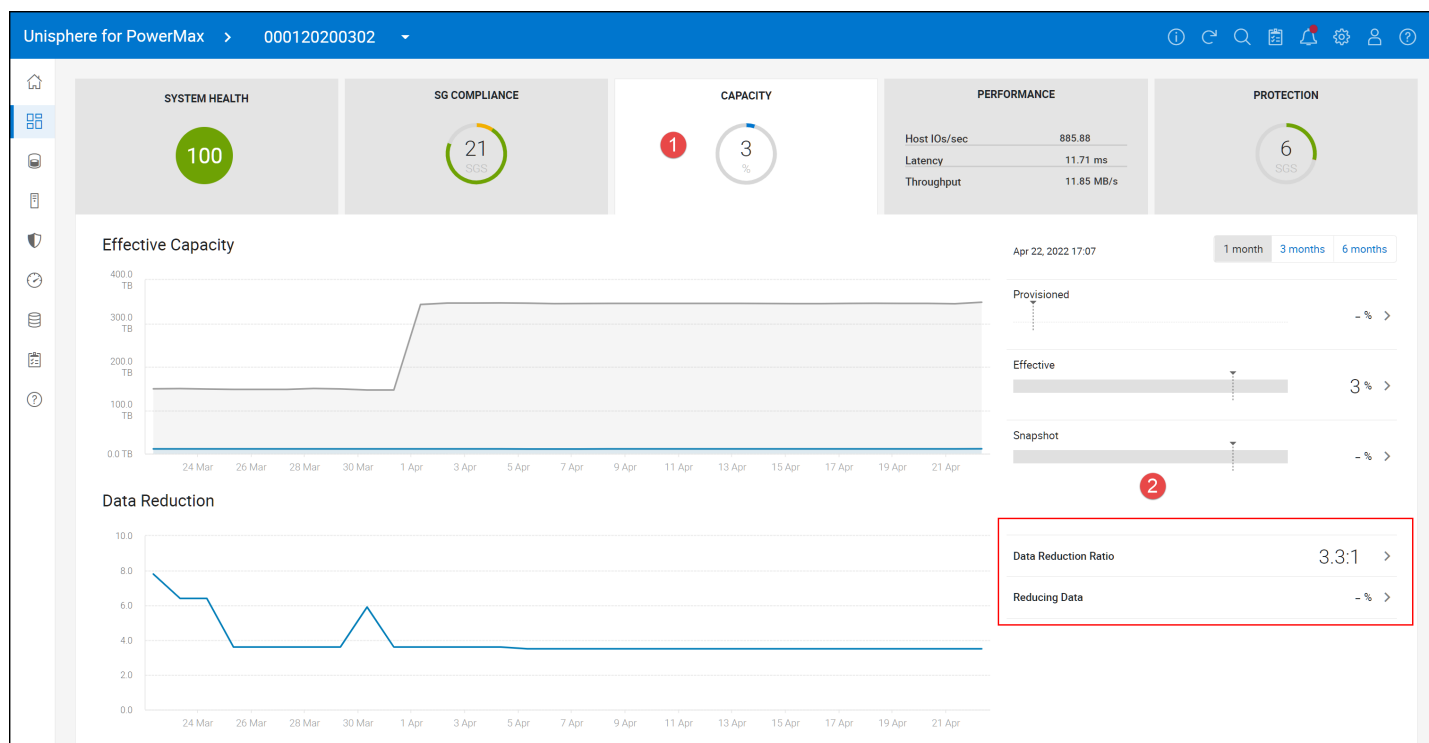


Figure 157. Steps 1-2 for accessing Storage Group Data Reduction information

In the Storage Group Data Reduction screen, the same information as the Storage Group Demand report is available.

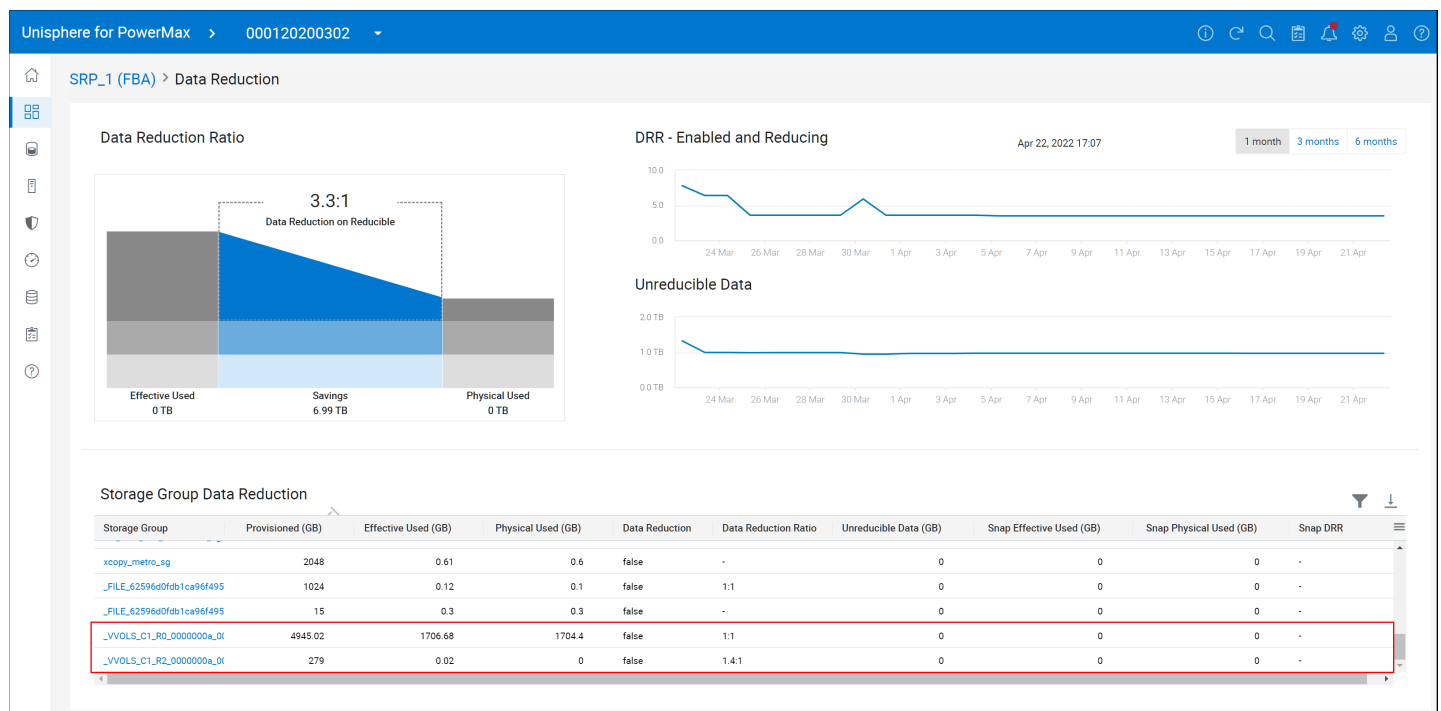


Figure 158. Storage Group Data Reduction information

14.4.3 Solutions Enabler Storage Group Demand Report

If Solutions Enabler is preferred, the following command will produce the same results as the Storage Group Demand Report. The command in Figure 159 is:

```
symcfg -sid xxx list -demand -srp -type sg -detail -gb
```

Note the image below is not from the same array as Figure 156.

```

10.228.246.17 - PuTTY
dsib2017:~ # symcfg -sid 355 list -demand -srp -type sg -detail -gb

STORAGE RESOURCE POOLS

Symmetrix ID : 0001976003

Name : SRP 1
Usable Capacity (GB) : 48283.6
Compression State : Enabled
Data Reduction Ratio : 15.3:1
SRDF DSE Allocated (GB) : 0.0

-----
SG Name                Subscribed    Allocated    Used    Comp    Snapshot    Snapshot    Snapshot
                        (GB)          (GB) (%)      Ratio   Allocated   Used        Comp
                        -----
temp_vcenter            2048.0        29.4    1      22.0  1.5:1      0.0         0.0      -
dsib1131_1132_1139_gk_sg 0.0          0.0    0       0.0   -         0.0         0.0      -
sra_testing_sg_1        410.0        90.5   22      80.9  1.2:1      0.0         0.0      -
sra_testing_sg_2        200.0        0.0    0       0.0   -         0.0         0.0      -
sra_vcenter15_sg_1      400.0        0.0    0       0.0   -         0.0         0.0      -
sra_vcenter15_sg_2      100.0        0.0    0       0.0   -         0.0         0.0      -
sra_vcenter15_sg_3       20.0        0.0    0       0.0   -         0.0         0.0      -
VROTEMP_MV              8.0          0.0    0       0.0  1.0:1      0.0         0.0      -
sra_92_sync_sg          200.0        21.4   10      13.4  1.6:1      0.0         0.0      -
sra_92_async_sg         100.0        3.8    3       2.9  1.6:1      0.0         0.0      -
dsib1131_1132_1139_gk_for_vvol_* 0.0          0.0    0       0.0   -         0.0         0.0      -
dsib1115_116_117_gk_vvol_sg 0.0          0.0    0       0.0   -         0.0         0.0      -
test_metro              100.0        1.9    1       1.4  1.7:1      0.0         0.0      -
INFRA_port_10_11_ig     2048.0        0.0    0       0.0   -         0.0         0.0      -
unisphere_gk            0.0          0.0    0       0.0   -         0.0         0.0      -
test_concurrent          125.0        0.1    0       0.1  1.0:1      0.0         0.0      -
vcf_workload_sg         2048.0        29.8    1      20.0  2.7:1      0.0         0.0      -
VVOLS_C0_R0_0003701e_0003701f 7660.3       1157.9  15     1086.3 1.1:1      0.0         0.0     16.0:1
VVOLS_C0_R1_0003701e_00038e2c 104.0        20.9   20      20.4  1.0:1      2.4         2.1     1.6:1
VVOLS_C1_R0_00038e2e_00038e2f 104.0        39.3   37      37.9  1.0:1      5.3         4.8     1.1:1
VVOLS_C1_R1_00038e2e_00038e30 16.0         3.0    18       1.6  1.9:1      0.0         0.0      -
<not_in_sg>            36336.2      163.2    0     163.2  1.0:1      0.0         0.0      -
-----
Total                    52027.7      1561.3    3     1450.1              7.7        6.9
dsib2017:~ #

```

Figure 159. Generating the Storage Group Demand Report in Solutions Enabler

To view Storage Group Data Reduction, run the same command on the V4 in Figure 160.

```

root@dsib2005:~# symcfg -sid 302 list -demand -srp -type sg -detail -gb

```

STORAGE RESOURCE POOLS

Symmetrix ID : 000120200302

Name : SRP_1
Physical Capacity (GB) : 88695.1
Compression State : Enabled
Data Reduction Ratio : 3.3:1
SRDF DSE Allocated (GB) : N/A

SG Name	Provisioned (GB)	Effective Used (GB)	Used (%)	Physical Used (GB)	Unreducible Used (GB)	Data Reduction Ratio	Snapshot Effective Used (GB)	Snapshot Physical Used (GB)	Snapshot Unreducible Used (GB)	Snapshot Resources Used (%)	Snapshot Reduction Ratio
VVOLs Cl R0 0000000a 0000000b	4945.0	1706.7	34	1704.4	0.0	1.0:1	0.0	0.0	0.0	0.0	-
vm infra 302_0180-0186_sg	24576.0	4258.0	17	1480.7	805.3	5.2:1	0.0	0.0	0.0	0.0	-
vm infra 302_0027-0078_sg	24576.0	734.7	2	275.1	121.2	4.0:1	0.0	0.0	0.0	0.0	-
infra vesam_sg	8192.0	674.4	8	264.8	18.5	2.7:1	0.0	0.0	0.0	0.0	-
unmap_testing_sg	12288.0	165.9	1	56.4	27.5	4.8:1	0.0	0.0	0.0	0.0	-
dsib0027_0049_0051_0078_gk_sg	0.0	0.0	0	0.0	0.0	-	0.0	0.0	0.0	0.0	-
dsib2232_gk_srm_sg	0.0	0.0	0	0.0	0.0	-	0.0	0.0	0.0	0.0	-
dsib1186_1187_infra_sg	2048.0	0.0	0	0.0	0.0	999.9:1	0.0	0.0	0.0	0.0	-
dsib1188_1189_infra_sg	2048.0	0.0	0	0.0	0.0	999.9:1	0.0	0.0	0.0	0.0	-
dsib0180_0182_0184_0186_tcp_sg	2048.0	0.0	0	0.0	0.0	37.5:1	0.0	0.0	0.0	0.0	-
star_sg	15.0	0.1	0	0.0	0.0	6.2:1	0.0	0.0	0.0	0.0	-
srm_migration_test_sg	100.0	0.0	0	0.0	0.0	31.4:1	0.0	0.0	0.0	0.0	-
xcopy_metro_sg	2048.0	0.6	0	0.6	0.0	-	0.0	0.0	0.0	0.0	-
dsib0027_0049_0051_0078_tcp_sg	6144.0	0.1	0	0.1	0.0	1.0:1	0.0	0.0	0.0	0.0	-
clone_demo_sg	200.0	0.1	0	0.1	0.0	1.6:1	0.0	0.0	0.0	0.0	-
test	100.0	0.0	0	0.0	0.0	-	0.0	0.0	0.0	0.0	-
test2	100.0	0.0	0	0.0	0.0	-	0.0	0.0	0.0	0.0	-
FILE_62596d0fdb1ca96f495302604*	15.0	0.3	1	0.3	0.0	-	0.0	0.0	0.0	0.0	-
FILE_62596d0fdb1ca96f495302604*	1024.0	0.1	0	0.1	0.0	1.0:1	0.0	0.0	0.0	0.0	-
concurrent_non_star_sg	100.0	20.9	20	20.9	0.0	1.0:1	0.0	0.0	0.0	0.0	-
VVOLs Cl R2 0000000a 00001848	279.0	0.0	0	0.0	0.0	1.4:1	0.0	0.0	0.0	0.0	-
<not_in_sg>	73356.1	3761.5	5	365.7	19.6	10.8:1	0.0	0.0	0.0	0.0	-
Total	164202.2			4169.2	992.1			0.0	0.0		

```

root@dsib2005:~#

```

Figure 160. Storage Group Data Reduction information

14.5 Queueing

As each ESXi host only has a single PE for all IO, customers may be concerned that queuing could become an issue. A SCSI PE has a default queue depth limit of 128, considerably higher than a normal VMFS or RDM device which defaults to 32 (max of 256), but it does represent all vVols on that host. The question then, is how much IO can that one PE handle?

Note: As NVMe uses a vPE, there is no device queue to adjust.

The following example will illustrate how much IO is possible using the default settings. The PowerMax array is an all-flash array, meaning their latency is very low, usually sub 1 millisecond. For this example, the latency will be rounded to 1 millisecond. If one assumes 1 millisecond this means 1000 IOs can be completed in a single second (1000 milliseconds in 1 second). Therefore with 1 outstanding IO in the queue 1000 IOs can be serviced. Now while the PE queue defaults to 128, if the HBA device queue depth limit is lower than that, and with fibre channel it usually is, it will be the actual queue size the vVols use. Using the previous calculation and assuming an average default HBA device queue depth limit of 64, this means 64,000 IOPS are possible to that one PE. If the HBA queue equaled the default PE queue, it would be twice that, 128,000 IOPS. While that is considerable for a single host, if the customer environment requires more IOPS the following section explains how to adjust the PE queue.

14.5.1 Adjusting the PE queue

The parameter which controls the PE queue is **Scsi.ScsvVolIPESNRO** and is shown in Figure 161.

Edit Advanced System Settings
dsib0180.lss.emc.com

⚠ Modifying configuration parameters is unsupported and can cause instability. Continue only if you know what you are doing.

Scsi.ScsiVVol

Name	Value
Scsi.ScsiVVolPESNRO	256

1 items

Min: 32 Max: 4096

Default schedNumReqOutstanding value for a PE LUN.

