





What's New with vSphere 8 Core Storage

VMware Storage

Table of contents

What's New with vSphere 8 Core Storage	5
vSphere 8 Update 3	5
Key Enhancements	5
vVols	5
vVols Stretched Storage Cluster SCSI Deployments (Uniform config)	5
Requirements	7
Additional UNMAP Support	7
Guest Clustering Application Support on NVMe-oF vVols	7
Update Host Authentication Reporting for VASA Provider	8
vVols Storage Provider Host Granularity	8
NVMe-oF	9
Provide NVMe reservation support for clustered VMDK with NVMe/TCP -VMFS	9
Enable NVMe-oF Support for NVMe Cross Namespace Copy.	9
VMFS	10
Reduce Time to Inflate EZT Disk on VMFS	10
Limiting the Number of vSphere Hosts Sending UNMAP to a Given VMFS Datastore	10
PSA	11
PSA Support for Fabric Notifications (FPIN, Link errors, Congestion)	11
NFS	11
NFS v4.1 vmk Port Binding	11
Add nConnect support to NFS v4.1	12
CNS/CSI	13
CNS Support of vSAN ESA File Service Supports 250 File Shares.	13
Enable File Volume in HCI Mesh Topology within a Single Center.	13
Use CNS on TKGs on Stretched vSAN	13
Enable PV Migration Across Non-shared Datastores within the Same VC	13
Use CNS on vSAN Max	13
vSphere 8 Update 2	14
vVols	14
 type=image,objattr=a:23:{s:10:"margin_top";d:0;s:13:"margin_bottom";d:0;s:11:"margin_left";d:0;s:12:"margin_right";d:0;s:11:"padding_top";i:0;s:14:"padding_bottom";i:0;s:12:"padding_left";i:0;s:13:"padding_right";i:0;s:5:"width";d:41.275;s:6:"height";d:41.010416666666664;s:10:"border_top";a:1:{s:1:"w";i:0;}s:13:"border_bottom";a:1:{s:1:"w";i:0;}s:11:"border_left";a:1:{s:1:"w";i:0;}s:12:"border_right";a:1:{s:1:"w";i:0;}s:10:"visibility";s:7:"visible";s:14:"vertical-align";s:2:"BS";s:4:"file";s:78:"https://images.core.vmware.com/sites/default/files/inline-images/image_409.png";s:4:"type";s:5:"image";s:5:"itype";s:3:"png";s:6:"orig_h";i:198;s:6:"orig_w";i:199;s:12:"image_height";d:41.010416666666664;s:11:"image_width";d:41.275;} 	14

- vVols Extend Online Shared Disks 14
- Support for vVols Online Disk Extension with Oracle RAC 15
- vVols NVMe in-band Migration 15
- Automated recovery from PDL for vVol PEs 15
- Enable 3rd party MPP support for NVMe vVols 15
- UNMAP Support for config vVol 15
- NVMeoF 16
- Clustered Applications 16
- Support vNVME controller type for MS WSFC 16
- Support Oracle RAC clustering with vVols NVMe (FC, TCP) 16
- Enable vNVMe support with Oracle RAC (vVols) 16
- VMFS 16
- Improved SE Sparse snapshot offline consolidation performance. 16
- NFS 16
- DNLC cache for NFS4.1 16
- nConnect 16
- vSphere 8 Update 1 18
 - vSphere Virtual Volumes (vVols) 18
 - Multi VC deployment for VASA Provider without Self-Signed Certificate 18
 - New vVols VASA 5.0 Spec 18
 - VASA 5.0 Feature Details 19
 - Move Sidecar vVols in config-vvol instead of another vVol object. 19
 - Config-vvol Enhancements, Support for VMFS6 config-vvol with SCSI vVols (instead of VMFS5) 19
 - Add NVMe-TCP support for vVols 20
 - NVMeoF, PSA, HPP 20
 - Infrastructure to support End to End NVMe 20
 - Increase Max Paths per NVMe-oF Namespaces from 8 to 32. 21
 - Increase WSFC clusters per ESXi host from 3 to 16. 21
 - VMFS 22
 - Enhanced XCOPY to Datastores Across Different Storage Arrays. 22
 - NFS 22
 - NFSv3 vmkPortBinding 22
- vSphere 8 23
 - New vSphere 8.0 Core Storage enhancements, features, and additions. 23
 - vSphere Virtual Volumes (vVols) 23
 - VM Swap improvements 23
 - Config vVol to remain bound. 23

NVMeoF vVols	23
Technical Details:	23
NVMeoF Enhancements	24
Support 256 namespaces and 2K paths with NVMe-TCP and NVMe-FC	25
Extend reservation Support for NVMe device	25
Auto-discovery of NVMe Discovery Service support in ESXi	25
Unmap Space Reclamation Enhancement	25
Lower minimum reclamation rate to 10MBPS	25
Dedicated Unmap Scheduling Queue	26
Container Storage CNS/CSI	26
VMFS and vSANDirect Disk Provisioning Storage Policy.	26
NFS Enhancements	27

What's New with vSphere 8 Core Storage

vSphere 8 Update 3

Each update in vSphere 8 has added numerous features to help customers with increased scale, resilience, and performance. In vSphere 8 Update 3, we have continued that momentum with some new enhancements and very exciting new vVols features! Continuing primary engineering focus on vVols and NVMe-oF, we also ensure both are supported in VMware Cloud Foundation. Some enhancements include additional support for MS WSFC guest clustering support on vVols and NVMe-oF, space reclamation enhancements, and enhancements for both VMFS and NFS.

Key Enhancements

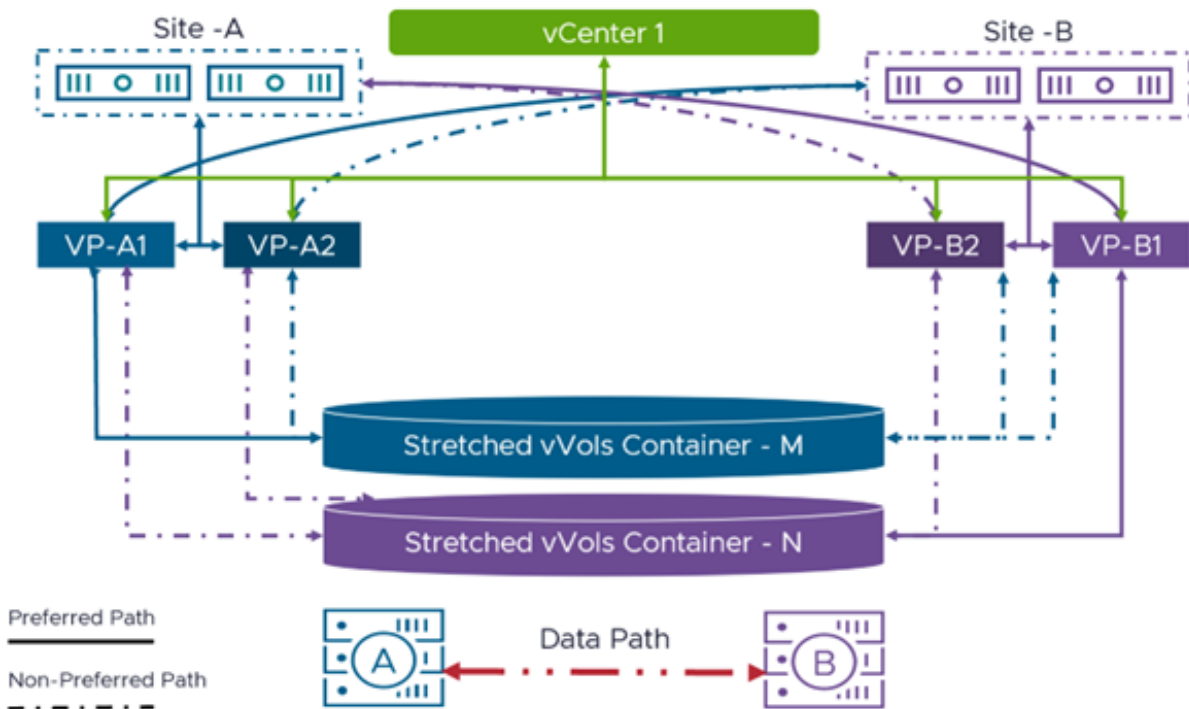
- **vVols Stretched Storage Cluster Support**
- New vVols VASA 6 Spec
- Clustered VMDK for NVMe/TCP
- Limit the maximum number of hosts submitting Unmap
- NFS 4.1 Port Binding and nConnect support
- Additional CNS/CSI vSAN support

vVols

vVols Stretched Storage Cluster SCSI Deployments (Uniform config)

vVols Stretched Storage Cluster (vVols-SSC) has been one of the number one requests for vVols for many years, especially in the European regions. We are excited to announce the wait is over! In vSphere 8 U3, we are adding support for stretched vVols storage on SCSI only (FC or iSCSI) for this release. This was a heavy lift and took quite a bit of engineering from both VMware and our storage partners. Initially, [Pure Storage](#), who was the design partner for this feature, will be supporting vVols-SSC, but many of our storage partners are actively working on adding support as well.

In vSphere 8 U2, we added a new VASA spec, [VASA 5](#), which laid much of the groundwork for vVols-SSC. This included support for multiple VCs to a single VASA, certificate management, and workload separation. In vSphere 8 U3, a new VASA 6 spec was added to support vVols-SSC.



Why did it take so long?

One of the reasons vVols-SSC took so long was the additional enhancements needed to VM Component Protection (VMCP). When using stretched storage, there needs to be a process when HA is enabled to handle storage events such as Permanent Device Loss (PDL) or All Paths Down (APD). VMCP is a vSphere HA feature that detects VM storage failures and provides automated recovery for affected virtual machines. It protects VMs against storage failures that don't necessarily cause VMs to power off but render them unusable (e.g., losing network/disk communication, causing corruption). It monitors storage for incidents, then kills affected VMs to kick in HA failover workflow. We also had to enable HA for the VASA. This meant making VASA able to handle either array/site failing but continue functioning for the remaining array/site.

The accessibility of a stretched storage container diverges from legacy containers in when they become inaccessible. This is changed in order to get semantics closely resembling VMFS and hence have VMCP working for vVols (it does not work for legacy containers). Another function needed for VMCP to work is that ESXi knows which PE a container uses - in the initial release, there can only be a single PE for a storage container, but later on, we will add support for more (and dynamically change this).

Summary of the failure/recovery scenarios

- **A vVols container/datastore becomes inaccessible.** A vVols container/datastore becomes inaccessible if either all the data paths (to the PE) or all the control path (access to the VPs exposing this container) becomes inaccessible or if all VASA Provider paths reporting the stretched container report it UNAVAILABLE (a new state a VP can report for stretched containers).
- **A PE of the vVols container may become inaccessible.** If the PE for the container goes into a PDL state, then the container also goes into PDL state. At that point, VMCP will be responsible for stopping VMs on affected hosts and restarting them on alternate hosts where the container is available.
- **A vVols container or a PE becomes accessible again or the VP connectivity is restored.** The container will come back to an accessible state from APD once one VP and the PE become accessible again. The container will only come back to accessible state from PDL state once all VMs using the containers have been powered off and any clients having open file handles to the vVols datastore closes them.

Behavior of a stretched container/datastore is fairly similar to VMFS and exiting from PDL require the PE to be destroyed and this can only happen when all vVols bound to the PE are released. Just like VMFS (or a PSA device) cannot exit out of PDL unless all clients of the VMFS volume (or the PSA device) close their handles.

Requirements

- SCSI (FC or iSCSI)
- Max vSphere host site to site RTT 11ms
- Max Storage array RTT 11ms
- vSphere vMotion network has a 250 Mbps dedicated bandwidth
- Single Virtual Center (Currently, vCenter HA is not supported with vVols)
- Storage IO Control is not supported on a vVol-SSC enabled datastore

Additional UNMAP Support

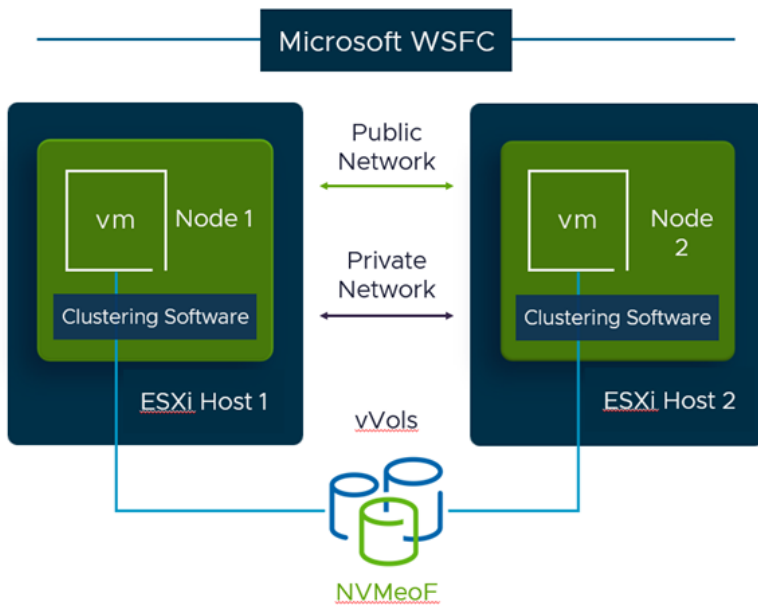
For vVols, starting in vSphere 8.0 U1, the config-vvol is now created with a 255GB VMFS vVol instead of a 4GB. This enabled new features but added to the need for space reclamation within the config-vvol. In 8.0 U3 we added support for both manual (CLI) and automatic UNMAP support for the config-vvol for both SCSI and NVMe. This ensures that as data is written and deleted within the config-vvol, we make sure to keep the space optimized for new writes. Starting with vSphere 8.0 U3 we also support Unmap for NVMe-oF Datastores, support for UNAMP with SCSI volumes was added in a prior release.

Reclaim Space on the vSphere Virtual Volumes Datastores



Guest Clustering Application Support on NVMe-oF vVols

In vSphere 6.7 we added support for SCSI3-PR and MS WSFC, in vSphere 8.0 U2 we added support for hot extending shared disks with MS WSFC. These features we limited to SCSI on vVols, for MS WSFC, but Oracle RAC multi-writer supported both SCSI and NVMe-oF. In vSphere 8.0 U3, we have extended the MS WSFC shared disks support to NVMe/TCP and NVMe/FC backed vVols. We also added support for VM's virtual NVMe (vNVME) controller as the frontend for guest clustering solution like MS WSFC. Please note for vNVMe controller as frontend for MS WSFC is currently only supported with NVMe as backend.