CANCEL
OK

Figure 161. Adjusting the PE queue

While the PE queue depth limit can be changed in the Advanced Settings, the new value will not take effect until the next reboot. In addition, it is important to remember that having a PE queue depth limit larger than the HBA device queue depth limit will not make any difference, though there is no concern leaving the PE at the default if it is larger.

If a reboot must be avoided, the PE queue depth limit can be altered through the following command line:

```
esxcli storage core device set -O <number> -d <naa>
```

While the value can be changed to 4096 through CLI, the number cannot be larger than the current device queue limit.

A change to the HBA device queue depth limit will require a reboot, so it may be easier to set the Advanced Parameter for the PE queue to a larger value before considering changing the HBA value, therefore if it is necessary to increase the HBA queue, the PE will already be set. Note that the value of **Scsi.ScsiVVolPESNRO** will not validate against the HBA value upon reboot.

In general, Dell does not recommend making any changes to the queues, even though there is a single Protocol Endpoint. Testing has shown the defaults are adequate for the vast majority of environments. If it is necessary, however, follow the previous instruction.

14.6 PowerMax code upgrade

During a PowerMax code upgrade, the VASA Provider may become unavailable for a short period of time. While this does not impact any ongoing IO to VMs, it may prevent management activity. Any requests made during the upgrade will be retried, but not indefinitely, and may eventually fail. In addition, VMs will go noncompliant, though will resolve to compliancy within ten minutes. Dell recommends close coordination between the VMware and storage administrators when planning for a code load so that VP management activities can be reduced or completely avoided during the upgrade.

14.7 vVol scaling

Although a PowerMax can theoretically support ~64,000 vVols⁷, as the number of VMs and their devices grow, the number of out of band VASA Provider commands increase. This is also compounded when multiple vCenters use the same array. While IO to the VMs itself is not impacted, operations which require the VASA Provider, do. Powering on and off VMs, adding and removing VMs or vmdks, DRS movements or even Storage vMotions are handled through the VASA Provider. When commands that require the VASA Provider are issued in bulk, they can impact the overall performance of the vVol VMs.

Currently Dell does not publish limits of the number of ESXi hosts, VMs, or vVols for non-replicated objects we support. Each customer environment is unique in how they use vVols and how large they can scale will be dependent on a number of factors; however, as noted the number of concurrent VASA Provider commands plays the most critical role in performance. Dell is constantly working with VMware to test and expand scalability.

⁷ Each engine of a PowerMax supports 32,000 devices, therefore a two engine PowerMax is required to reach 64,000.

15 Using Virtual Volumes with VMware Live Site Recovery

15.1 VMware Live Site Recovery

VMware Live Site Recovery⁸ leverages storage array-based replication such as Dell Symmetrix Remote Data Facility (SRDF) to protect virtual machines in VMware vSphere environments. The interaction between VMware SRM and storage array replication is governed by a well-defined set of specifications. These VMware-defined specifications are implemented by the storage array vendor as either a lightweight application referred to as the storage replication adapter (SRA) for VMFS/NFS/RDMs, or within the VASA Provider for virtual volumes.

The Dell embedded VASA Provider (EVASA) enables VMware SRM to interact with a Dell storage environment running vVols. It allows VMware SRM to automate storage-based disaster restart operations on PowerMax arrays in an SRDF configuration. Unlike the SRDF SRA, there is no configuration required outside of SRM for vVols. There are no specialized settings in XML files, or advanced configurations. The main reason for this is that the VASA Provider for SRM only supports SRDF/Asynchronous replication with an RPO of 300 seconds.

The installation and general configuration of SRM will not be covered in this section. Follow the VMware documentation for installing SRM as there is nothing about its configuration that requires changes for vVols. No additional software must be installed (e.g., an SRA) to use SRM with vVols, save for registering the VASA Provider which is necessary to run vVols even without replication. Where necessary, however, specific SRM vVol information will be included below.

15.1.1 VMware PowerCLI

Although it is possible to use VMware PowerCLI to manage the SRM environment, Dell strongly discourages its use. There are known issues and not all SRM GUI operations are replicated in PowerCLI. Internal testing has led to undesirable SRM states which are difficult to recover from.

15.2 Protocols

SRM supports all protocols that the PowerMax offers with vVols:

- Fibre Channel (FC)
- iSCSI
- NVMe over Fabrics (TCP)

15.3 SRM restrictions

VMware and Dell have a number of restrictions when using SRM with vVols which are important to understand prior to configuration.

15.3.1 VMware

Three of the most important restrictions from VMware are:

⁸ Referred to as VMware SRM herein.

- Site Recovery Manager does not support protection of virtual machines that have non-replicated virtual disks with vVols protection groups.
- Site Recovery Manager does not support placing replication groups from different fault domain pairs in the same vVol protection groups. For example, if Array A has VMs replicating to Array B in replication group 1, and other VMs replicating to Array C in replication group 2, there cannot be a single protection group with both RG 1 and RG 2. Each RG would have to be in a different protection group, though they could be in the same recovery plan.
- Site Recovery Manager does not support the protection of virtual machines with different vVols-based disks, replicated by different storage policies or different vVols replication groups.
- vVols does not support the recovery of template virtual machines.

Be sure to review the VMware SRM documentation for more details around other limitations, particularly maximums for object like protection groups and VMs. Also, while bi-directional protection is supported, maximums apply to the sites together not separately.

15.3.2 Dell

Dell limits the following objects to these soft values:

- 250 VMs supported with SRM
 - 2000 vVols – average of 8 vVols per VM; however, the 2000 vVols may allocated as required across VMs
- 25 VASA Replication Groups

For environments that exceed these requirements, please contact your local Dell sales team.

15.4 VM Replicated Objects

The PowerMax will only replicate hard disks and the configuration files, it will not replicate VM snapshots (i.e., array) or the swap (memory) file. Any snapshots (with or without memory) that exist on the source VM will be lost when failing over to the remote vCenter. However, when sizing the remote storage container, be sure to account for both the SnapVX targets as well as the swap file if you plan on powering on the VMs during the test.

Note: Before running a testfailover or failover with vVols, be sure the remote vVol datastore is large enough to account for duplicating any VMs in the Recovery Plan. Both testfailover and failover use target snapshot devices. The Dell implementation of vVols does not failover directly to the remote device (R2), rather the failover is akin to testfailover. New vVols will be created at the time of failover, so though the R2s will be removed in the background, there must be enough space in the datastore to account for them, along with a copy. A Reprotect, therefore, creates new pairs and requires a full synchronization.

15.4.1 Compliancy

When using vVols with SRM, the idea of VM compliancy plays a key role. When the user initially creates a replicated VM, the VM will show as Noncompliant as in [Figure 162](#).

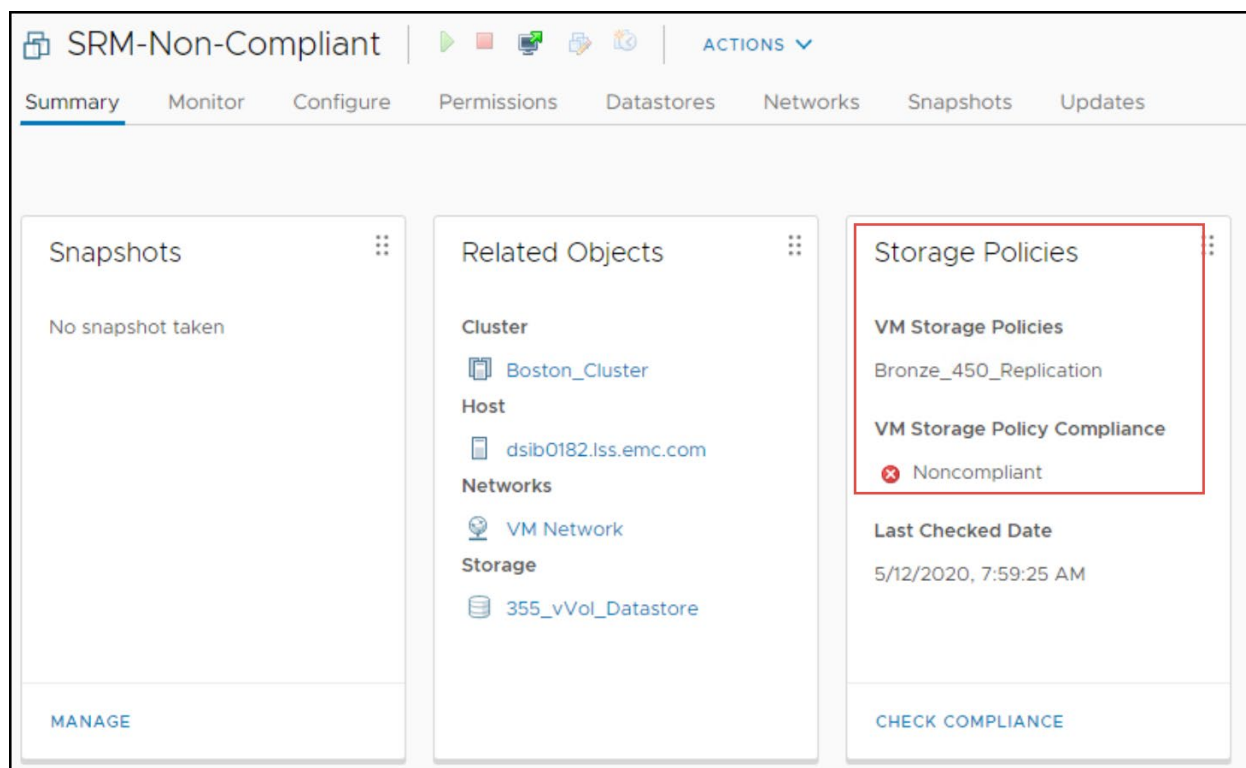


Figure 162. Storage Policy Compliance

This is informing the user that the VM does not meet the conditions set forth in the storage policy, namely that the VM is replicated. A VM will remain Noncompliant until the devices (vVols) making up that VM are fully synched with the remote devices, i.e., that the pair(s) are in a consistent state (SRDF/A). Use the **CHECK COMPLIANCE** link in the bottom of that widget to have VMware re-evaluate compliance. When the VM is Compliant there will be a green check mark as in Figure 163.

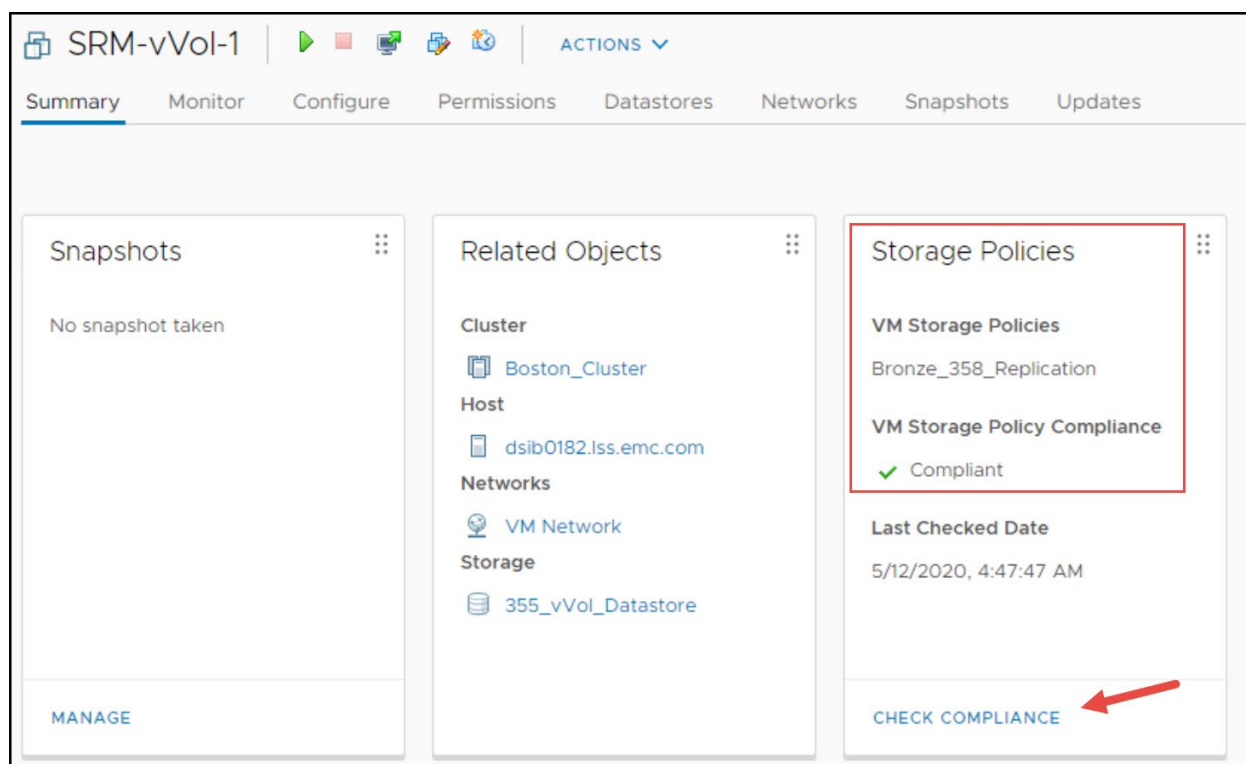


Figure 163. Storage Policy Compatibility – CHECK COMPLIANCE

If there is more than one VM in a Replication Group, all VMs should reach compliancy at the same time. The user should check that VMs are Compliant before attempting any SRM activity as SRM will prevent the user from adding Noncompliant VMs to a Protection Group. Since a VASA Replication Group must failover all devices at once, when new VMs are added to the same replication group, all VMs will go Noncompliant until the new pairs are synchronized. To check multiple VMs for compliancy, it is easiest at the VM Storage Policy level. Navigate to Policies and Profiles and select VM Storage Policies in the left-hand panel. Highlight the desired VM Storage Policy and select the Check Compliance button. All VMs associated with that policy will be checked, and the status will show under the VM Compliance tab in the bottom panel. This is shown in Figure 164.

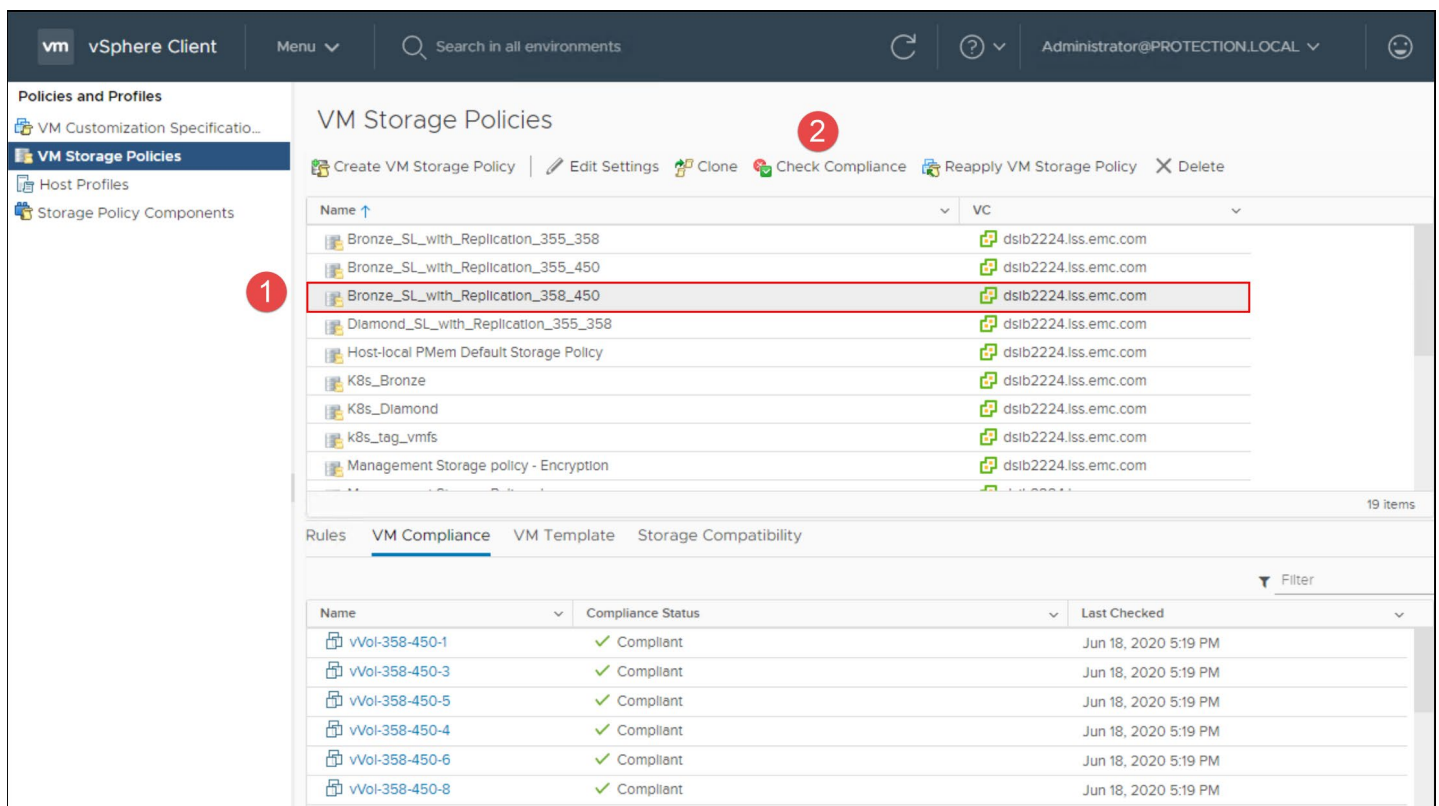


Figure 164. Checking compliancy for multiple VMs

VMs that are in a test state (after testfailover) will always be Noncompliant. It is expected. If a VM policy is viewed, the Replication group will show it is in a test configuration similar to Figure 165.

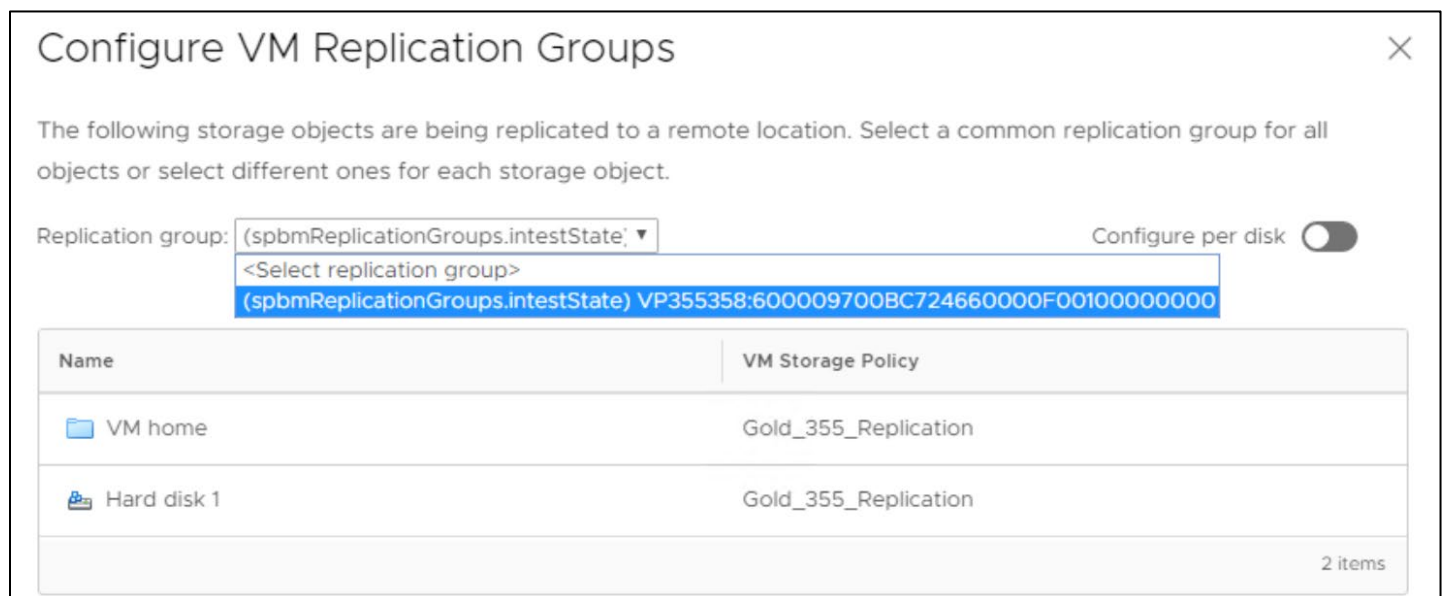


Figure 165. Replication group in a test state

15.5 Configuration

15.5.1 Placeholder datastores and Storage Policy Mappings

Once the vVol environment is configured as detailed above, and a VM created with a replication storage policy, it is possible to use SRM for testing and failover. There are a couple settings in SRM that are specific to vVol testing. Start by pairing the sites and then the mappings for Network, Folder, and Resources. Be sure to create reverse mappings either at the same time through the checkbox, or manually. Next, configure the Placeholder Datastores. When using vVols, Placeholder Datastores cannot be vVol datastores. Any VMFS or NFS datastore can be used, however. Be aware that VMware will not prevent the user from adding a vVol datastore as a placeholder.

The last mapping to complete is Storage Policy Mappings. This is essential with vVols to ensure the failed-over environment runs with the desired service level. Here the user needs to map any storage policies that will be used with VMs and replication, with its counterpart in the recovery site (be sure these are created prior). Dell requires configuring Storage Policy Mappings at both sites, so that SRM can work in a bi-directional manner. The wizard is included here for clarity.

Note: Failure to configure Storage Policy Mappings can result in unpredictable behavior and errors during test and failover. SRM expects the mapping to exist, and if it is not present, may attempt to create it resulting in failure. It is critical, therefore, they are setup before any test/failover or failover is executed in SRM.

Begin within SRM and select the menu option Storage Policy Mappings on the left-hand panel. Next select **NEW** to begin the mapping in Figure 166.

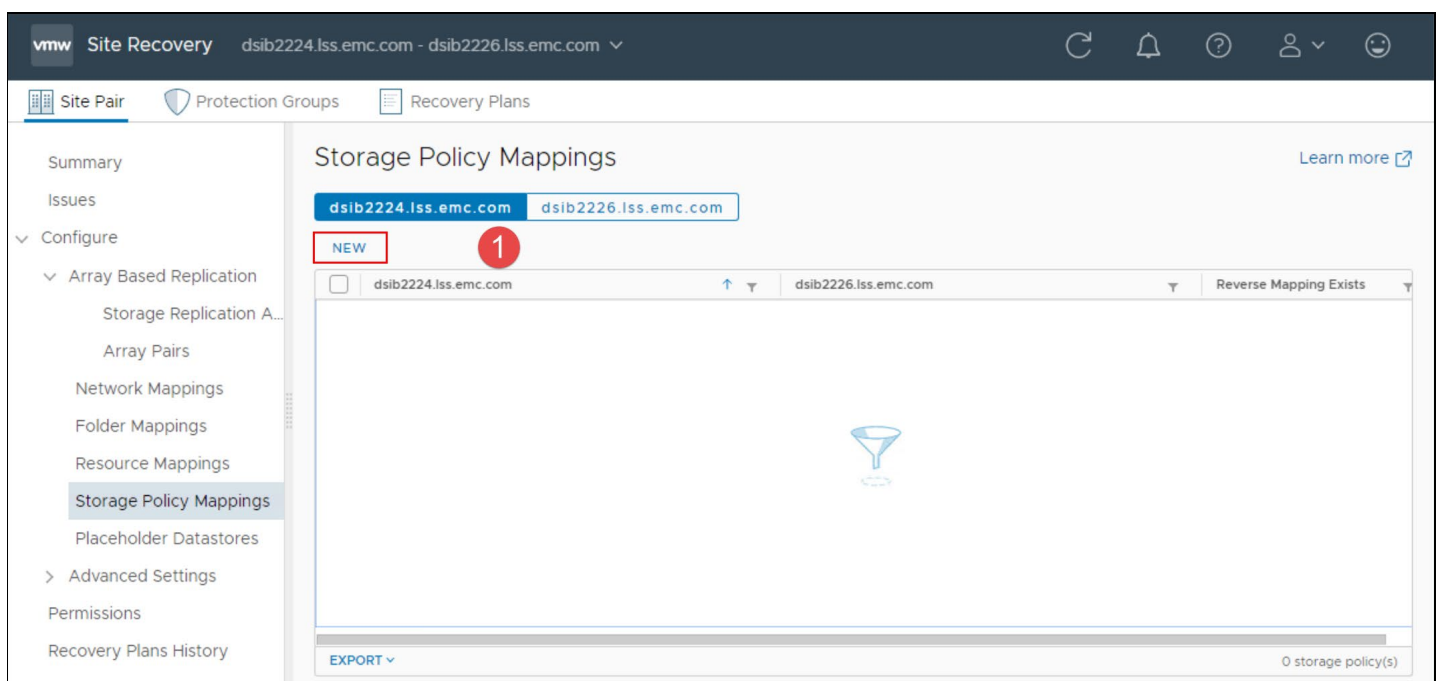


Figure 166. SRM Storage Policy Mappings – step 1

In step 2 in Figure 167 leave the default to automatically prepare the mappings. Select **NEXT**.

New Storage Policy Mappings

- 1 Creation mode
- 2 Recovery storage policies
- 3 Reverse mappings
- 4 Ready to complete

Creation mode

Select the way you want to create mappings.

☒ Automatically prepare mappings for storage policies with matching names
The system automatically prepares mappings for storage policy with matching names under the selected policy containers.

☐ Prepare mappings manually
Manually select which exact storage policies to map.

CANCEL NEXT

Figure 167. SRM Storage Policy Mappings – step 2

In the next screen all the available storage policies will be listed for each site. Using the radio buttons, select the storage policy on the left and its match on the right. For each pairing, select the **ADD MAPPINGS** button to move it to the bottom. Dell recommends using the same storage policy name at each site to make mapping easier. Using the same name also ensures if a mapping is missed in SRM, VMware knows which policy to use at the recovery site. Though Dell recommends using the same service levels at each site, it is not required. If different service levels are used at each site, for instance the recovery site uses Platinum instead of Diamond, it is perfectly acceptable to map a Diamond policy to a Platinum policy instead. Note this could impact performance of the VMs after failover as compared to the protection site. These steps are shown in Figure 168 and Figure 169.

New Storage Policy Mappings

- 1 Creation mode
- 2 Recovery storage policies
- 3 Reverse mappings
- 4 Ready to complete

Recovery storage policies

Configure recovery storage policy mappings for one or more storage policies.

Search...

dsib2224.lss.emc.com

- ☒ Bronze_SL_with_Replication_355_358
- ☐ Diamond_SL_with_Replication_355_358
- ☐ Host-local PMem Default Storage Policy
- ☐ K8s_Bronze
- ☐ K8s_Diamond
- ☐ k8s_tag_vmfs

dsib2226.lss.emc.com

- ☒ Bronze_SL_with_Replication_355_358
- ☐ Diamond_SL_with_Replication_355_358
- ☐ Host-local PMem Default Storage Policy
- ☐ Management Storage policy - Encryption
- ☐ Management Storage Policy - Large
- ☐ Management Storage Policy - Regular

ADD MAPPINGS

dsib2224.lss.emc.com dsib2226.lss.emc.com

0 mapping(s)

CANCEL BACK NEXT

Figure 168. SRM Storage Policy Mappings – step 3

New Storage Policy Mappings

1 Creation mode
2 Recovery storage policies
3 Reverse mappings
4 Ready to complete

Recovery storage policies
Configure recovery storage policy mappings for one or more storage policies.

Search...

dsib2224.lss.emc.com

- ☐ Bronze_SL_with_Replication_355_358 *
- ☐ Diamond_SL_with_Replication_355_358 *
- ☐ Host-local PMem Default Storage Policy
- ☐ K8s_Bronze
- ☐ K8s_Diamond
- ☐ k8s_tag_vmfs

dsib2226.lss.emc.com

- ☐ Bronze_SL_with_Replication_355_358
- ☐ Diamond_SL_with_Replication_355_358
- ☐ Host-local PMem Default Storage Policy
- ☐ Management Storage policy - Encryption
- ☐ Management Storage Policy - Large
- ☐ Management Storage Policy - Regular

ADD MAPPINGS

dsib2224.lss.emc.com	dsib2226.lss.emc.com
Bronze_SL_with_Replication_355_358	Bronze_SL_with_Replication_355_358
Diamond_SL_with_Replication_355_358	Diamond_SL_with_Replication_355_358

2 mapping(s)

CANCEL BACK NEXT

Figure 169. SRM Storage Policy Mappings – step 4

With the mappings configured, in step 5 in [Figure 170](#), check the box so that SRM will automatically create the reverse mappings. Again, this ensures if replication is used bi-directionally, or if a failover and reprotect is executed, the mappings are there.

New Storage Policy Mappings

1 Creation mode
2 Recovery storage policies
3 Reverse mappings
4 Ready to complete

Reverse mappings
Select configured mappings for which to automatically create reverse mappings. This might overwrite existing mappings.

dsib2226.lss.emc.com	dsib2224.lss.emc.com
<input checked="" type="checkbox"/> Bronze_SL_with_Replication_355_358	Bronze_SL_with_Replication_355_358
<input checked="" type="checkbox"/> Diamond_SL_with_Replication_355_358	Diamond_SL_with_Replication_355_358

☒ 2

2 mapping(s)

CANCEL BACK NEXT

Figure 170. SRM Storage Policy Mappings – step 5

In step 6, complete the wizard in [Figure 171](#).

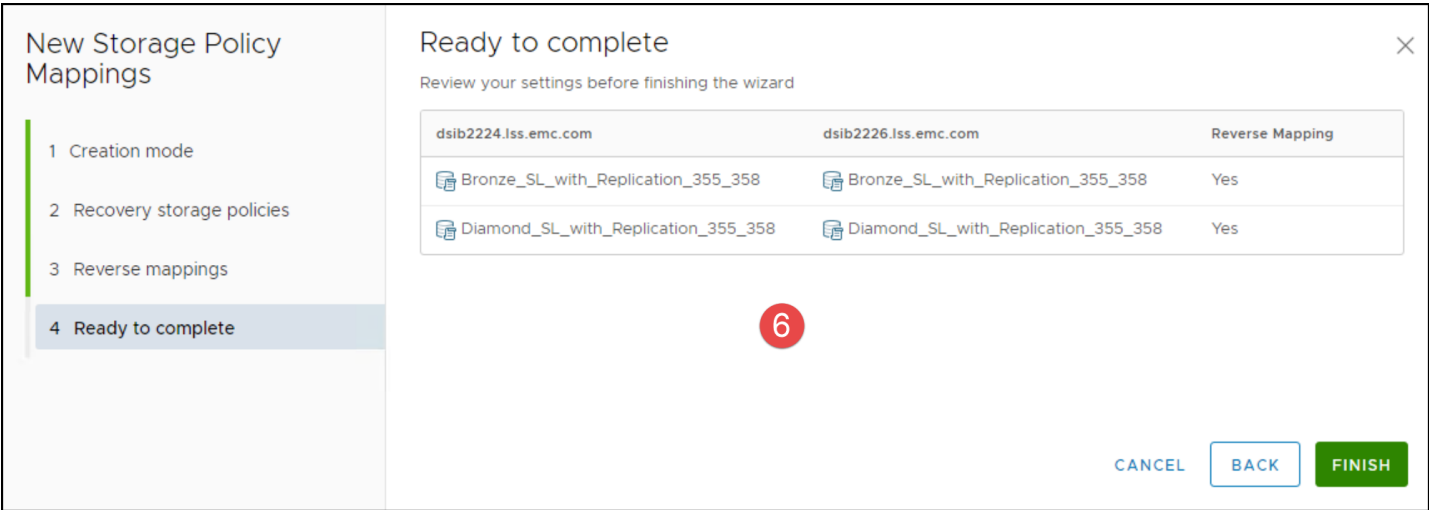


Figure 171. SRM Storage Policy Mappings – step 6

The final policies are seen in Figure 172.