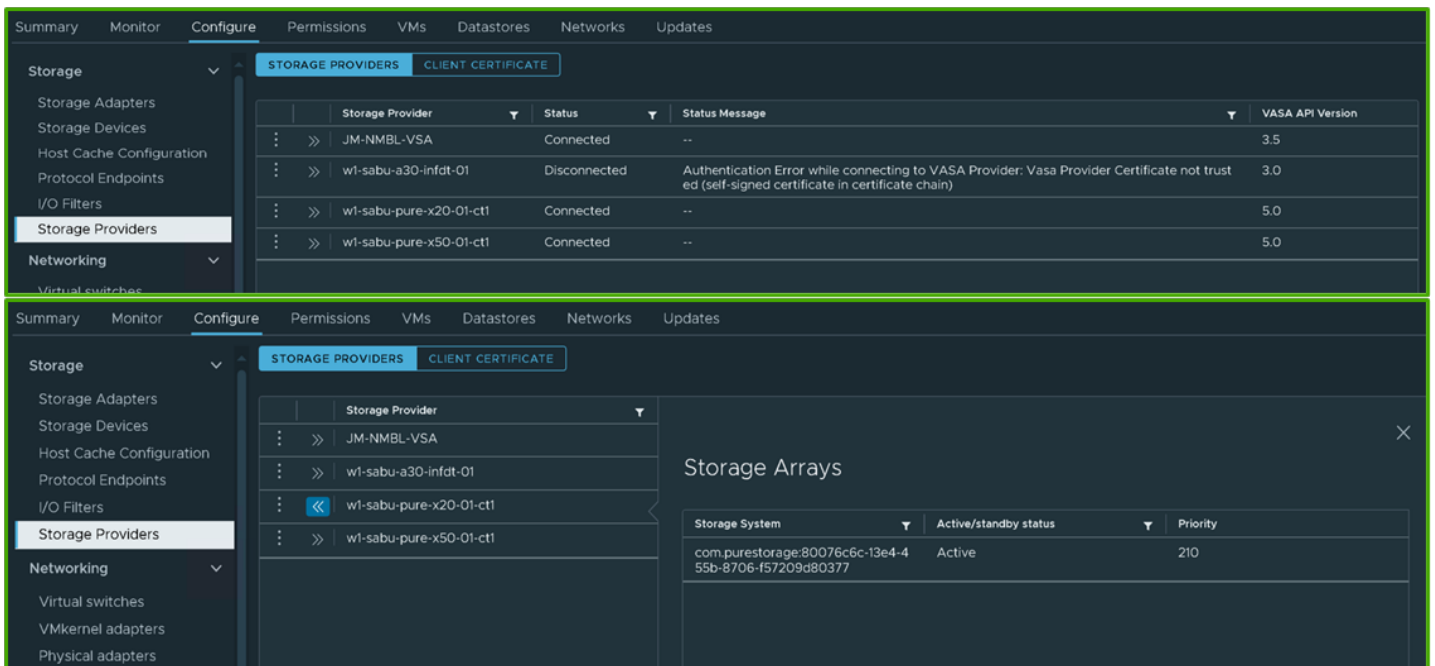


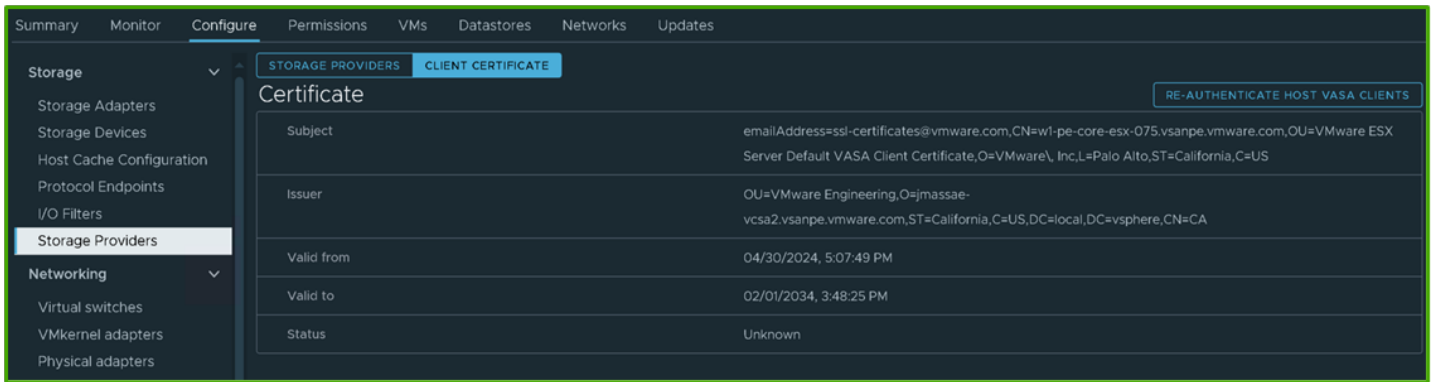
Update Host Authentication Reporting for VASA Provider

Occasionally, when configuring the Storage Provider for vVols, some hosts may fail to authenticate and can be challenging to debug. In 8.0 U3, we have added the capability for vCenter to notify users about specific hosts authentication against the VASA provider failed and we have provided a mechanism to re-authenticate the hosts in the vCenter UI. This simplifies detecting and resolving Storage Provider authentication issues.

vVols Storage Provider Host Granularity

With vSphere 8 U3, there is now additional host-level vVols Storage Provider and certificate information. This provides additional vVols details to our customers and support when troubleshooting.

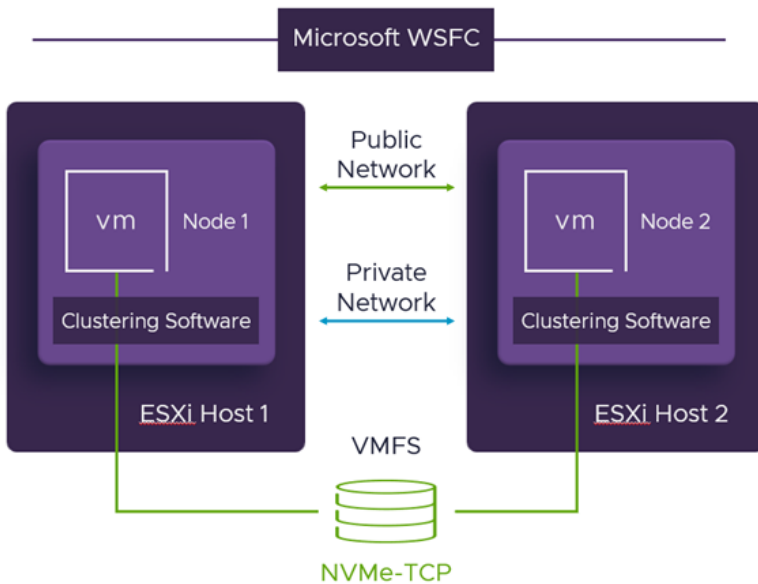




NVMe-oF

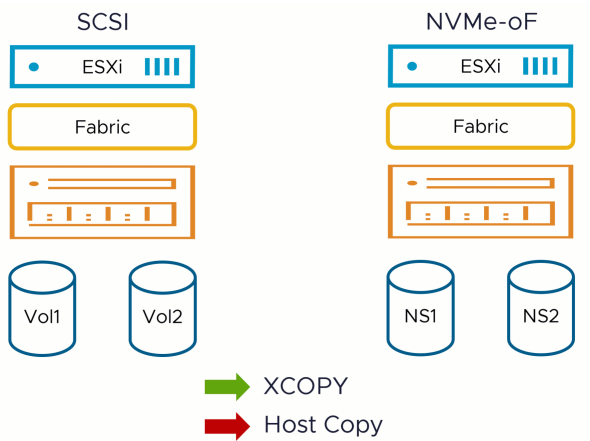
Provide NVMe reservation support for clustered VMDK with NVMe/TCP -VMFS

In vSphere 8.0 U3, we have extended Guest Clustering support to NVMe/TCP, initially only NVMe/FC was supported. This gives customers more options when using NVMe-oF and wanting to move guest clustering applications such as MS WSFC and Oracle RAC to NVMe-oF datastores.



Enable NVMe-oF Support for NVMe Cross Namespace Copy.

In prior releases, the VAAI function, Hardware Accelerated Copy, or XCOPY, was supported with SCSI but wasn't supported with NVMe-oF. This meant that copies across NVMe namespaces used the host resources for the data transfer. With the release of vSphere 8.0 U3, Cross Namespace Copy for NVMe-oF, is now available for supported arrays. The use case here is the same as SCSI XCOPY across logical units/namespaces. Data transfers like disk copy/clone or VM clones can now be significantly faster. This capability offloads the data transfer to the array, subsequently reducing the host utilization.




VMFS

Reduce Time to Inflate EZT Disk on VMFS

In vSphere 8.0 U3, a new VMFS API was implemented to accelerate the inflation of blocks on a VMFS disk while the disk is in-use. This API can be up to 10x faster than existing methods when used to inflate an EZT disk on VMFS.


Virtual disks on VMFS have a provisioning type that defines how the underlying storage is reserved. Thin, Lazy Zeroed Thick, or Eager Zeroed Thick. EZT is typically chosen on VMFS for faster run-time performance because all blocks are fully allocated and zeroed up-front during the creation of the disk. If a user previously chose to provision a thin disk and wanted to inflate the disk to EZT, previously, this process was slow. The new API enables inflation to be substantially faster.



Limiting the Number of vSphere Hosts Sending UNMAP to a Given VMFS Datastore

In vSphere 8.0 U2, we added the capability to limit the UNMAP rate to 10MB/s from 25MB/s. This is intended for customers with high churn or power-off storms, to help mitigate the impact of space reclamation on the array.

The default is all hosts in a cluster which is up to 128 hosts. In 8.0 U3, there's a new advanced reclaim parameter called Reclaim Max Hosts. This can be set with a value between 1 and 128 and is a per-datastore setting. To change this setting, you use ESXCLI. The way the algorithm works is once a new value is set, the number of hosts submitting UNMAP at any one time is limited to that number. For example, if you set the max number to 10, in a cluster of 50, only 10 hosts will send UNMAP to the datastore at a time. If other hosts need to reclaim space, once one of the 10 is finished, a slot will be open for another host to reclaim space.



Space Reclamation on vSphere VMFS Datastores

Usage: `esxcli storage vmfs reclaim config set -l <Datastore> -n <number_of_hosts>`

Here is an example of changing the max number of hosts to submit UNMAP at one time.


```
[root@wl-pe-core-esx-075:~] esxcli storage vmfs reclaim config get -l PS-VMFS-x50
Reclaim Granularity: 1048576 Bytes
Reclaim Priority: low
Reclaim Method: fixed
Reclaim Bandwidth: 100 MB/s
Reclaim Max Hosts: 128
[root@wl-pe-core-esx-075:~] esxcli storage vmfs reclaim config set -l PS-VMFS-x50 -n 10
[root@wl-pe-core-esx-075:~] esxcli storage vmfs reclaim config get -l PS-VMFS-x50
Reclaim Granularity: 1048576 Bytes
Reclaim Priority: low
Reclaim Method: fixed
Reclaim Bandwidth: 100 MB/s
Reclaim Max Hosts: 10
```

PSA

PSA Support for Fabric Notifications (FPIN, Link errors, Congestion)

We have added support for Fabric Performance Impact Notification (FPIN) in vSphere 8.0 U3. With FPIN the vSphere infrastructure layer can now handle notifications from SAN switches or targets to learn about degraded SAN links to make sure to use the healthy paths to storage devices. It can notify the hosts of link congestion, and errors. FPIN is an industry-standard that provides a means to notify devices of link and other issues with a connection or possibly a path through the fabric.

You can use the command `esxcli storage FPIN info set -e= <true/false>` to activate or deactivate the Fabric Performance Impact Notification (FPIN) frame.