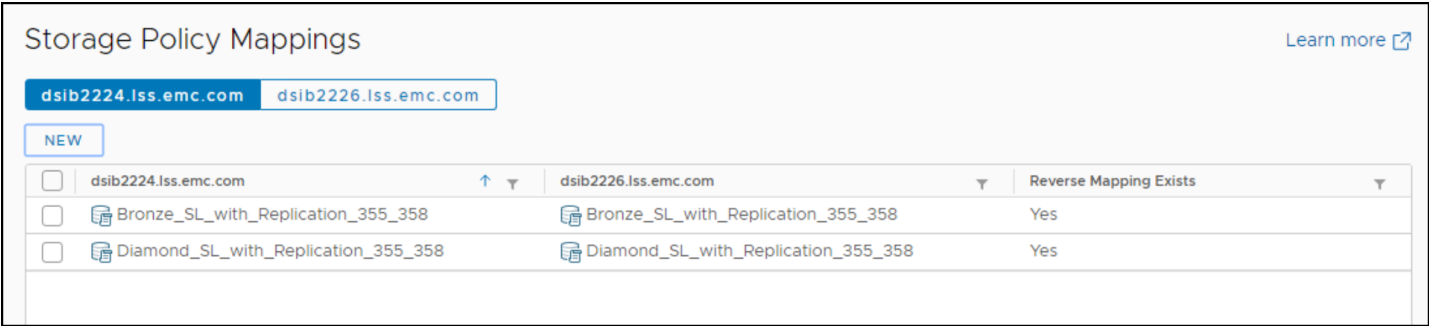


Figure 172. SRM Storage Policy Mappings

If more policies are added in the future, be sure they are added to both sites and then a new mapping created. Failure to add a new storage policy on the recovery site and thus a mapping, can lead to issues during SRM operations.

After the mappings are complete, setup a protection group(s) and recovery plan(s). Remember that with vVols, all management is conducted through the VASA Provider, so there is no Storage Replication Adapter, nor a need to add array managers. The protection site is already aware of the pairings by way of the previously created VASA Replication Group(s).

15.6 Create a Protection Group and Recovery Plan

The following is a quick run-through of setting up a protection group and recovery plan. For more detail, please see the VMware documentation.

First name the group and select Next in Figure 173. Be sure the direction is correct.

New Protection Group

1 Name and direction

Name and direction

Name: vVol Test
71 characters remaining

Description:
4096 characters remaining

Direction:
☒ Boston → London
☐ London → Boston

Location:

 Protection Groups

CANCEL **NEXT**

Figure 173. Create SRM Protection Group – step 1

Next, in [Figure 174](#), select the new Virtual Volumes radio button. Note that despite the fact that Storage Policies are used with vVols, that is not the Type of Protection Group to be used with vVols. The Storage policies Type is used with the SRDF SRA. You will be given an option to choose a Fault Domain. If there is more than one target array, both Fault Domains will be listed. Select the desired one and then Next.

New Protection Group

2 Type

Select the type of protection group you want to create:

☐ Datastore groups (array-based replication)
 Protect all virtual machines which are on specific datastores.

☐ Individual VMs (vSphere Replication)
 Protect specific virtual machines, regardless of the datastores.

☒ Virtual Volumes (vVol replication)
 Protect virtual machines which are on replicated vVol storage.

☐ Storage policies (array-based replication)
 Protect virtual machines with specific storage policies.

Select fault domain.

Fault Domain	Description	Status
<input checked="" type="radio"/> 000197600355	600009700BC724630000F00100000000	✓ OK

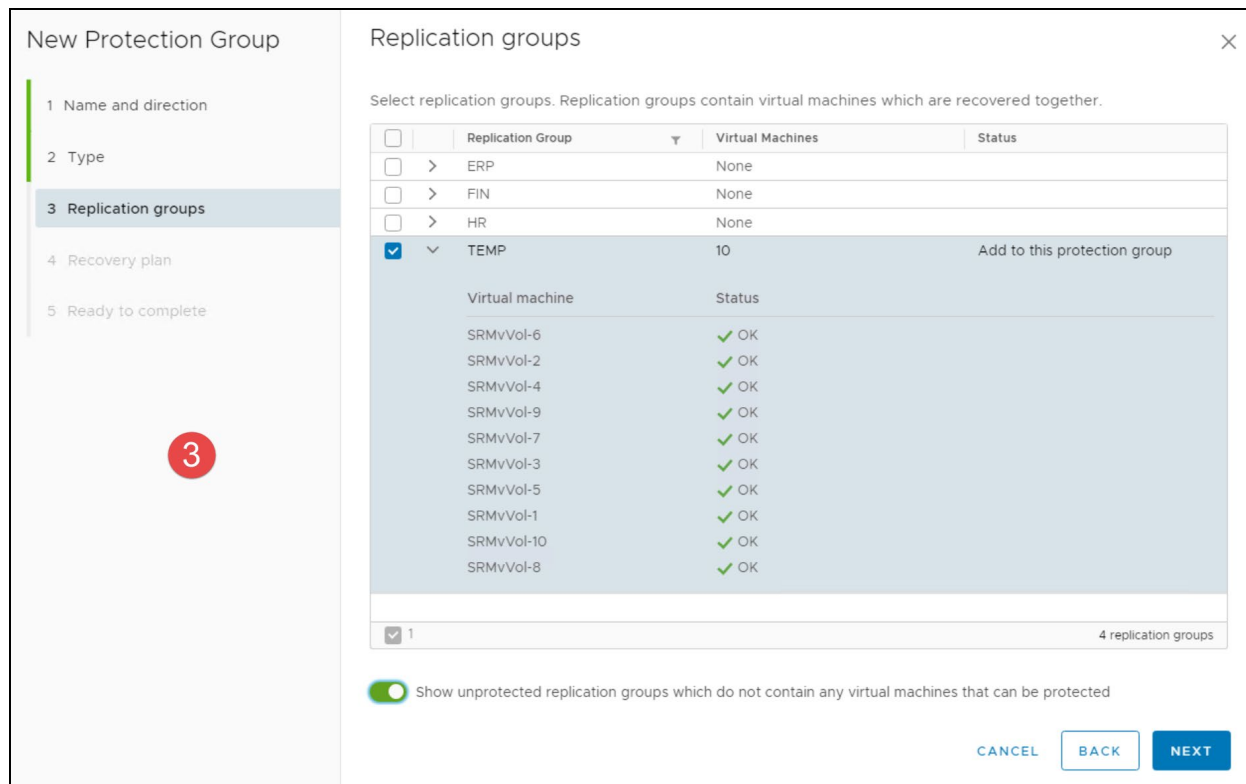
1 domains

CANCEL **BACK** **NEXT**

Figure 174. Create SRM Protection Group – step 2

Next, select the desired VASA Replication Group in [Figure 175](#). By selecting the arrow to the left of the Replication group, any VMs in the VRG will be listed. Note that each Replication Group is a single Failover entity, meaning all VMs are selected. You cannot do a partial Failover. If a VM

needs to be failed over separately, it must be in its own Replication Group. SRM will not list any VMs that are in a Noncompliant state. Then select Next.



New Protection Group

1 Name and direction
2 Type
3 Replication groups
4 Recovery plan
5 Ready to complete

Replication groups

Select replication groups. Replication groups contain virtual machines which are recovered together.

<input type="checkbox"/>	Replication Group	Virtual Machines	Status
<input type="checkbox"/>	> ERP	None	
<input type="checkbox"/>	> FIN	None	
<input type="checkbox"/>	> HR	None	
<input checked="" type="checkbox"/>	▼ TEMP	10	Add to this protection group

Virtual machine	Status
SRMvVol-6	✓ OK
SRMvVol-2	✓ OK
SRMvVol-4	✓ OK
SRMvVol-9	✓ OK
SRMvVol-7	✓ OK
SRMvVol-3	✓ OK
SRMvVol-5	✓ OK
SRMvVol-1	✓ OK
SRMvVol-10	✓ OK
SRMvVol-8	✓ OK

☒ 1 4 replication groups

☒ Show unprotected replication groups which do not contain any virtual machines that can be protected

CANCEL BACK NEXT

Figure 175. Create SRM Protection Group – step 3

Note: If the displayed replication groups represent more than one array (fan-out), do not combine replication groups from different arrays in the same Protection Group. Dell does not support consistency across arrays. SRM will not check that RGs are from the same array and running a Recovery Plan for a Protection Group with multiple arrays will lead to issues that may be very difficult to recover from. The error “Cannot find target replication groups for source replication groups...” will be produced. To resolve, place the replication groups in separate protection groups and then combine them in the recovery plan.

Finally, in [Figure 176](#) type in a new recovery plan name and select Next.

New Protection Group

1 Name and direction
2 Type
3 Replication groups
4 Recovery plan
5 Ready to complete

Recovery plan

You can optionally add this protection group to a recovery plan.

☐ Add to existing recovery plan
☒ Add to new recovery plan
☐ Do not add to recovery plan now

Recovery plan name: vVol Test
71 characters remaining

CANCEL BACK NEXT

Figure 176. Create SRM Protection Group – step 4 – Recovery Plan

Review in Figure 177 and execute Finish.

New Protection Group

1 Name and direction
2 Type
3 Replication groups
4 Recovery plan
5 Ready to complete

Ready to complete

Review your selected settings.

Name	vVol Test
Description	
Protected site	Boston
Recovery site	London
Location	Protection Groups
Protection group type	Virtual Volumes (vVol replication)
Replication groups	TEMP
Total virtual machines	10
Recovery plan	<input type="checkbox"/> vVol Test (new)

CANCEL BACK FINISH

Figure 177. Create SRM Protection Group – step 5

Note that if new VMs are added to an existing VASA Replication Group after the protection group is created, VMware will automatically configure and add the VM to the existing protection group.⁹

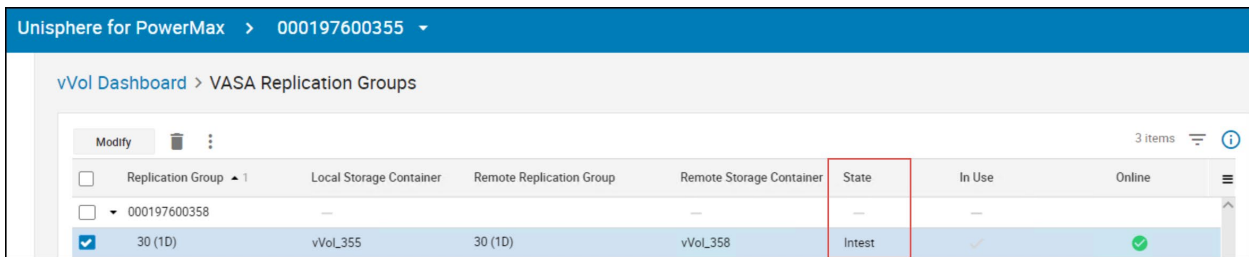
15.7 TestFailover

The PowerMax automatically takes SnapVX snapshots of all the devices in a VASA Replication Group on a 5-minute cycle, always retaining five copies on the recovery array. When the user runs a testfailover on a Recovery Plan, the latest snapshot will be used for that test. The VASA Provider will create new vVol target devices and then link the snapshot to those devices. The production vVols are never impacted. If the user attempts to run a testfailover before there is a viable snapshot

⁹ The opposite, however, does not hold true. If a VM that is in a protection group is deleted, protection must be manually removed for that VM from the protection group.

(not enough time between VM creation and SRM testing), SRM will not warn the user; however, the testfailover will immediately fail with a storage provider error. This failure may not even require a Cleanup, however, please wait another 5 minutes before trying again.

While an SRM test is running, the Replication Group State on the array will change from *Target* to *InTest*, seen in Figure 178.



Replication Group	Local Storage Container	Remote Replication Group	Remote Storage Container	State	In Use	Online
000197600358	—	—	—	—	—	—
30 (1D)	vVol_355	30 (1D)	vVol_358	InTest	✓	✓

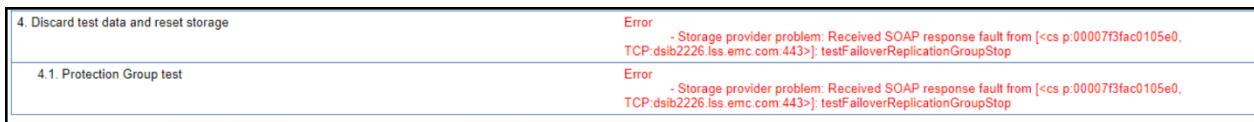
Figure 178. VASA Replication Groups - State

This does not restrict the ability to add more VMs to the same Replication Group if desired. Note that this State will also change for other SRM tasks. A failover will change the state to *FailedOver* while a Reprotect will change to *Source* from *FailedOver*.

15.8 Cleanup

Cleanup runs essentially the same as for VMFS/NFS, with the sole difference that the State of the Replication Group will change back to *Target* from *InTest*.

During Cleanup the user may see the following error in Figure 179.



4. Discard test data and reset storage	Error - Storage provider problem: Received SOAP response fault from [<cs p:00007f3fac0105e0, TCP:dsib2226.iss.emc.com:443>]: testFailoverReplicationGroupStop
4.1. Protection Group test	Error - Storage provider problem: Received SOAP response fault from [<cs p:00007f3fac0105e0, TCP:dsib2226.iss.emc.com:443>]: testFailoverReplicationGroupStop

Figure 179. Cleanup error

This is an expected VMware error indicating some vVols are still bound and is simply a timing issue between the VASA Provider and SRM. The error is innocuous, but it does require that Cleanup is re-run with the **Force** option.

15.9 Failover (Planned Migration or Failover)

A failover with vVols follows a path similar to testfailover with vVols in that the target devices for the failover will be snapshot targets. The actual remote replicated devices (R2s) are not used for the failover. The VASA Provider will use the last consistent snapshot available, just as in a testfailover. This guarantees a consistent copy of the data, within the configured 5-minute RPO. After a failover is executed, the array will begin the process of deleting the original replicated pairs in the background, thus reclaiming space from the storage container. If the protection site is available during this failover, those vVols on that array will also be cleaned up.

15.10 Reprotect

This process again differs from traditional array replication with SRM. Because snapshot targets are used during failover, there is no ability to reverse replication and begin synchronizing the remote copy back to the source copy. Instead, when Reprotect is run, the array will create new vVol RDF pairs and begin synchronizing them. The Reprotect process itself can take a long time depending

on the number and size of devices as it will wait for the pair(s) synchronization to complete. Once the Reprotect finishes successfully, there will still need to be a snapshot taken by the array before another testfailover or failover is run, so it is beneficial to wait at least five minutes before trying. Note that all VMs in the Replication Group should also be compliant before running a testfailover or failover (followed by another Reprotect if desired).

15.10.1 Reprotect error

During a Reprotect an occasional error is expected. This is because of the status calls that are made by the VASA Provider and how VMware responds. These errors can be ignored and will be in the form of the following in [Figure 180](#).




Figure 180. Reprotect error

The Reprotect will still complete successfully.

15.11 Configuration Monitoring

VMware offers a browser-based report that keeps track of changes to the vVol SRM environment. The most pertinent information is seen on the protection site since it includes vVols and VMs, though the recovery site will still show replication groups and fault domains. The URLs to access the monitoring information depend on the SRM platform and are:

- Windows: `https://<IP>:9086/vvol`
- PhotonOS: `https://<IP>/drserver/vvol`

Login with SRM credentials. A sample report is seen in [Figure 181](#).

```

-----
Fault domains      : 1
Date              : 2020-08-17T16:33:55.892451Z
-----
FD1               : 600009700BC724630000F00100000000
                  Name: 000197600355
Provider          : b3129609-9d5c-4458-abf4-3a88c71e7082
                  Name: VASA_0_355
                  Version: 9.2.0.732
                  VASA Version: 3.0
                  VendorID: Dell EMC
                  ModelID: PowerMax
Target domains    : [600009700BC724C20000F00100000000]
Source domains    : [600009700BC724C20000F00100000000]
-----
ArrayID           : VmaxVVolVasaProvider:60000970000197600355F00000000000
Datastore         : datastore-2039
                  Name: 355_vVol
SC                : vvol:600009700bc72463-380101360003701e
                  Name: 355_vVol
Datastore         : datastore-2240
                  Name: Demo_vVol_355
SC                : vvol:600009700bc72463-3801013600038e2e
                  Name: Demo
-----
Replication groups : 9
                  : Name
RG6               : 600009700BC724630000F00100000000:VP30
                  Name: VP30
                  State: SOURCE
                  Peer1: 600009700BC724C20000F00100000000:VP30
                  Device1: naa.600009700BC724633801003600000F32
                  Device2: naa.600009700BC724633801003600000F34
                  Device3: naa.600009700BC724633801003600000F35
                  Device4: naa.600009700BC724633801003600000F36
                  Device5: naa.600009700BC724633801003600000F3A
                  Device6: naa.600009700BC724633801003600000F46
-----
Virtual machines   : 5
VM1               : vm-4035
                  Name: vVol-355-450-1
                  RG: [600009700BC724630000F00100000000:VP30]
                  Device1: naa.600009700BC724633801003600000F32
                  Device2: naa.600009700BC724633801003600000F46
VM2               : vm-4036
                  Name: vVol-355-450-2
                  RG: [600009700BC724630000F00100000000:VP30]
                  Device1: naa.600009700BC724633801003600000F34
                  Device2: naa.600009700BC724633801003600000F48
VM3               : vm-4037
                  Name: vVol-355-450-3
                  RG: [600009700BC724630000F00100000000:VP30]
                  Device1: naa.600009700BC724633801003600000F35
                  Device2: naa.600009700BC724633801003600000F4A
-----

```

Figure 181. vVol monitoring report

16 Conclusion

VMware's storage paradigm, vVol, brings storage management from the VMFS datastore level, down to the virtual machine. Virtual volumes on the PowerMax are individual TDEVs which are mapped directly to a vmdk on a VM. Such granularity permits the VMware administrator the ability to customize a VM according to the services offered by the PowerMax – Service Levels, Data Reduction, and Replication.

Managing storage tiers, provisioning, migrating, cloning virtual machines and correct virtual machine placement in vSphere deployments have become more efficient and user friendly with VASA and vVols. It removes the need for maintaining complex and tedious spreadsheets and validating compliance manually during every migration or creation of a virtual machine or virtual disk.

This white paper discussed how to install, configure and use Virtual Volumes in VMware vSphere 8.x and 9.x environments with PowerMax storage arrays. An understanding of the principles that were exposed here should allow the reader to deploy and utilize VMware vSphere in the most effective manner with vVols.

A Appendix: Virtual Volume Troubleshooting

Before attempting any troubleshooting, refer to the VASA Provider Release Notes for all current limitations in case the issue is known. The following sections discuss varied issues that have been seen in both lab and customer environments. They are presented below in no particular order.

A.1 Directories

Any directory created in a vVol datastore is a config vVol. This means that a directory is automatically 255 GB in size and cannot be changed. On the PowerMax it is possible to create up to a 16 TB directory. Unfortunately, this capability has not been incorporated into any UI or even PowerCLI as of the initial release. In order to create the directory of a size larger than the default, the Managed Objects (MOB) interface, accessible through **https ://<VC_SERVER>/mob**, must be used. The following is a walkthrough of using the MOB interface to create the directory. Note that there are many navigation paths to the same location in MOB and this will only cover one possible option.

A.1.1 Creating a config vVol greater than the default

The following vVol datastore in [Figure 182](#), **CNS_vVol**, is where the new config vVol will be created. Part of the process of creating the directory requires knowing the reference id of the datastore. This id will also be obtained from MOB.

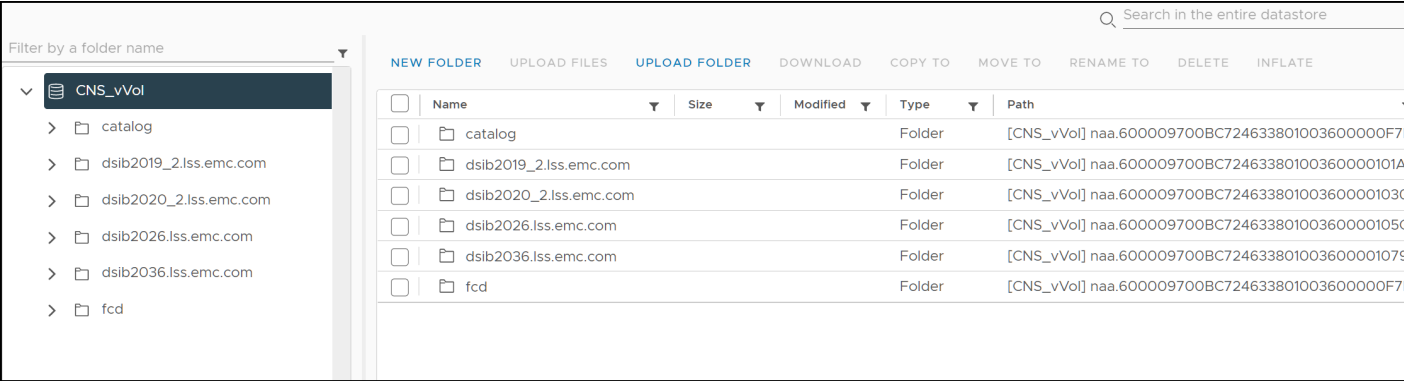


Figure 182. CNS_vVol vVol datastore

Navigate to **https ://<VC_SERVER>/mob** and login as the administrator in [Figure 183](#).

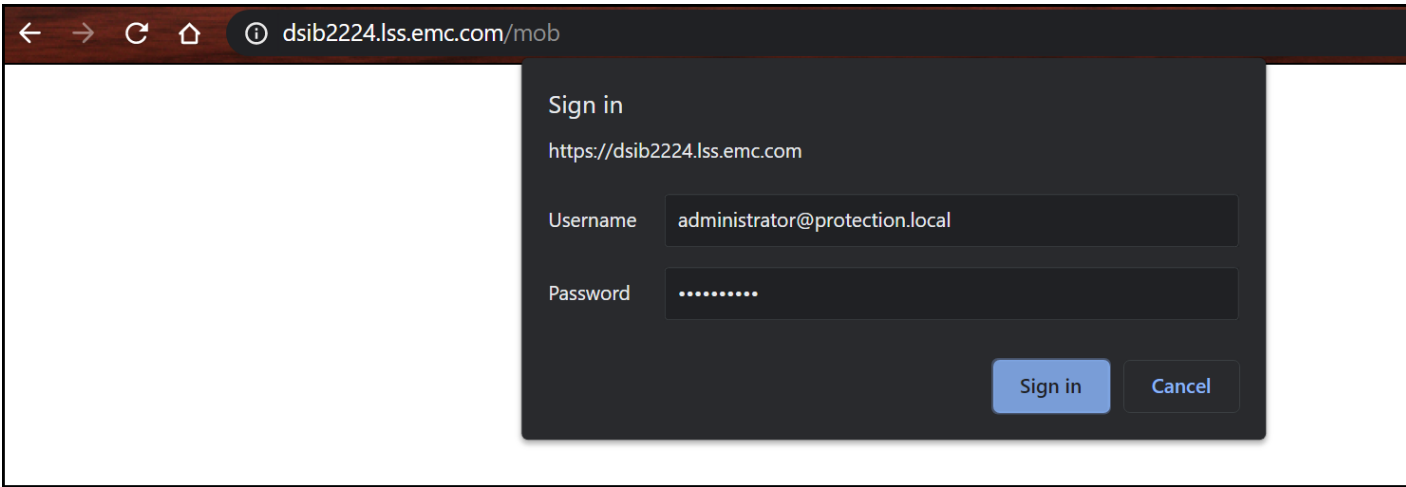


Figure 183. Accessing MOB

Select **RetrieveServiceContent** from under the **Methods** table in Figure 184.

[Home](#)[Logout](#)

Managed Object Type: ManagedObjectReference:ServiceInstance
Managed Object ID: ServiceInstance

Properties

NAME	TYPE	VALUE
capability	Capability	capability
content	ServiceContent	content
serverClock	dateTime	"2021-03-17T19:54:40.444736Z"

Methods

RETURN TYPE	NAME
dateTime	CurrentTime
HostVMotionCompatibility[]	QueryVMotionCompatibility
ServiceContent	RetrieveServiceContent
ProductComponentInfo[]	RetrieveProductComponents
Event[]	ValidateMigration

Figure 184. RetrieveServiceContent

Then select Invoke Method in Figure 185.

Managed Object Type:
ManagedObjectReference:ServiceInstance
Managed Object ID: **ServiceInstance**
Method: **RetrieveServiceContent**

ServiceContent RetrieveServiceContent

Parameters

NAME	TYPE	VALUE
<input type="text"/> <input type="button" value="Invoke Method"/>		

Figure 185. Invoke Method

In the next screen, scroll down to the **rootFolder** in the **Name** column. In Figure 186 select the value on the right, in this example it is named **group-d1 (Datacenters)**.

ioFilterManager	ManagedObjectReference:IoFilterManager	IoFilterManager
ipPoolManager	ManagedObjectReference:IpPoolManager	IpPoolManager
licenseManager	ManagedObjectReference:LicenseManager	LicenseManager
localizationManager	ManagedObjectReference:LocalizationManager	LocalizationManager
overheadMemoryManager	ManagedObjectReference:OverheadMemoryManager	OverheadMemoryManager
ovfManager	ManagedObjectReference:OvfManager	OvfManager
perfManager	ManagedObjectReference:PerformanceManager	PerfMgr
propertyCollector	ManagedObjectReference:PropertyCollector	propertyCollector
rootFolder	ManagedObjectReference:Folder	group-d1 (Datacenters)
scheduledTaskManager	ManagedObjectReference:ScheduledTaskManager	ScheduledTaskManager
searchIndex	ManagedObjectReference:SearchIndex	SearchIndex
serviceManager	ManagedObjectReference:ServiceManager	ServiceMgr
sessionManager	ManagedObjectReference:SessionManager	SessionManager
setting	ManagedObjectReference:OptionManager	VpxSettings
siteInfoManager	ManagedObjectReference:SiteInfoManager	SiteInfoManager
snmpSystem	ManagedObjectReference:HostSnmpSystem	SnmpSystem
storageQueryManager	ManagedObjectReference:StorageQueryManager	StorageQueryManager

Figure 186. rootFolder

In the subsequent screen, in the third row of the table, select the datacenter. In Figure 187, this example is **datacenter-3 (Boston)**.

[Home](#)
[Logout](#)

Managed Object Type: **ManagedObjectReference:Folder**
Managed Object ID: **group-d1**

Properties

NAME	TYPE	VALUE
alarmActionsEnabled	boolean	true
availableField	CustomFieldDef[]	availableField[201] CustomFieldDef
childEntity	ManagedObjectReference:ManagedEntity[]	datacenter-3 (Boston)
childType	string[]	"vim.Folder" "vim.Datacenter"
configIssue	Event[]	
configStatus	ManagedEntityStatus	"gray"
customValue	CustomFieldValue[]	
declaredAlarmState	AlarmState[]	declaredAlarmState["alarm-1.group-d1"] AlarmState declaredAlarmState["alarm-10.group-d1"] AlarmState declaredAlarmState["alarm-100.group-d1"] AlarmState declaredAlarmState["alarm-101.group-d1"] AlarmState

Figure 187. Datacenter selection

Next, expand the **datastore** name row and scroll down to the selected vVol datastore in Figure 188. Note down or highlight and copy the name of the object. This is the reference id. In this example it is **datastore-10011** and is required for the next step.

Home

Logout

Managed Object Type: **ManagedObjectReference:Datacenter**

Managed Object ID: **datacenter-3**

Properties

NAME	TYPE	VALUE
alarmActionsEnabled	boolean	true
availableField	CustomFieldDef[]	availableField[201] CustomFieldDef
configIssue	Event[]	
configStatus	ManagedEntityStatus	"gray"
configuration	DatacenterConfigInfo	configuration
customValue	CustomFieldValue[]	
datastore	ManagedObjectReference:Datastore[]	<div> <div>datastore-9009 (INFRA_357_DEV_24)</div> <div>datastore-17 (355_DS_1)</div> <div>datastore-18 (355_DS_2)</div> <div>datastore-19 (355_share)</div> <div>datastore-20 (SRA_92_VMFS_DS_1)</div> <div>datastore-21 (SRA_92_VMFS_DS_2)</div> <div>datastore-22 (SRA_92_VMFS_DS_3)</div> <div>datastore-23 (SRA_92_VMFS_DS_4)</div> <div>datastore-13005 (VeeamBackup_dsib2014.lss.emc.com)</div> <div>datastore-10010 (355_vVol)</div> <div>datastore-10011 (CNS_vVol)</div> <div>datastore-14036 (355_Demo)</div> <div>datastore-14037 (test)</div> </div>

Figure 188. Datastore selection

Now return to the home screen in Figure 184 and select **content** in the **Properties** table. Next, select **DatastoreNamespaceManager** then **CreateDirectory** in Figure 189 and Figure 190.

[Home](#)
[Logout](#)

Data Object Type: ServiceContent
 Parent Managed Object ID: **ServiceInstance**
 Property Path: **content**

Properties

NAME	TYPE	VALUE
about	AboutInfo	about
accountManager	ManagedObjectReference:HostLocalAccountManager	Unset
alarmManager	ManagedObjectReference:AlarmManager	AlarmManager
authorizationManager	ManagedObjectReference:AuthorizationManager	AuthorizationManager
certificateManager	ManagedObjectReference:CertificateManager	certificateManager
clusterProfileManager	ManagedObjectReference:ClusterProfileManager	ClusterProfileManager
complianceManager	ManagedObjectReference:ProfileComplianceManager	MoComplianceManager
cryptoManager	ManagedObjectReference:CryptoManagerKmp	CryptoManager
customFieldsManager	ManagedObjectReference:CustomFieldsManager	CustomFieldsManager
customizationSpecManager	ManagedObjectReference:CustomizationSpecManager	CustomizationSpecManager
datastoreNamespaceManager	ManagedObjectReference:DatastoreNamespaceManager	DatastoreNamespaceManager
diagnosticManager	ManagedObjectReference:DiagnosticManager	DiagMgr
dvSwitchManager	ManagedObjectReference:DistributedVirtualSwitchManager	DVSwitchManager

Figure 189. DatastoreNamespaceManager

[Home](#)
[Logout](#)

Managed Object Type: ManagedObjectReference:DatastoreNamespaceManager
 Managed Object ID: **DatastoreNamespaceManager**

Properties

NAME	TYPE	VALUE
------	------	-------

Methods

RETURN TYPE	NAME
string	ConvertNamespacePathToUuidPath
string	CreateDirectory
void	DeleteDirectory

Figure 190. CreateDirectory

Finally, fill in the appropriate values into the boxes of the **Parameters** table. Note the placeholder for the **datastore** reference id in Figure 191 is MOID. This will be replaced with the reference id collected in Figure 188. The **displayName** is the directory name. While VMware provides an option for the policy, Dell does not support putting a value here. Any attempt to use it will fail. Instead, the config vVol will be created with the default policy, or in the lowest performant service level (storage resource). Note if that storage resource is full, the next lowest will be used. The last property, size, is in MB, in this case 1 TB. Select **Invoke Method**.