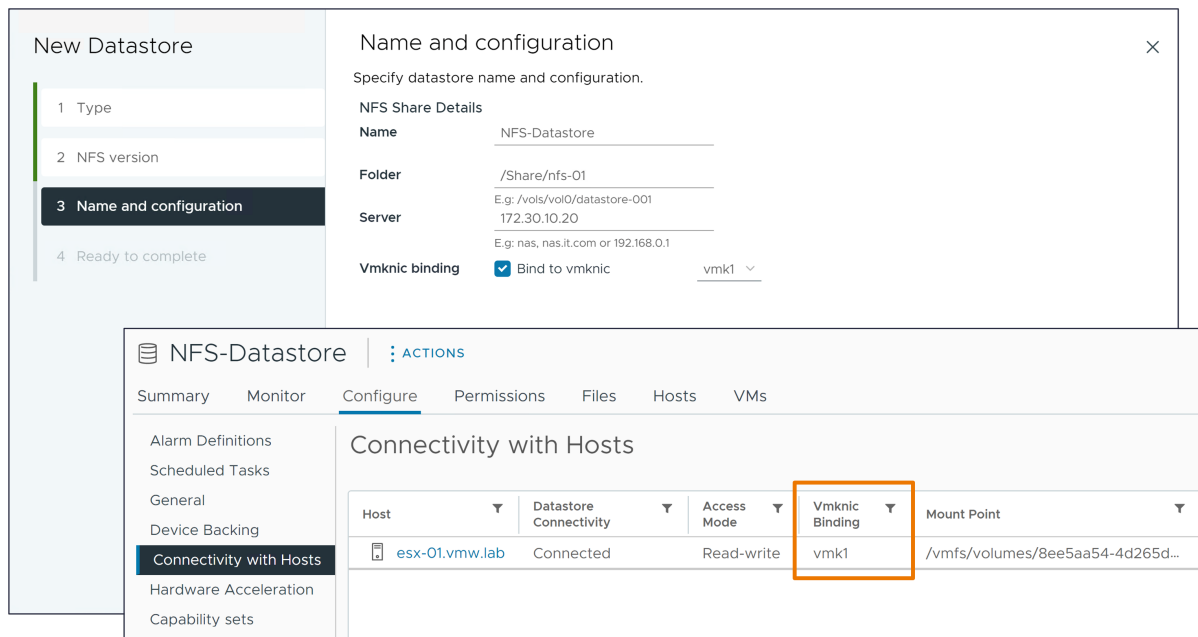


NFS

NFS v4.1 vmk Port Binding

This feature adds the ability to bind an NFS v4.1 connection to a specific vmknic to provide path isolation. When using multipathing, multiple vmknics can be provided. This provides path isolation and helps with security by directing NFS traffic across a specified subnet/VLAN and ensures NFS traffic does not use the mgmt or other vmks. Support for NFS v3 added in vSphere 8.0 U1. Currently this feature is supported using esxcli interface only and may be used in conjunction with nConnect.

Configure VMkernel Binding for NFS 4.1 Datastore on ESXi Host



Add nConnect support to NFS v4.1

Starting with 8.0 U3 nConnect support has been added for NFS v4.1 datastores. nConnect provides multiple connections using a single IP in a session thus extending session trunking functionality to that IP. With this feature, multipathing and nConnect coexist. Customers can configure datastores with multiple IPs to the same server and also multiple connections with the same IP. Currently, the maximum number of connections is limited to 8, with the default value being 1. Current versions of vSphere NFSv4.1 implementations create a single TCP/IP connection from each host to each datastore. Being able to add multiple connections per IP can greatly increase performance.

When adding a new NFSv41 datastore, the number of connections can be specified at the time of the mount using the command:

```
esxcli storage nfs41 add -H <host> -v <volume-label> -s <remote_share> -c <number_of_connections>
```

The maximum number of connections per session are limited to "4" by default, however it can be increased to "8" using advanced NFS option.

- `esxcfg-advcfg -s 8 /NFS41/MaxNConnectConns`
- `esxcfg-advcfg -g /NFS41/MaxNConnectConns`


The total number of connections used across all mounted NFSv41 datastores are limited to 256.

For an existing NFSv41 datastore, the number of connections can be increased or decreased at run time using the following command:

```
esxcli storage nfs41 param set -v <volume-label> -c <number_of_connections>
```

There is no impact to multipathing. NFS41 nConnect and multipaths can coexist. The number of connections are created for each of the multipathing IPs.

```
esxcli storage nfs41 add -H <IP1,IP2> -v <volume-label> -s <remote_share> -c <number_of_connections>
```



CNS/CSI

CNS Support of vSAN ESA File Service Supports 250 File Shares.

Currently, CNS only supports 100 file share volumes. In vSphere 8.0 U3, we have increased the limit to 250 volumes. This will help with scale for customers needing additional file share volumes for K8's PVs or PVCs.

Enable File Volume in HCI Mesh Topology within a Single Center.

Enable file volume in HCI Mesh with topology within a single VC.

Use CNS on TKGs on Stretched vSAN

Support Stretched vSAN Cluster for TKGs to ensure High Availability.

Enable PV Migration Across Non-shared Datastores within the Same VC

Ability to move a PV either attached or detached from vSAN to vSAN where there is no common host. An example for this would be the ability to move K8s workload from a vSAN OSA cluster to a vSAN ESA cluster.

Use CNS on vSAN Max

Enable support for vSphere Container Storage Plug-in consumers to deploy CSI Volumes to vSAN Max deployments.

vSphere 8 Update 2

vSphere 8 Update 2 has some significant announcements, and the storage side is no exception. As you have hopefully noticed, VMware is focusing on vSAN, vVols, and NVMeoF and with that, it is definitely the year for vVols. There have been so many great updates and enhancements in vSphere 8. For example, new vVols VASA specs, better performance and resilience, enhanced certificate mgmt., and support for NVMeoF to name a few. We have also made sure vVols is supported throughout the VMware ecosystem.

Although there aren't a lot of new core storage features in vSphere 8 Update 2, there are some incredibly significant updates nonetheless. vVols, NVMeoF, VMFS, and NFS all received enhancements and features many customers will appreciate.

You can see the latest in [What's New with vSphere 8 U2 here](#).

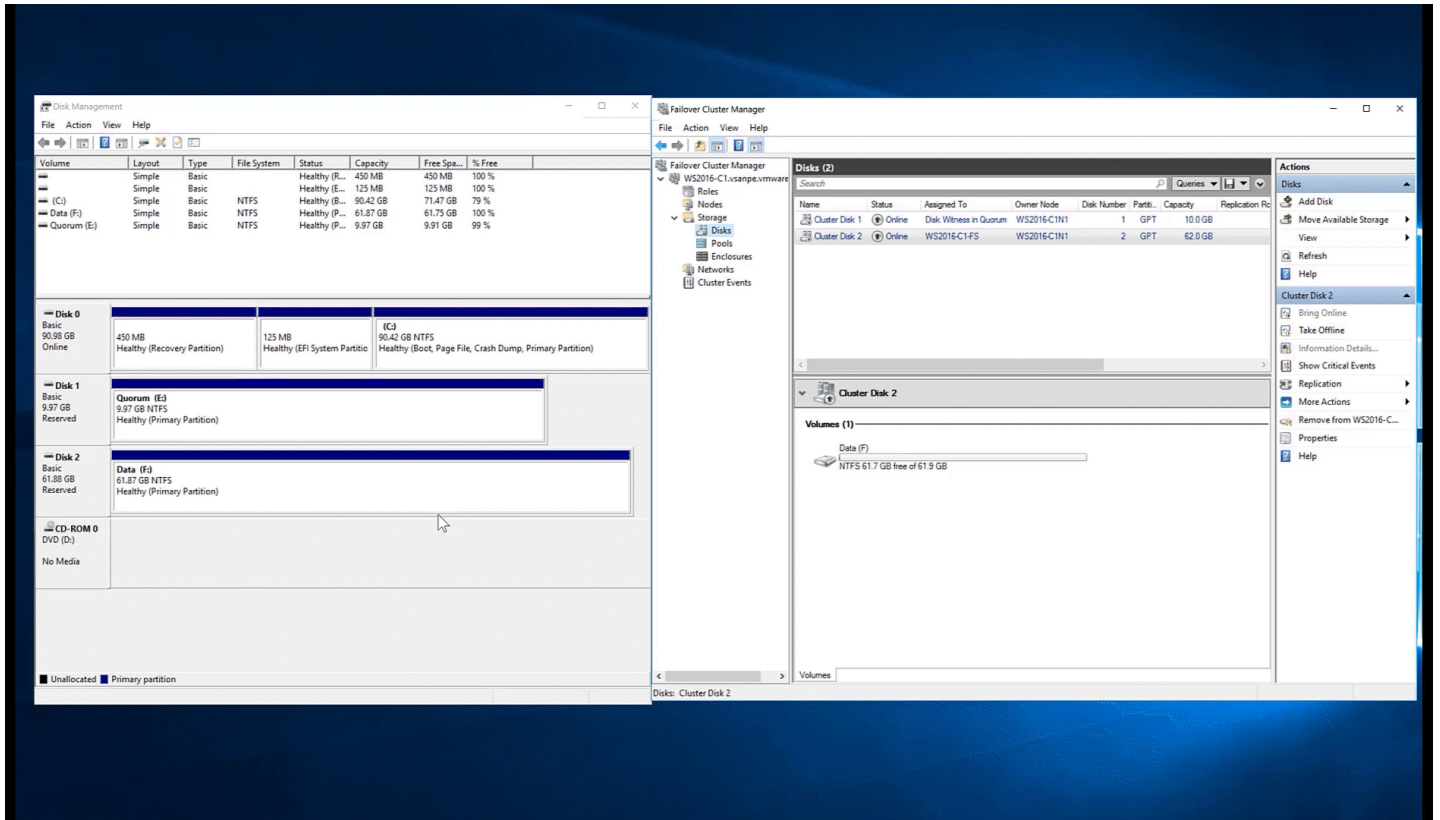
vVols



vVols Extend Online Shared Disks



With the release of vSphere 6.7, we added support for SCSI3-Persistent Reservations for vVols. This feature allows clustering applications to control the locking of shared disks. Microsoft WSFC is an example of an application requiring SCSI3-PR for the shared disks within the cluster. But one of the last features RDMs had over vVols was the ability to extend an online shared disk in applications like Microsoft WSFC. This has been a feature that has been requested at every VMworld and Explore as well as from numerous customers. **I'm very excited to announce vVols now supports the extension of online shared disk applications using SCSI3-Persistent Reservations!** This provides the customer with the ability to extend shared disks without having to shut down the application cluster. RDMs no longer have an edge over vVols. Customers can now migrate their MS WSFC applications to vVols and get rid of RDMs. This makes the virtual storage environment much easier to manage and reduces complexity while maintaining array-based features.