Managed Object Type: **ManagedObjectReference:DatastoreNamespaceManager**
Managed Object ID: **DatastoreNamespaceManager**
Method: **CreateDirectory**

string CreateDirectory

VALUE

<datastore type="Datastore">MOID</datastore>

Parameters

NAME	TYPE	VALUE
datastore (required)	ManagedObjectReference:Datastore	<datastore type="Datastore">datastore-10011</datastore>
displayName (optional)	string	1TB_Config_vVol
policy (optional)	string	
size (optional)	long	1048576

Invoke Method

Figure 191. Directory creation

If the task was successful, VMware will indicate the mobility safe ID of the vVol in Figure 192.

Managed Object Type: **ManagedObjectReference:DatastoreNamespaceManager**
Managed Object ID: **DatastoreNamespaceManager**
Method: **CreateDirectory**

string CreateDirectory

Parameters

NAME	TYPE	VALUE
datastore (required)	ManagedObjectReference:Datastore	<datastore type="Datastore">datastore-10011</datastore>
displayName (optional)	string	1TB_Config_vVol
policy (optional)	string	
size (optional)	long	1048576

Invoke Method

Method Invocation Result: string

NAME	TYPE	VALUE
name	string	"Return value"
val	string	"/vmfs/volumes/vvol:600009700bc72463-3801013600079d2d/naa.600009700BC7246338010036000010DB"

Figure 192. Successful directory creation

In Figure 193 is the newly created directory, along with a 12 GB uploaded ISO.

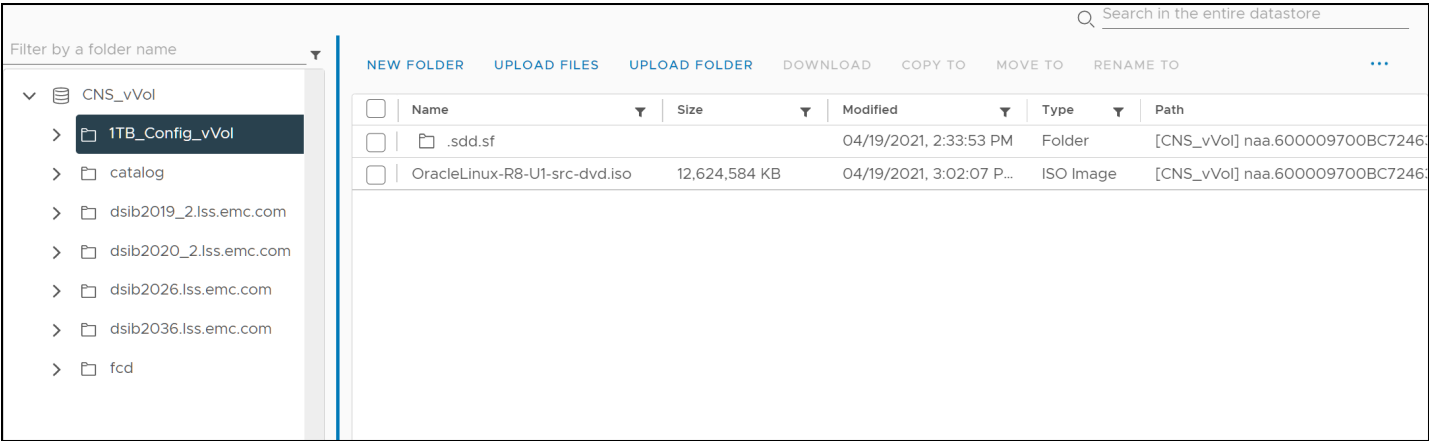


Figure 193. Newly created directory in vSphere

A.2 VASA Provider registration issues

In addition to the common invalid user/password when registering, there are a couple common errors that prevent VASA Provider registration.

A.2.1 Clock synchronization

VMware relies on synchronization of the clocks of ESXi, vCenter and the VASA Provider (EVASA). A time skew between them will result in the following error in Figure 194 when registering the VASA Provider.

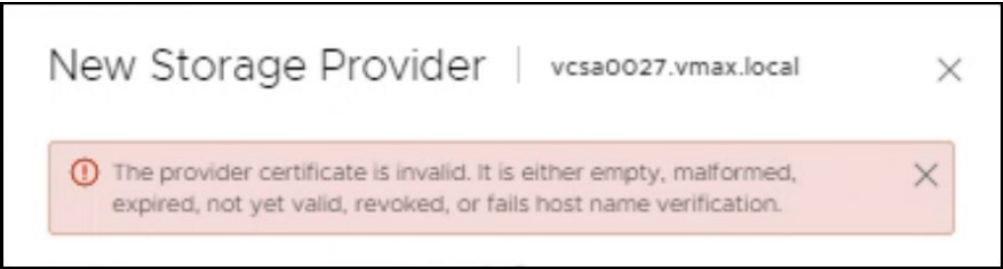


Figure 194. Failure to register VASA Provider due to time skew

The VASA Provider clocks should already be properly synchronized to an NTP server as this is handled during the installation of PowerMax. This information is available in the embedded Unisphere instance by navigating to **Serviceability -> Serviceability -> NTP Server** and checking the field is filled in. This is shown in Figure 195.

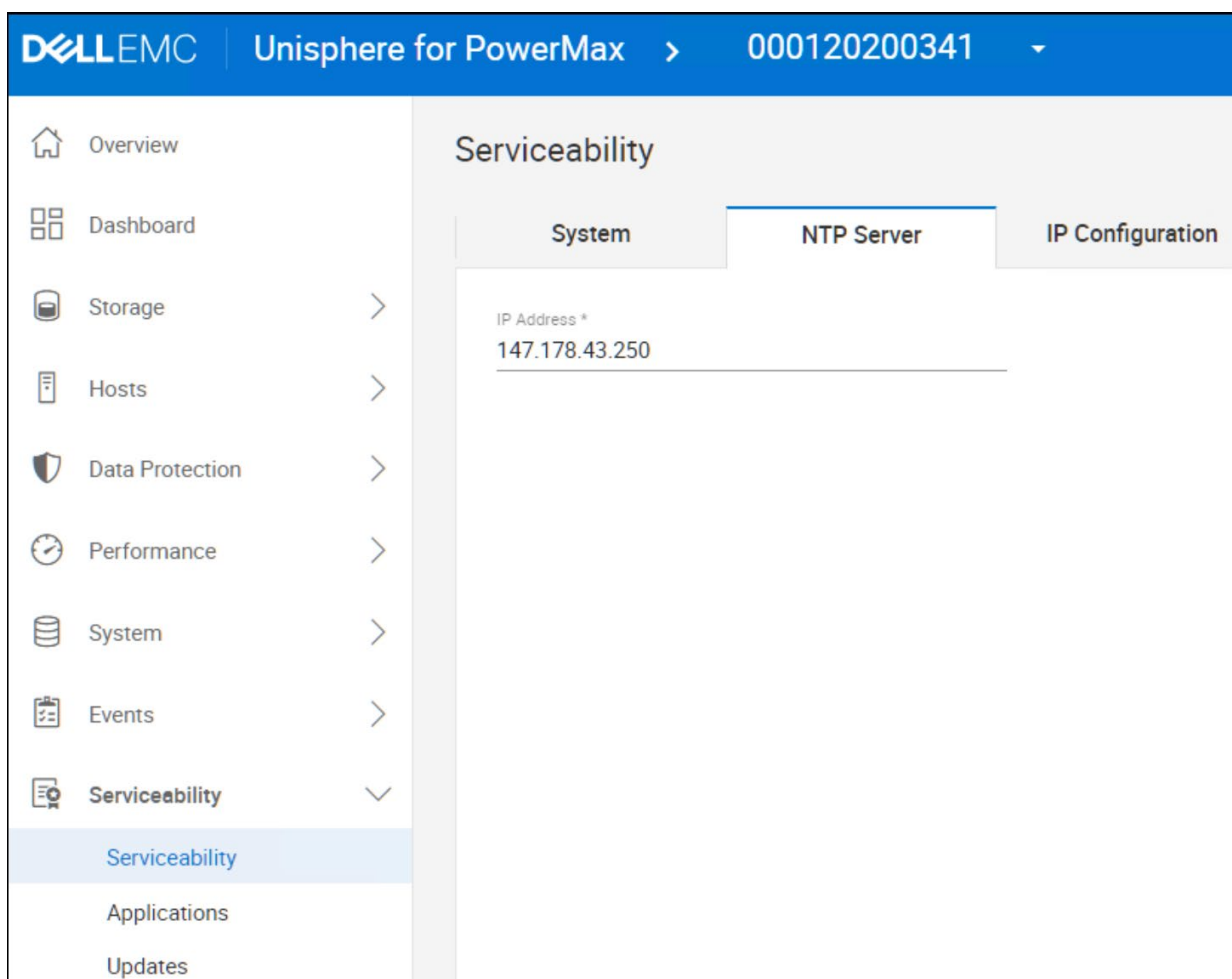


Figure 195. NTP server in embedded Unisphere

The ESXi server(s) and vCenter should be checked by logging in via SSH and running the `date` command. It is always preferable to use the same NTP server for all three components if the networks allow it. The clock times need to be within seconds of each other. Note that the time zone is not relevant.

Note: The invalid provider certificate error may also appear if a new self-signed certificate is generated and the user does not modify the Serial Number to a non-zero value.

A.2.2 Firewall

If the registration is simply rejected, be sure that the firewall between the vCenter appliance and the array is open on port 5989. In addition, each ESXi host must also be able to communicate with the VASA Providers.

A.3 Cascaded Groups/Multi-host initiator groups

As noted in the paper, currently there is no support for cascaded auto-provisioning groups for the Protocol Endpoint. If the PE is part of a masking view with cascaded objects, there will be various issues with the implementation on the vCenter. Neither Unisphere for PowerMax nor Solutions

Enabler will prevent the user from provisioning the PE with cascaded groups so it is important to check this first if issues are encountered.

In addition to the impermissibility of cascaded initiator groups, an initiator group used in a PE masking view may not contain initiators from multiple hosts. Each ESXi host in a cluster (for example) must have a unique PE. To avoid any issue, it is advisable to have a separate initiator group for each ESXi host. Note, however, it is still possible to have a parent initiator group for use in non-PE masking views.

A.4 vVol datastore access issues

Since the creation of a vVol datastore only requires the VASA Provider registration and array communication with the vCenter, there are a couple of error states the user may experience with the datastore. They are covered here.

A.4.1 vVol datastore reports inaccessible

After creation, if a vVol datastore reports as “inaccessible” (usually a few minutes after creation), this is indicative of a protocol endpoint (PE) issue. This will be one of two situations:

- No PEs are provisioned to the ESXi host(s). Note that if only some of the ESXi hosts have a PE, the others will show the datastore as inaccessible.
- A single PE is provisioned to more than one ESXi host. This is typical when customers use a parent initiator group for the entire ESXi cluster and provision only a single masking view/PE. Parent initiator groups are not supported. Each ESXi host must see its own unique PE and have its own masking view. The number of PE masking views should always be the same as the number of hosts in the cluster.

A.4.2 vVol shows as zero bytes

After creation, a vVol datastore reports as 0 (zero) bytes. In addition, if HA is enabled on the ESXi cluster, there will be HA configuration errors because VMware will be unable to create a folder/file on the vVol datastore. A zero-byte datastore is a communication issue between the ESXi hosts and the array VASA Providers. Examine the `/var/log/vvold.log` file and look for entries similar to the following, where the name **vasaip** is replaced with the IP of the VASA Provider:

```
esx.host.fqdn VVold: info vvold[135445092] [Originator@6876 sub=Default]
Initialize: Failed to establish connection https://vasaip:5989/vasa-providers.xml
[vasaip]
esx.host.fqdn VVold: warning vvold[135445094] [Originator@6876
sub=IO.Connection] Failed to connect; <io_obj p:0x0000007653d64df8, h:13,
<TCP 'esxip : 23145'>, <TCP 'vasaip : 5989'>>, e: 110(Connection timed out),
duration: 75005msec
```

Failure to connect is usually a firewall issue. Each ESXi host must be able to communicate with both VASA Providers. Add new firewall rules to open the ports between the hosts and VASA Providers. Then refresh the datastore to retrieve the actual datastore size. Also if HA was enabled, it should be disabled and re-enabled to clear the configuration error.

A.5 vCenter failure to recognize replication or service level

There is a known issue where the vCenter does not recognize the vVol datastore is capable of replication though the array steps have been completed (VASA replication group exists). If rescanning the VASA Provider does not resolve the issue, stop and restart the VASA Providers.

A.6 VASA Provider recovery

If both VASA Providers become unavailable, contact Dell support immediately for assistance.

A.7 Orphaned virtual volumes

While unlikely, it is possible that a vVol fails to be removed when a VM is deleted. It is important to understand that this is not the same as removing a vVol (vmdk) from a VM but choosing not to select the box to delete the file. In that case, it is possible to manually delete the vmdk from the datastore, or to add the vmdk back to a VM. The case of an orphaned vVol means that the vmdk is deleted, but the backing vVol on the array is not. Therefore, no vmdk exists to delete. As a Storage Resource cannot be deleted while a vVol(s) exists, this is the most common way a user might discover an orphaned vVol. As this is an unexpected condition, there is currently no process for a user to remove an orphaned vVol. If such a condition arises, please contact Dell support who can resolve the issue through an internal process.

A.8 Changing Storage Policies with multi-writer flag

When changing the Storage Policy on VMs that have shared vmdks (multi-writer), the reconfigure may fail with the error in [Figure 196](#).




Figure 196. Changing Storage Policies with multi-writer flag

This appears to be a VMware issue. The workaround is to change the Storage Policy for non-shared vmdks first, then do a second reconfiguration for the shared vmdks. Change the VMs in sequence, one at a time, not in parallel.

A.9 SRM TestFailover errors

Errors occurring during SRM testfailover are not always straightforward and do not provide the exact issue encountered. A generic VASA error will take the following form in [Figure 197](#).




Figure 197. Generic VASA error

The actual errors encountered, however, are not generally fatal. Two of the most common ones would be:

- No snapshots yet exist for the replication group being tested (noted previously)
- There is not enough space in the remote storage container

The first issue would only occur if a testfailover is attempted immediately after creating the VMs and associated protection group/recovery plan or reprotect. If this is suspected, run a cleanup operation and wait some time before trying again. Be sure the replication status of the RDF pairs is “Consistent”, indicating that a snapshot is possible.

The second issue may not be as clear since the VASA Provider will not attempt to create SnapVX target devices unless there is enough space available. This means the remote storage container will not be full. Instead, the size of the VMs in the recovery plan would need to be summed and then deducted from the available space in the container to determine if this is the issue experienced in the testfailover. Note that the VASA Provider does not add the size of the memory to the total space required. Therefore, a testfailover could succeed the recovery plan to the point the VMs are powered on, at which point the lack of space error would be reported by VMware.

A.10 SRM Failover

The errors experienced during SRM failover are obviously more problematic than a test. Proper setup will avoid them, but even small misses can cause problems. For example, this error in Figure 198 occurred during a failover. It is another generic error because the storage issue is unknown.