If you are interested in the process of migrating your WSFC from RDMs to vVols, see this article: [Migrating your WSFC from RDMs to vVols](#)

Support for vVols Online Disk Extension with Oracle RAC

In this release, in addition to the MS WSFC, we've added hot extend support for Oracle RAC disks, using Multi-writer Mode. There's no downtime required to extend the clustered disks. This can be performed on both SCSI and NVMe vVols disks. This again enables customers to migrate off RDMs.

vVols NVMe in-band Migration

Enables support for in-band migration of NVMe vVol namespaces between ANA groups. This functionality provides equivalence to SCSI rebind primitive and allows target/storage admin to balance IO loads across PEs.

Automated recovery from PDL for vVol PEs

Previously, when a PE went PDL and was later brought back, the PSA stack needed to recycle the device (destroy and detect paths again). This could only happen once the vVol objects using the PE are closed. With this release, the VMs will be auto-terminated in the cases where we detect that a PE to which the VM is bound vVols is PDL. This further enhances resilience and recovery during certain storage failures.

Enable 3rd party MPP support for NVMe vVols

Allows 3rd party Multi Path Policy (MPP) to support NVMe vVols. Enables our storage partners to use their custom MPP.

UNMAP Support for config vVol

Starting at vSphere 8.0U1, the config vVols are now created with thin provisioned size of 255GB and formatted with VMFS-6. With this release, we support command line (esxcli) based unmap support for config vVols.

NVMeoF

Clustered Applications



Support vNVME controller type for MS WSFC

In vSphere 7, we added a new VM vNVMe storage controller, but it was initially not supported for use with WSFC. In vSphere 8, we added support for clustered applications on NVMeoF Datastores. In vSphere 8 U2, our cluster validation team has now certified the vNVMe controller to be used with Microsoft WSFC. With this, you can now have end-to-end NVMe for your WSFC application. Initially, this will only be supported with SCSI vVols.

Support Oracle RAC clustering with vVols NVMe (FC, TCP)

We have added support for Oracle RAC clustered vVol disks to be hosted on NVMe backend. (NVMe-FC vVols, NVMe-TCP vVols).

Enable vNVMe support with Oracle RAC (vVols)

We added support for using the vNVMe controller as the frontend for multi-writer clustering solution (Oracle RAC)

VMFS

Improved SE Sparse snapshot offline consolidation performance.

The performance of SE Sparse snapshot offline consolidation has been improved by multiple factors in this release. This helps enhance RPO and RTO for customers using offline consolidation for DR purposes.

NFS

DNLC cache for NFS4.1

Some environments using NFS4.1 have large datastores housing hundreds of VMs, and doing searches, powering on, or listing VMs, can be slower than NFSv3. Directory Name Lookup Cache or DNLC is intended to reduce the number of NFS LOOKUP operations by caching some of this data. With this release, we've added DNLC support for NFS4.1. This will benefit operations like "ls" on a directory with a large number of VMs or files in a datastore.

nConnect

in vSphere 8.0 U1, we added support for nConnect which enables the ability to add multiple connections to an NFSv3 datastore. This can help reduce latency and increase performance. With this release, we've added support for dynamically increasing and decreasing the number of connections used with nConnect. Currently, this is configurable via esxcli only.

[Preliminary support for nConnect feature added in ESXi NFS client \(91497\) \(vmware.com\)](#)

When adding a new NFSv3 datastore, the number of connections can be specified at the time of the mount using the command:

- `esxcli storage nfs add -H <host> -v <volume-label> -s <remote_share> -c <number_of_connections>`

The maximum number of connections per session is limited to "4" by default. However, it can be increased to "8" using the advanced NFS option.

- `esxcfg-advcfg -s 8 /NFS/MaxConnectionsPerDatastore`
- `esxcfg-advcfg -g /NFS/MaxConnectionsPerDatastore`

The total number of connections used across all mounted NFSv3 datastores is limited to 256.

For an existing NFSv3 datastore, the number of connections can be *increased or decreased* at run time using the following command:

- `esxcli storage nfs param set -v <volume-label> -c <number_of_connections>`
-

vSphere 8 Update 1

New vSphere 8 Update 1 vVols and Core Storage enhancements, features, and additions.

vSphere 8 U1 storage has several storage enhancements in vVols, NVMeoF, VMFS, and NFS.

vSphere Virtual Volumes (vVols)



Continuing with vVols as a priority, we have added additional features, capabilities, and enhancements in vSphere 8 Update1.

One of the feature enhancements is a new certificate management framework. This will simplify the ability to register multiple vCenters to a single VASA provider. This lays the groundwork for potential future capabilities such as vVols vMSC. Some of the other features focus on scalability and performance. Because vVols can scale much larger than traditional storage, we want to ensure vVols will perform at scale as well.

Multi VC deployment for VASA Provider without Self-Signed Certificate

The new VASA 5 spec was developed to enhance vVols certificate management enabling self-signed certificates for multi-vCenter deployments. The solution also addresses certificate management where independent vCenter deployments running with different certificate management can work together. For example, it might be possible that one vCenter uses a 3rd party CA and another vCenter uses the VMCA signed certificate. This kind of deployment can be useful in shared VASA Provider deployment. This new capability will utilize Server Name Indication (SNI).

Server Name Indication (SNI) is an extension to the Transport Layer Security (TLS) computer networking protocol by which a client indicates which hostname it is attempting to connect at the start of the handshaking process. This enables a server to present multiple certificates on the same IP address and TCP port. Subsequently, allows multiple secure (HTTPS) websites (or other services over TLS) to be served by the same IP address without requiring all those sites to use the same certificate. It is the conceptual equivalent to HTTP/1.1 name-based virtual hosting, but for HTTPS. This also allows a proxy to forward client traffic to the right server during TLS/SSL handshake.

New vVols VASA 5.0 Spec

Specific features added to vSphere 8 U1

- **Container isolation** - Specific to the vendor's VASA provider.
- **Uptime** - Support alarm of certificate change and VASA 5.0 workflow which allows the VMCA cert to be refreshed in multi-VC systems.
- **Better Security for Shared Environment** - Specific to the vendor's VASA provider.
- **Backward Compatibility** - ESX which supports VASA 5.0, also resolves the Self Signed Certificate issues and downtime.
- **Heterogeneous Certificate Configuration** - Specific to the vendor's VASA provider.
- **Zero User Intervention** - Multi-VC support VMCA provisioning which does not require the user to install and manage certificates for the VP. SMS would take care of certificate provisioning.
- **Security Compliance** - No Self Signed Certificate in the trusted roots solving security compliance issues.

VASA 5.0 Feature Details

Container isolation - Allow per vCenter access control policy, even allowed to share containers in selected vCenter (Cross VC migration). Allow better isolation at the container level, VASA Provider can manage the access rights for a container per vCenter.

Uptime - Invalid/Expired Certificate causes downtime and there is the possibility of downtime when multiple vCenters try to register with the same VASA Provider. In a multi-VC setup, a certificate could be refreshed without any downtime.

Better Security for Shared Environment - All operations could be authenticated in the context of a vCenter and each vCenter would be having its own ACL (Access Control List). No self-signed certificate in the trust store. VASA Provider could be shared in a cloud-like environment. VASA 5.0, access role to do that.

Backward Compatibility - VASA 5.0 remains backward compatible and gives control to upgrade as per security needs. VASA 5.0 can coexist with earlier releases if the vendor will support it.

Heterogeneous Certificate Configuration - Uses only VMCA signed certificate and no additional CA is needed. Isolate the trust domain of vSphere. VASA 5.0 allows different configurations for each vCenter (i.e., Self-Managed 3rd Party CA Signed Certificate and VMCA managed Certificate).

Zero User Intervention - Plug and Play with automated certificate provisioning without any additional user intervention, no manual steps to use the VASA Provider with any vCenter.

Workload Separation - Allow load balancing by redirecting the transport for vCenter-specific virtual hosts. Each vCenter can be running in separate transport layer configurations. This creates more flexibility to isolate the workflow for each vCenter.