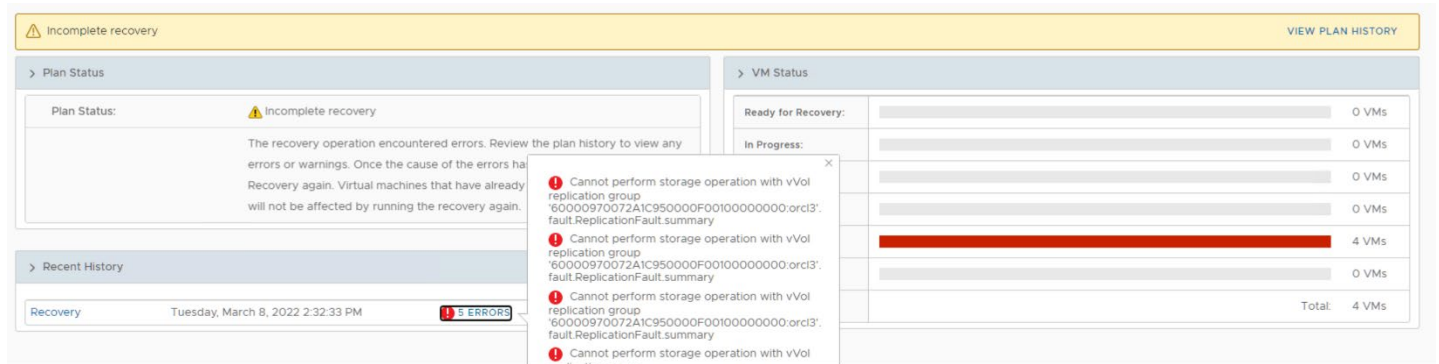


Figure 198. SRM failover error

Careful review of the logs revealed that though the setup appeared to be done correctly, the recovery site did not have the Service Level in the vVol container that was being used in the Storage Policy. When the mappings are setup in SRM, SRM does not validate that there is a vVol container that satisfies a Storage Policy before executing the failover and it is unable to tell the VASA Provider to use a different Service Level.

A.11 VASA Provider rescan errors

There are situations that may arise when the primary VASA Provider experiences a rescan error. This does not indicate the provider is down, so the standby VASA Provider does not take over. An example is shown in [Figure 199](#).

The screenshot shows the VMware vSphere Client interface for a host named 'dsib2226.lss.emc.com'. The 'Configure' tab is selected, and the 'Storage Providers' section is expanded. A table lists the storage providers, with one provider, 'VASA_1_450', showing a 'Rescan Error' status. Below the table, the 'General' tab for this provider is selected, displaying details such as the provider name, status, active/standby status, activation method, URL, provider version, VASA API version, default namespace, provider ID, and supported profiles.

Storage Provider/Storage Sys...	Status	Active/Stand...	Priority	URL	Last Rescan Time	VASA API Version	Certific...
▲ VASA_0_450	Online	--	--	https://10.228.2...	08/03/2020, 8:...	3.0	299 day
000197600450 (1/2 online)		Standby	255				
▲ VASA_1_450	Rescan Error (details...)	--	--	https://10.228.2...	08/12/2020, 3:1...	3.0	299 day
000197600450 (1/2 online)		Active	255				
▲ IOFILTER Provider dsib0184.lss...	Online	--	--	https://dsib018...	08/12/2020, 3:3...	1.5	1694 da

General | Supported Vendor IDs | Certificate Info

Provider name	VASA_1_450
Provider status	Rescan Error
Active/standby status	--
Activation	Explicit
URL	https://10.228.244.240:5989/vasa-providers.xml
Provider version	9.2.0.732
VASA API version	3.0
Default namespace	VmaxVVolVasaProvider
Provider ID	VASA2_000197600450_VASA-1
Supported profiles	Storage Profile Based Management ReplicationProfile

Figure 199. Rescan error with VASA Provider

Unfortunately, the rescan error renders the vVol datastore inaccessible. Running VMs continue operation (IO), but all management activities are unavailable. VMware has published a number of different KB articles on this which are summarized here: <https://kb.vmware.com/s/article/57863>. The best way to resolve this is to unregister the provider that has the rescan error at which point the standby will take over and the vVol datastore(s) will become available. Then re-register the problematic provider which will then become the standby.

A.12 Timeouts

In vVol environments that consist of five or more storage containers involved in an SRM configuration, it is possible for SRM tasks to exceed the vCenter timeouts. These can be adjusted higher per the following VMware KB: <https://kb.vmware.com/s/article/1017253>.

B Appendix: Virtual Volume Operational Detail

This appendix contains additional detail related to the vVol environment.

B.1 Adding a TCP software adapter in vCenter

Before creating the masking view for the vPE, create the NVMe/TCP initiator group for the ESXi host. When running ESXi 8.0 U1 and higher there are two NVMe Qualified Names or NQNs per host. One NQN is for traditional NVMe/TCP with VMware, the other for NVMeoF vVols.

Initially, after adding the TCP software adapter in VMware to the ESXi host, there are no NQNs available in the Create Host dialog in Unisphere shown in [Figure 200](#).

The screenshot shows the 'Create Host' dialog in the Unisphere interface. At the top, there is a 'Host Name' field with a red asterisk indicating it is required. Below this is the 'Initiator Type' section, which contains three radio buttons: 'Fibre', 'iSCSI', and 'NVMe/TCP'. The 'NVMe/TCP' option is selected and highlighted with a red rectangular box. Underneath, the 'Select Initiators' section contains two side-by-side tables. The left table is titled 'Available Initiators' and the right table is titled 'Initiators in Host'. Both tables have columns for 'Name', 'Host ID', and 'IP Addresses'. Both tables are currently empty. Between the two tables are navigation arrows. At the bottom left of the dialog is a 'Set Host Flags' button. At the bottom right are 'Cancel' and 'Run Now' buttons.

Figure 200. Create Host

The reason it is empty is because VMware has not communicated with the array yet and transmitted the NQNs. Therefore, one needs to add the array or controller in vSphere.

B.1.1 Add array controller

This step is required whether using regular NVMe/TCP or NVMe/TCP vVols. Navigate to the vmhbaXX for NVMe/TCP, select the **Controllers** tab and select **ADD CONTROLLER** in [Figure 201](#).

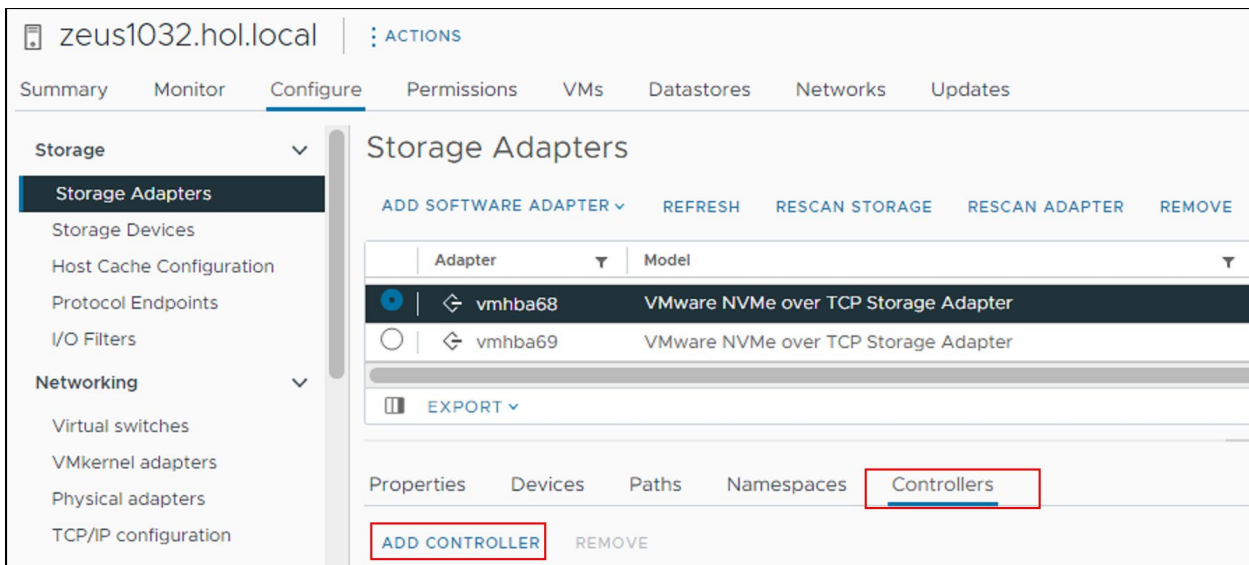


Figure 201. Add array controller – step 1

Enter in the IP of the interface assigned to the virtual port and the port number 8009 (TCP default, though 4420 will also work) in Figure 202.

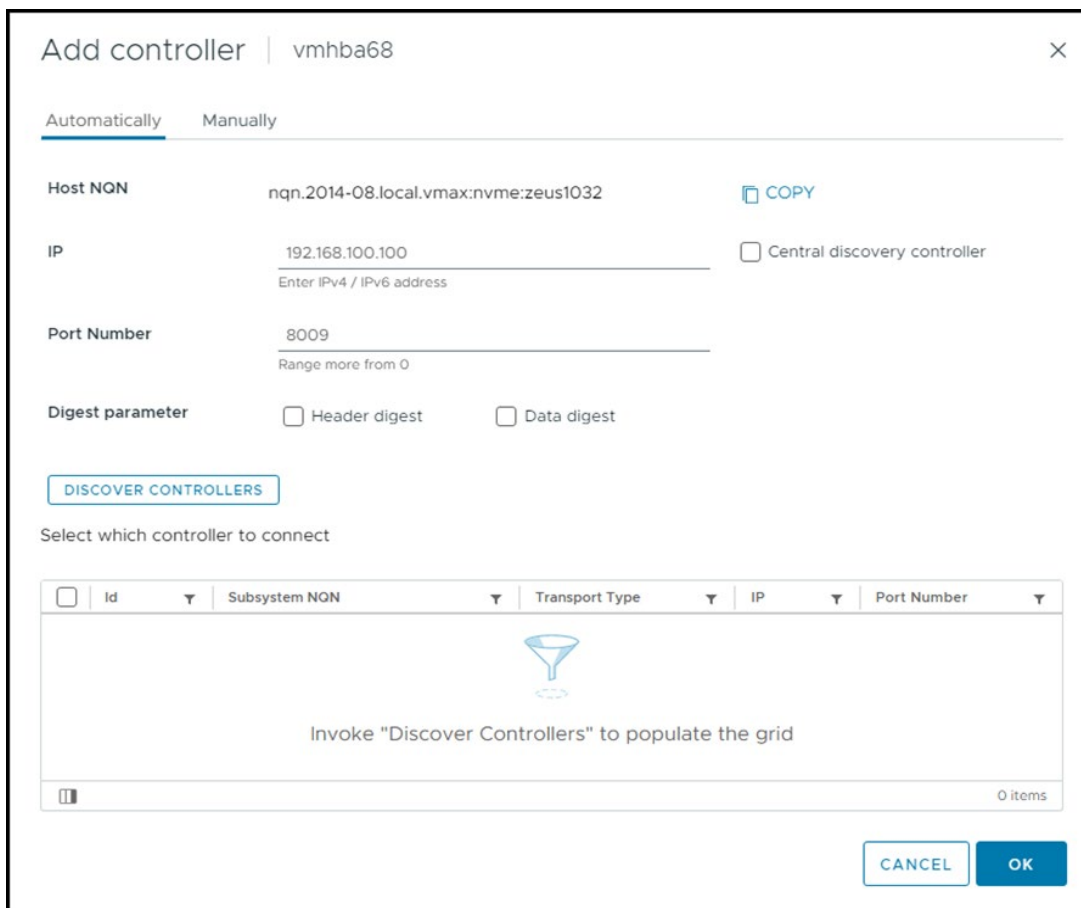


Figure 202. Add array controller – step 2

Select **DISCOVER CONTROLLERS** to discover the Subsystem NQN in [Figure 203](#). Note that depending on how many virtual ports are configured on the array, the user may see other ports (configured IPs), even if they are on different subnets/VLANs. Here, two different IPs are shown, though this adapter can only connect with one as each IP is on a different VLAN.

Add controller | vmhba68

Automatically | Manually

Host NQN: nqn.2014-08.local.vmax:nvme:zeus1032 COPY

IP: 192.168.100.100 Central discovery controller
Enter IPv4 / IPv6 address

Port Number: 8009
Range more from 0

Digest parameter: ☐ Header digest ☐ Data digest

DISCOVER CONTROLLERS

Select which controller to connect

<input type="checkbox"/>	Id	Subsystem NQN	Transport Type	IP	Port Num
<input type="checkbox"/>	65535	nqn.1988-11.com.dell:PowerMax_2500:00:000120000640	nvm	192.168.100.100	4420
<input type="checkbox"/>	65535	nqn.1988-11.com.dell:PowerMax_2500:00:000120000640	nvm	192.168.101.100	4420

2 items

CANCEL OK

Figure 203. Add array controller – step 3

As seen in [Figure 204](#) below, only check the box of the controller that matches the VMkernel associated with the vmnic/TCP software adapter and hit OK. This process can be repeated to achieve multi-pathing if there is a second TCP software adapter that is configured to communicate with the other IP.

Add controller | vmhba68

Automatically

Manually

Host NQN

nqn.2014-08.local.vmax:nvme:zeus1032

COPY

IP

192.168.100.100

Enter IPv4 / IPv6 address

☐ Central discovery controller

Port Number

8009

Range more from 0

Digest parameter

☐ Header digest

☐ Data digest

DISCOVER CONTROLLERS

Select which controller to connect

<input type="checkbox"/>	Id	Subsystem NQN	Transport Type	IP	Port Number
<input checked="" type="checkbox"/>	65535	nqn.1988-11.com.dell:PowerMax_2500:00:000120000640#vmhba68#192.168.100.100:4420	nvm	192.168.100.100	4420
<input type="checkbox"/>	65535	nqn.1988-11.com.dell:PowerMax_2500:00:000120000640#vmhba68#192.168.101.100:4420	nvm	192.168.101.100	4420

☒ 1

2 items

CANCEL

OK

Figure 204. Add array controller – step 4

The controller is now added in Figure 205:

zeus1032.hol.local

ACTIONS

Summary

Monitor

Configure

Permissions

VMs

Datastores

Networks

Updates

Storage

Storage Adapters

Storage Devices

Host Cache Configuration

Protocol Endpoints

I/O Filters

Networking

Virtual switches

VMkernel adapters

Physical adapters

TCP/IP configuration

Virtual Machines

VM Startup/Shutdown

Agent VM Settings

Default VM Compatibility

Swap File Location

System

Licensing

Host Profile

Time Configuration

Authentication Services

Certificate

Storage Adapters

ADD SOFTWARE ADAPTER

REFRESH

RESCAN STORAGE

RESCAN ADAPTER

REMOVE

Adapter	Model	Type	Status	Identifier	Targets	Devices
vmhba64	Emulex LPe35000/LPe36000 Fibre Channel Adapter	Fibre Channel	Online	20:00:00:10:9b:65:ee:25 10:...	2	2
vmhba65	Emulex LPe35000/LPe36000 Fibre Channel Adapter	Fibre Channel	Online	20:00:00:10:9b:65:ee:26 10:...	2	2
vmhba70	iSCSI Software Adapter	iSCSI	Online	iscsi_vmk(lqn.1998-01.com.v...	0	0
vmhba0	Lewisburg SATA AHCI Controller	Block SCSI	Unknown	--	0	0
vmhba1	Lewisburg SATA AHCI Controller	Block SCSI	Unknown	--	0	0
vmhba68	VMware NVMe over TCP Storage Adapter	NVME over T...	Online	--	0	0
vmhba69	VMware NVMe over TCP Storage Adapter	NVME over T...	Online	--	0	0

EXPORT

14 items

Properties

Devices

Paths

Namespaces

Controllers

ADD CONTROLLER

REMOVE



Name	Subsystem NQN	Transport Type	FUSE Support	Model
nqn.1988-11.com.dell:PowerMax_2500:00:000120000640#vmhba68#192.168.100.100:4420	nqn.1988-11.com.dell:PowerMax_2500:00:000120000640#vmhba68#192.168.100.100:4420	tcp	true	EMC PowerM...

Figure 205. Add array controller – step 5

Note that after creating and mounting the first vVol datastore with NVMe/TCP, you will see two additional controllers listed for the TCP adapter as in [Figure 206](#). This is expected:

Storage Adapters

ADD SOFTWARE ADAPTER ▾ REFRESH RESCAN STORAGE RESCAN ADAPTER REMOVE

Adapter	Model	Type	Status	Identifier	Targets	Devices
 vmhba68	VMware NVMe over TCP Storage Adapter	NVME over T...	Online	--	0	0
 vmhba69	VMware NVMe over TCP Storage Adapter	NVME over T...	Online	--	0	0

EXPORT ▾ 14 items

Properties Devices Paths Namespaces Controllers

ADD CONTROLLER REMOVE

<input type="checkbox"/>	Name	Subsystem NQN	Transport Type	FUSE Support	Model	Firmware Version
<input type="checkbox"/>	nqn.2014-08.org.nvmexpress.disc...	nqn.2014-08.org.nvmexpress.disc...	tcp	true	EMC PowerMax_2500	60790225TCP
<input type="checkbox"/>	nqn.1988-11.com.dell:PowerMax_25...	nqn.1988-11.com.dell:PowerMax_25...	tcp	true	EMC PowerMax_2500	60790225TCP
<input type="checkbox"/>	nqn.1988-11.com.dell:PowerMax_25...	nqn.1988-11.com.dell:PowerMax_25...	tcp	true	EMC PowerMax_2500	60790225TCP

3 items

Figure 206. Additional controllers

B.1.2 Add host

With successful communication, run a discover in Unisphere and start the **Create Host** dialog. Now the user should see two available initiators for NVMe/TCP, vVol and non-vVol. Expand the IP Addresses column as in [Figure 207](#). The initiator for vVols will have the IP of the VMkernel interface, in this case 192.168.100.10. Note that there is only one host NQN so even if multiple controllers are configured, there will only be a single initiator.

Note: If the initiators do not appear, see [Obtaining the Host NQN and Host ID](#) where one can obtain the initiator names through CLI.

Create Host

Host Name *
This field is required

Initiator Type

☐ Fibre
 ☐ iSCSI
 ☒ NVMe/TCP

Select Initiators

Available Initiators			2 Items
<input type="checkbox"/> Name ↑	Host ID	IP Addresses	
<input type="checkbox"/> nqn.2014-08.local.vn	615F55EC69...	(1)	0.0.0.0
	—		
<input type="checkbox"/> nqn.2021-01.com.vmr	52C6C14734...	(1)	192.168.1
	—		

Initiators in Host			0 Items
<input type="checkbox"/> Name ↑	Host ID	IP Addresses	

Figure 207. Create Host – step 1

If one hovers over the initiator, as below in [Figure 208](#), one can see the entire name. The vVol initiator has **vvol** in the name. Select the check box and move it over the right-hand side. Be sure to provide a name for the initiator group. Leave the regular NVMe/TCP initiator on the left-hand side. It is only used for regular devices, not vVols. Select **Run Now**.

Create Host

Host Name *
zeus1032_vvol_tcp_ig

Initiator Type
☐ Fibre ☐ iSCSI ☒ NVMe/TCP

Select Initiators

Available Initiators			1 Items
<input type="checkbox"/> Name ↑	Host ...	IP Ad...	
<input type="checkbox"/> nqn.2014-08.local.vmax:nvme:zeus	615F...	> (1)	
<input type="checkbox"/> nqn.2021-01.com.vmware:615f55ec-693c-0cbc-ad9c-b026 nqn.2021-01.com.vmware:615f55ec-693c-0cbc-ad9c- b02628da4920 vvol:30e8b0ee3999			

Initiators in Host			1 Items
<input type="checkbox"/> Name ↑	Host ID	IP Addresses	
<input type="checkbox"/> nqn.2021-01.com.	52C6C...	> (1)	
		192.168.100.10	

Set Host Flags

Cancel Run Now

Figure 208. Create Host – step 2

The task is successful in Figure 209.

Task in progress

✓ Success

Starting Tasks...

Succeeded

Create new Host zeus1032_vvol_tcp_ig...

Succeeded

OK

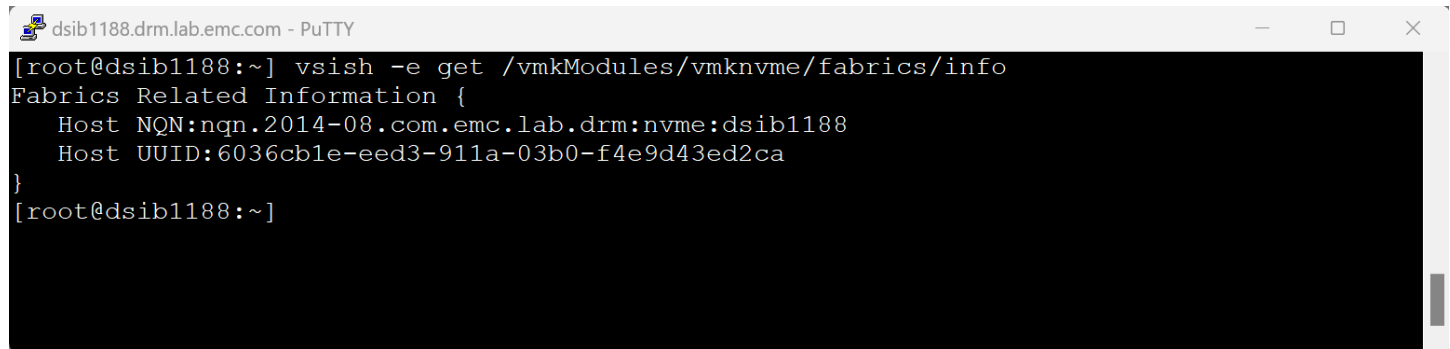
Figure 209. Create Host – step 3

B.2 Obtaining the Host NQN and Host ID

In the event the initiator for either vVol NVMe/TCP or regular NVMe/TCP does not appear in Unisphere, each one can be obtained through CLI on the ESXi host. When adding an initiator manually in Unisphere, one must supply both the host NQN and the Host ID. The Host ID is different for regular NVMe/TCP and vVol NVMe/TCP.

The names of the regular NVMe/TCP host NQN and Host ID are obtained with `vsish` in [Figure 210](#):

```
vsish -e get /vmkModules/vmknvme/fabrics/info
```

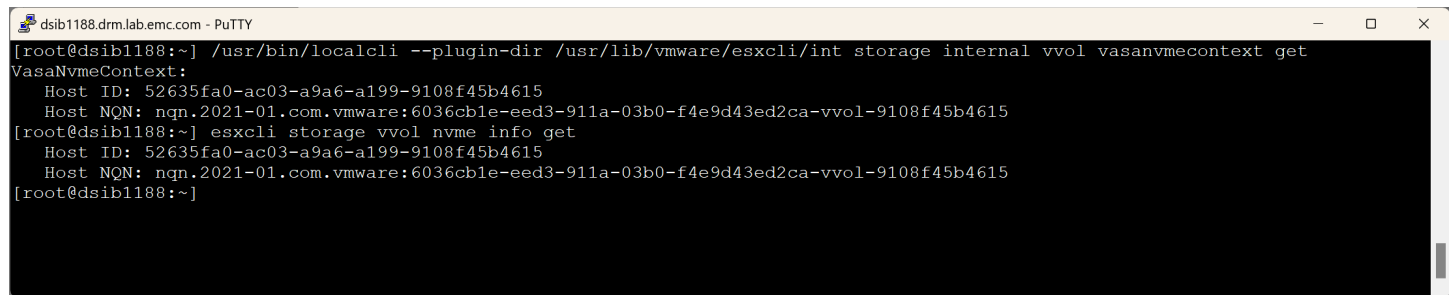


```
dsib1188.drm.lab.emc.com - PuTTY
[root@dsib1188:~] vsish -e get /vmkModules/vmknvme/fabrics/info
Fabrics Related Information {
  Host NQN:nqn.2014-08.com.emc.lab.drm:nvme:dsib1188
  Host UUID:6036cb1e-eed3-911a-03b0-f4e9d43ed2ca
}
[root@dsib1188:~]
```

Figure 210. Host NQN and Host ID for NVMe/TCP

The names of the vVol NVMe/TCP host NQN and Host ID can be obtained with `localcli` prior to vSphere 8.0.2, or `esxcli` with vSphere 8.0.2. Both are show in [Figure 211](#):

```
/usr/bin/localcli --plugin-dir /usr/lib/vmware/esxcli/int storage internal
vvol vasanvmecontext get
esxcli storage vvol nvme info get
```



```
dsib1188.drm.lab.emc.com - PuTTY
[root@dsib1188:~] /usr/bin/localcli --plugin-dir /usr/lib/vmware/esxcli/int storage internal vvol vasanvmecontext get
VasaNvmeContext:
  Host ID: 52635fa0-ac03-a9a6-a199-9108f45b4615
  Host NQN: nqn.2021-01.com.vmware:6036cb1e-eed3-911a-03b0-f4e9d43ed2ca-vvol-9108f45b4615
[root@dsib1188:~] esxcli storage vvol nvme info get
  Host ID: 52635fa0-ac03-a9a6-a199-9108f45b4615
  Host NQN: nqn.2021-01.com.vmware:6036cb1e-eed3-911a-03b0-f4e9d43ed2ca-vvol-9108f45b4615
[root@dsib1188:~]
```

Figure 211. Host NQN and Host ID for vVol NVMe/TCP

Enter the results as below in Unisphere in [Figure 212](#). Note that the user must remove the dashes from the Host ID before entering into the field (copy and paste will fail to do anything if the dashes are kept).

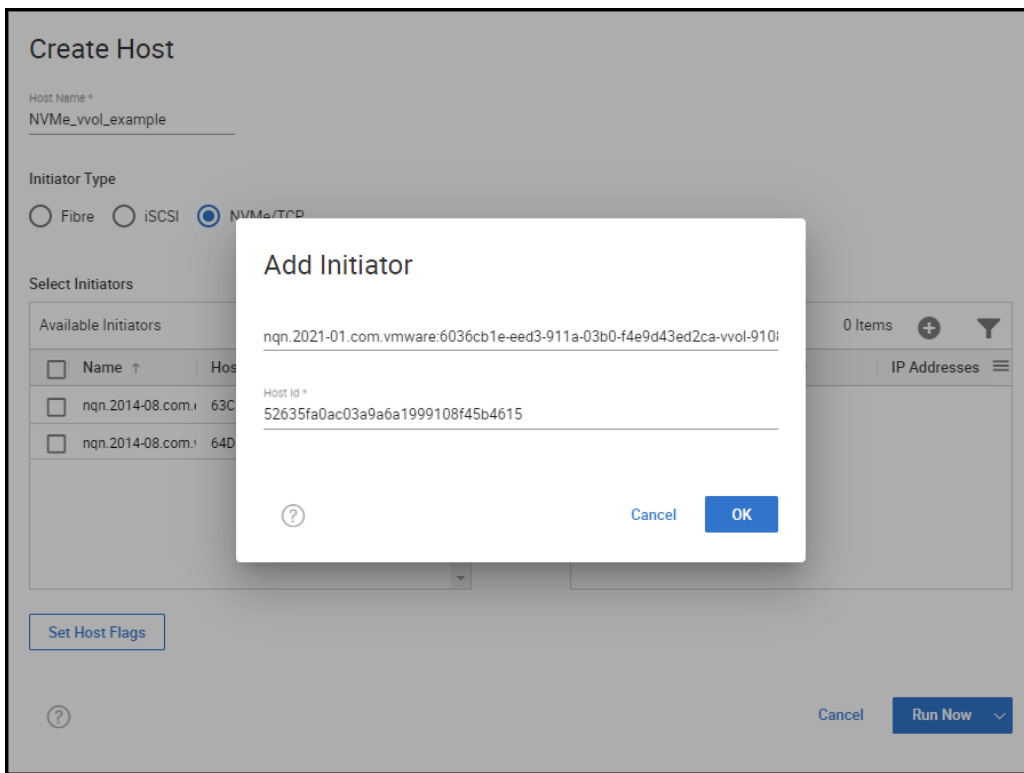


Figure 212. Manual host creation

B.3 vVol Config UNMAP

Beginning with vSphere 8.0.2, VMware supports unmapping on a config vVol. Config vVols are formatted with VMFS 6 so in essence they are datastores. The syntax is:

```
esxcli storage vvol virtualvolume unmap -c <container-id> -u <meta vVol UUID>
```

To issue manual unmap navigate to the folder of the VM whose config vVol requires unmapping:

```
[root@dsib1188:~] cd /vmfs/volumes/NVMe_vVol_Datastore/NVMe_vVol_VM
[root@dsib1188:/vmfs/volumes/vvol:60000970072a1c95-
4dee014c0001e78a/naa.60000970072A1C954DEE004C00000063]
```

From here, gather the container id which is the first folder, vvol:60000970072a1c95-4dee014c0001e78a, and the meta vVol UUID which is the second folder (the config vVol of the VM), naa.60000970072A1C954DEE004C00000063.

The UUID can be compared against a `symdev` listing of vVols in [Figure 213](#):

```

root@dsib2235:~# symdev list -vvol -sid 341
Symmetrix ID: 000120200341

```

Sym	Physical	SA :P	Config	Attribute	Sts	Cap (GB)
10007	Not Visible	???:???	VVOL	N/Grp'd	WD	255.0
10008	Not Visible	???:???	VVOL	N/Grp'd	WD	16.0
10009	Not Visible	???:???	VVOL	N/Grp'd	WD	255.0
1000A	Not Visible	***:***	VVOL	N/Grp'd	RW	255.0
1000B	Not Visible	***:***	VVOL	N/Grp'd	RW	90.0
1000C	Not Visible	***:***	VVOL	N/Grp'd	RW	4.0

```

root@dsib2235 ~# symdev list -vvol -sid 341 -wnn
Symmetrix ID: 000120200341

```

Sym	Physical	Config	Attr WWN
10007	Not Visible	VVOL	60000970072713C15437002900000001
10008	Not Visible	VVOL	60000970072713C15437002900000013
10009	Not Visible	VVOL	60000970072713C15437002900000011
1000A	Not Visible	VVOL	60000970072A1C954DEE004C000000063
1000B	Not Visible	VVOL	60000970072A1C954DEE004C000000064
1000C	Not Visible	VVOL	60000970072A1C954DEE004C000000065

```

root@dsib2235 ~#

```

Figure 213. Obtaining meta vVol UUID for unmap

Note how NVMe/TCP vVols are represented by the “*” symbol for director and port.

In this unmap example below, a large ISO was added to the folder, then deleted, so there was something to unmap:

```
esxcli storage vvol virtualvolume unmap -c vvol:60000970072a1c95-4dee014c0001e78a -u naa.60000970072A1C954DEE004C000000063
```

As there is no response from the command, view the hostd.log file for confirmation in [Figure 214](#):

```

dsib1188.drm.lab.emc.com - PuTTY
root@dsib1188:~# esxcli storage vvol virtualvolume unmap -c vvol:60000970072a1c95-4dee014c0001e78a -u naa.60000970072A1C954DEE004C000000063
root@dsib1188:~# grep "Total Unmapped blocks from vmfs" /var/run/log/hostd.log
2023-10-13T17:55:35.936Z In(166) Hostd[2099845]: [Originator@6876 sub=Libs opID=esxcli-69-0495 sid=52e92ab2 user=root] Unmap: Done : Total Unmapped blocks from vmfs6 volume naa.60000970072A1C954DEE004C000000063 : 195584 (LFB Pa
ss)
2023-10-13T17:57:26.851Z In(166) Hostd[2099845]: [Originator@6876 sub=Libs opID=esxcli-69-0495 sid=52e92ab2 user=root] Unmap: Done : Total Unmapped blocks from vmfs6 volume naa.60000970072A1C954DEE004C000000063 : 158217 (SFB Pa
ss)
root@dsib1188:~#

```

Figure 214. Running unmap on a config vVol

In general, unless the config vVol is being used for a content library or to store ephemeral files, it is unnecessary to run manual unmap.

C Appendix: Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

C.1 Dell

- The Dell PowerMax and VMware vSphere Configuration Guide

<https://infohub.delltechnologies.com/t/dell-powermax-and-vmware-vsphere-configuration-guide-1/>

- Unisphere for PowerMax

<https://www.dell.com/support/product-details/en-us/product/unisphere-powermax/overview>

C.2 Broadcom/VMware

- VMware vSphere

<https://techdocs.broadcom.com/>