Security Compliance - Non-CA certificates are no longer part of the vSphere Certificate Trust Store. VASA 5.0. This enforces the VASA Provider to use CA signed certificate for VASA communication.

Move Sidecar vVols in config-vvol instead of another vVol object.

vVols Sidecars were created as vVol objects which introduced the overhead of VASA operations like bind/unbind. Solutions like First Class Disks (FCD) create a large number of small sidecars, which can degrade vVols performance. Also, since numerous sidecars can be created, this can count towards the total number of vVols objects supported on the storage array. To improve performance and scalability, these are now treated as files in the config-vvol and where normal file operations may be performed. Remember, in vSphere 8 we updated how the config-vvol was bound. Subsequently, with the config-vvol remaining bound, it reduces operations and delays, thus improving performance as well as scale.

With this new functionality introduced in this release, there are a few constraints regarding how the VM is created. Older vSphere releases will work with this new change, meaning existing VMs can run in the new format with updated hosts, but the new format will not work with older vSphere releases. **Newly created VMs or Virtual Disks using the new config-vVol/namespaces created by ESXi 8 U1 hosts are not supported on ESXi hosts of previous versions.**

Please refer to this knowledge base article to learn more about this functionality - [New config vVols objects created on ESXi 8.0U1 \(90791\) \(vmware.com\)](#)

Config-vvol Enhancements, Support for VMFS6 config-vvol with SCSI vVols (instead of VMFS5)

The config-vvol, which acts as a directory for the vVols datastore, and VM home contents, was capped at 4GB. This restricts the usage of folders with the vVols Datastore as content repositories. To overcome this, the config-vvol will now be created as 255GB thin-provisioned objects. Additionally, VMFS-6 will be used as the format for these objects instead of VMFS-5. This will enable the sidecar files, other VM files, and content libraries to be placed in the config-vvol.

In the image below, you can see the different sized config-vvols. For the Win10-5 VM, the config-vvol is using the original 4GB format. The Win10-vVol-8u1 VM is using the new 255GB config-vvol format.

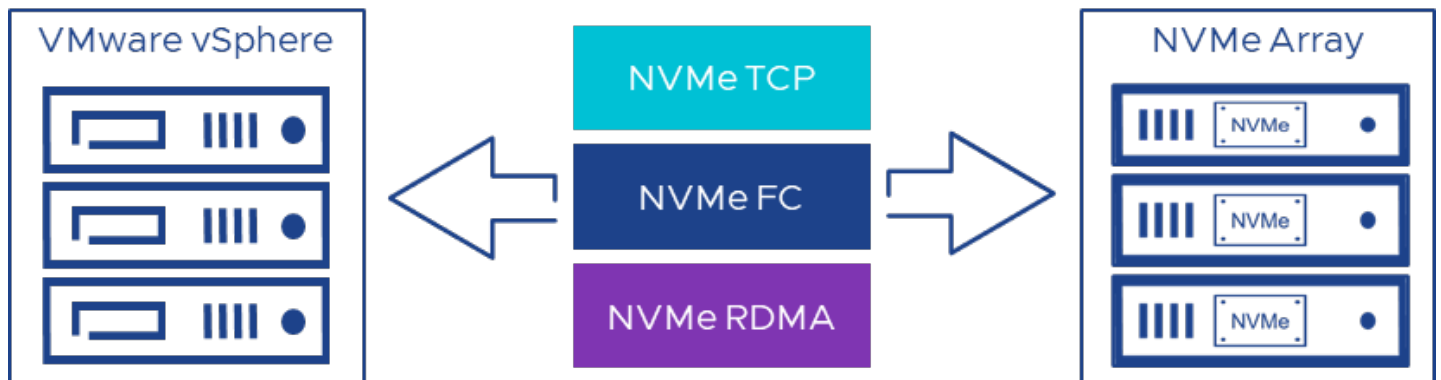
NAME	VM	THIN	TYPE	SIZE	ALLOCATED
Config-17ac7	...	Yes	MASTER	4 GiB	860 KiB
Config-1b328	Win10-vVol-8u1	Yes	MASTER	255 GiB	8.3 MiB
Config-a738	Win10-5	Yes	MASTER	4 GiB	6.1 MiB

Add NVMe-TCP support for vVols

NVMe over TCP support for vVols: In vSphere 8 we added NVMe-FC support for vVols. With vSphere 8 U1, we have validated NVMe-TCP further enabling the perfect union of NVMeoF and vVols. [See the article vVols with NVMe - A Perfect Match | VMware](#)

NVMeoF, PSA, HPP

Infrastructure to support **End to End NVMe**

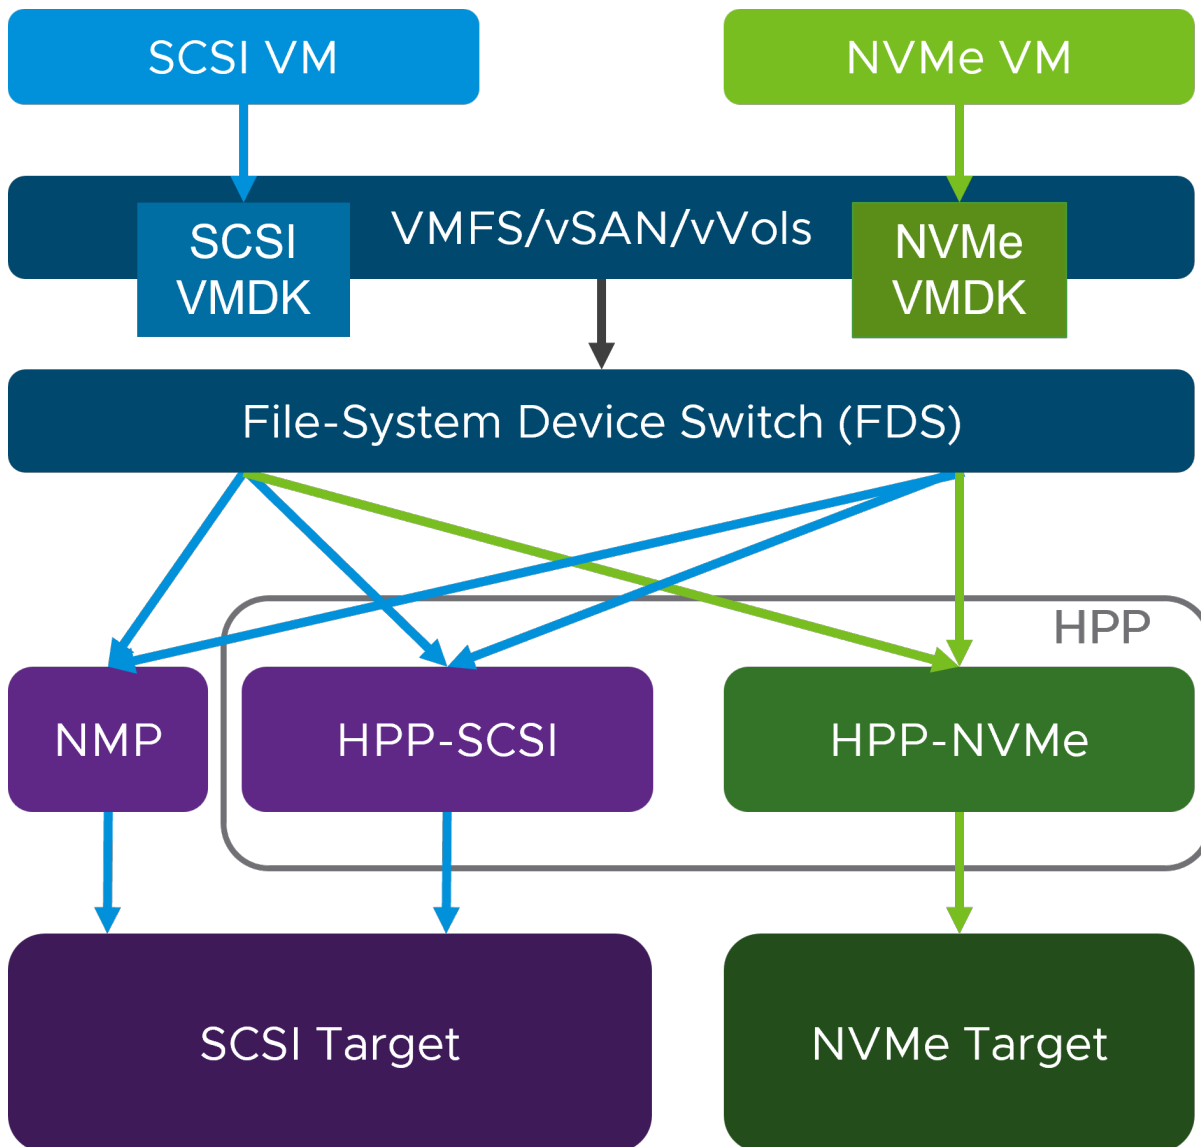


Extend NVMe capabilities to support end-to-end NVMe stack without any SCSI to NVMe translation in any of the ESXi layers. Another important aspect of supporting end-to-end NVMe is to allow third-party multi-pathing plugins to control and manage NVMe arrays.

This is a substantial final step in enabling the full capability of NVMe. Now, with supporting GOS, the NVMe protocol can be used from GOS to the end target.

A significant aspect of how VMware implemented the VM storage translation is full reverse compatibility support. We have made sure that with any combination of VM using either a SCSI or vNVMe controller and the target being either SCSI or NVMe, we can translate the storage stack path. This is a key design enabling customers to move between SCSI and NVMe Datastores without having to change the VM's storage controller. Similarly, if a VM has either a SCSI or vNVMe controller, it will work on both a SCSI and NVMeoF Datastore.

Simplified storage stack diagram.



For more information on NVMeoF for vSphere, see the [NVMeoF Resource page](#).

Increase Max Paths per NVMe-oF Namespaces from 8 to 32.

Increasing the number of paths helps scale with multiple paths to NVMe namespaces. This is needed in HA and scalability use cases where hosts can have multiple ports and the appliance could have multiple nodes per appliance and multiple ports per node.

Increase WSFC clusters per ESXi host from 3 to 16.

Increase the maximum number of WSFC clusters (multi-cluster) running on the same set of ESXi hosts from 3 to 16. With this increase, it can reduce the number of Microsoft WSFC licenses required by increasing the number of clusters that can run on a single host.

For more information on Microsoft WSFC on vSphere, here are some resources:

- [About Setup for Windows Server Failover Clustering on VMware vSphere](#)
- [Getting Started with WSFC on VMware vSphere](#)
- [Hardware and Software Requirements for WSFC on vSphere](#)
- [Microsoft Windows Server Failover Clustering \(WSFC\) with shared disks on VMware vSphere 8.x: Guidelines for supported configurations \(89327\)](#)

- [Microsoft Windows Server Failover Clustering \(WSFC\) with shared disks on VMware vSphere 7.x: Guidelines for supported configurations \(79616\)](#)

VMFS

Enhanced XCOPY to Datastores Across Different Storage Arrays.

ESXi now supports **Extended XCOPY** which optimizes the data copy between Datastores **across different arrays**. This will help customers offload migrate and clone workloads to the storage array arrays. While vSphere 8 U1 enables this feature, the actual data migration across the arrays must be supported on the storage array side.

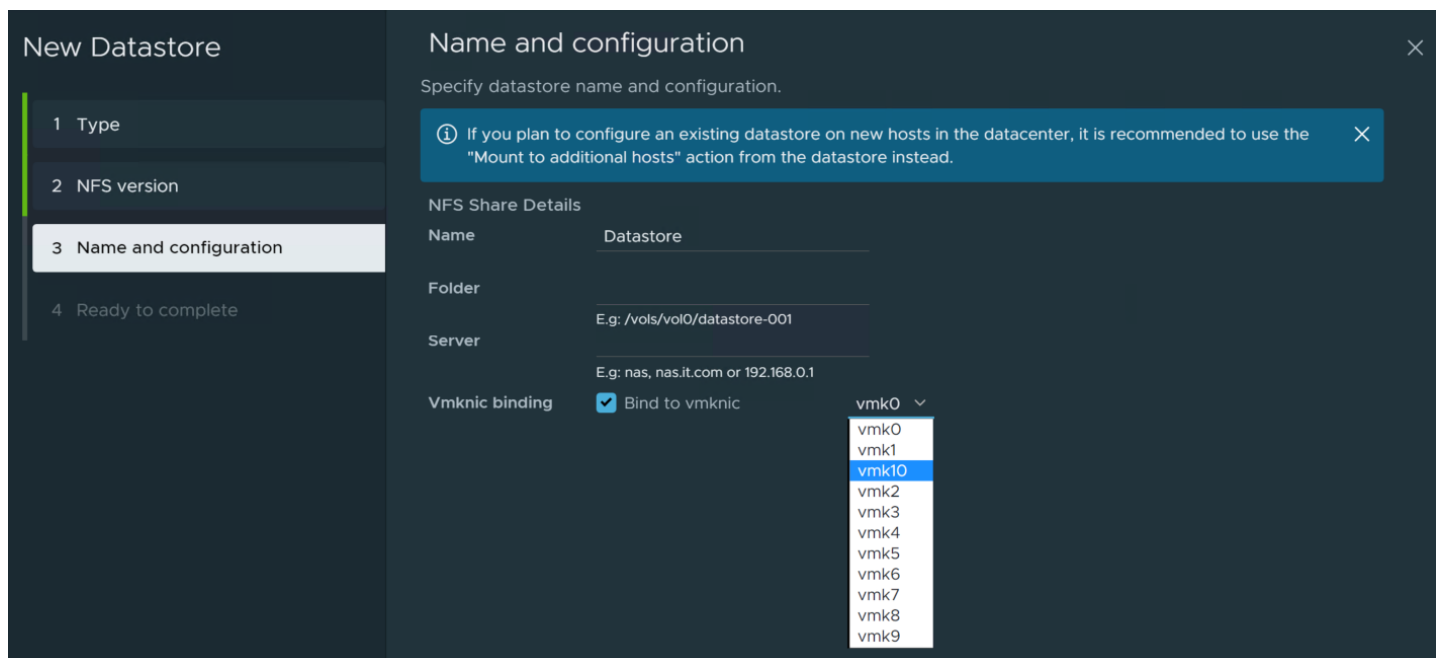
- [Using Extended XCOPY with vSphere 8.0 Update 1](#)

NFS

NFSv3 vmkPortBinding

This feature provides the ability to bind an NFS connection for a volume to a specific vmkernel. This helps with security by directing NFS traffic across a dedicated subnet/VLAN and ensures NFS traffic does not use the mgmt. or other vmkernels.

Previous NFS mounts will not have these values stored in config store. During the upgrade when the configuration is read from config store, *vmknic* and *bindToVmknic* values, if present, will be read. Upgrades from previous versions will not have these values since they are optional there will not be any impact.



- [ESXCLI command to update VMKernel binding of an NFSv3 datastore](#)

vSphere 8

New vSphere 8.0 Core Storage enhancements, features, and additions.

vSphere Virtual Volumes (vVols)



VM Swap improvements

- Faster Power on/off performance
- Faster vMotion performance

Changes in how the vVols swap is provisioned/destroyed have helped to improve power on/off as well as vMotion and svMotion performance.

Config vVol to remain bound.

- Helps reduce query times when looking for VM information.
- Caching various vVol attributes size, name, etc.

The config-vvol is where the VM's home data resides. (vmx, nvram, logs, etc.) and is usually only accessed at boot or change. Previously we did what was called a lazy unbind and unbound the config-vvol when not in use. In some cases, some applications periodically access the config-vvol and a new bind operation was required. Keeping the config-vvol bound reduces latency in accessing the VM home data.

NVMeoF vVols

vVols has been the primary focus of VMware storage engineering for the last few releases, and with vSphere 8.0, it is no different. The biggest announcement in vSphere 8.0 core storage is adding vVols support in NVMeoF. Initially, we will support FC only but will continue to validate and support other protocols supported with vSphere NVMeoF. This is a new vVols Spec, VASA/VC framework - [VASA 4.0/vVols 3.0](#).

The reason for adding vVols support to NVMeoF is many of the array vendors, and the industry, are moving towards using or at least adding NVMeoF support for better performance and throughput. Subsequently, VMware is making sure vVols remains current with the latest storage technologies.

Another benefit of NVMeoF vVols is the setup. When deploying, once you register the VASA, the underlying setup is completed in the background, you only need to create the datastore. The virtual Protocol Endpoints (vPEs) and connections are all handled by the VASA, simplifying the setup.

Technical Details:

ANA Group (Asymmetrical Namespace Access)

With NVMeoF, the implementation of vVols is a bit different. With traditional SCSI based vVols, the Storage Container is the logic grouping of the vVol objects themselves. With NVMeoF, this will vary depending on how the array vendor implements. But in general, at the array, an ANA Group is a grouping of vVol namespaces. The array decides on the number of ANA Groups, each having a unique ANAGRPID within the NVM subsystem. Namespaces are allocated and active only on BIND request to the VASA Provider (VP). Namespaces are also be added to an ANA Group on a BIND request to VP. A namespace remains allocated/active

until the last host UNBINDs the vVol.

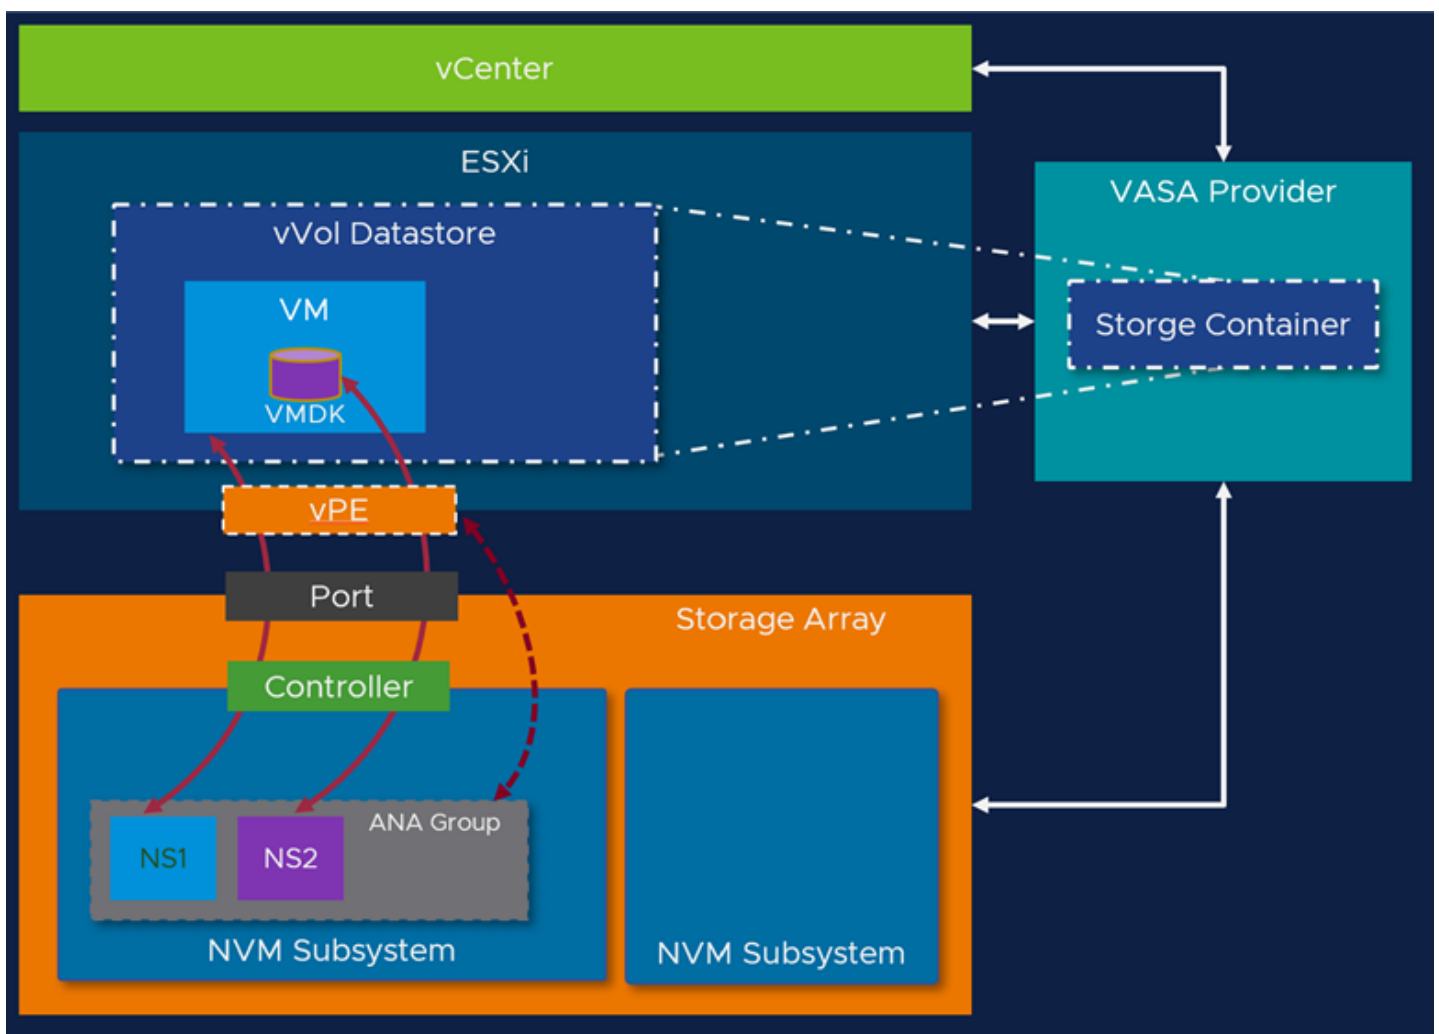
vPE (virtual Protocol Endpoint)

With traditional SCSI based vVols, the Protocol Endpoint (PE) is a physical LUN or volume on the array and shows up in the storage devices on the hosts. With NVMeoF vVols, there is no physical PE, the PE is now a logical object representation of the ANA group where the vVols reside. In fact, until a VM is powered on, the vPE doesn't exist. Array decides on the number of ANA Groups, each having a unique ANAGRPID within the NVM subsystem. Once a VM is powered on, the vPE is created so the host can access the vVols in the ANA group. You can see in the diagram the vPE points to the ANA group on the array.

NS (Namespace, NVMe equivalent to a LUN)

Each vVol type (Config, Swap, Data, Mem) created and used by a VM will create an NS that resides in an ANA group. It's a 1:1 vVol to NS ratio. This allows vendors to scale vVols objects easily. Typically, vendors support thousands to even hundreds of thousands of NS. The NS limits will be based on the array vendor.

Here in the diagram, you can see the VM itself is a NS, this would be the Config vVol, and the disk is another NS, a Data vVol.



Learn more about the technical details of NVMe-FC and vVols in this blog article [vVols with NVMe - A Perfect Match | VMware](#)


NVMeoF Enhancements



Support 256 namespaces and 2K paths with NVMe-TCP and NVMe-FC

NVMe over Fabrics (NVMeoF) continues to gain popularity for obvious reasons. Higher performance and throughput over traditional SCSI or NFS connectivity. Many storage partners are also moving to NVMe arrays, and using SCSI to access NVMe flash is a bottleneck for the potential gains.

•Continuing to add features and enhancements, VMware has increased the supported namespaces and paths for both NVMe-FC and TCP.



Extend reservation Support for NVMe device

Support NVMe reservation commands to enable solutions such as WSFC. This will allow customers to use Clustered VMDK capability for use with Microsoft WSFC with NVMeoF Datastores. Initially FC only.



Auto-discovery of NVMe Discovery Service support in ESXi

- Advanced NVMe-oF Discovery Service support in ESXi enables dynamic discovery of standards compliant NVMe Discovery Service.
- ESXi will use mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network.

ESXi sends a multicast DNS (mDNS) query requesting information from entities providing (NVMe) discovery service (DNS-SD). If such an entity is active on the network (on which the query was sent), it will send a (unicast) response to the host with the requested information - IP address and port number where the service is running.

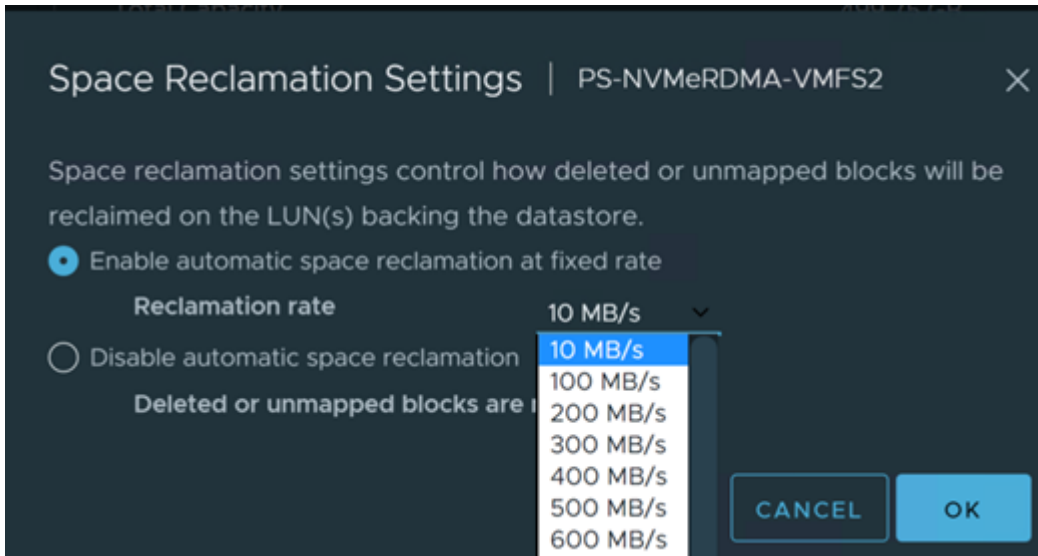
Unmap Space Reclamation Enhancement

Lower minimum reclamation rate to 10MBPS

Starting with vSphere 6.7 we added a feature to make the unmap rate configurable at the datastore level. With this enhancement, customers may change the unmap rate best suited for their array's capability and vendor recommendation. The higher unmap rate has benefited many array vendors quickly reclaim space. But we heard from some customers, even with the lowest unmap rate of 25 MB/sec, the rate can be disruptive when multiple hosts send unmap commands concurrently. The disruption can increase when scaling the number of hosts per datastore.

Example of potential overload: 25 MB/s * 100 datastores * 40 hosts ~ **104GB/s**

To help customers in situations where the 25MB/s unmap rate can be disruptive, we have reduced the minimum rate to 10MB/s, configurable per datastore.



This allows customers to reduce the potential impact of numerous unmap commands being sent to a single datastore. If needed, you can also disable space reclamation completely for a given datastore.

Dedicated Unmap Scheduling Queue

The dedicated unmap scheduling queue allows high priority VMFS metadata IOs to be separated and served from separate schedQueues to prevent them from getting starved behind UNMAP commands.

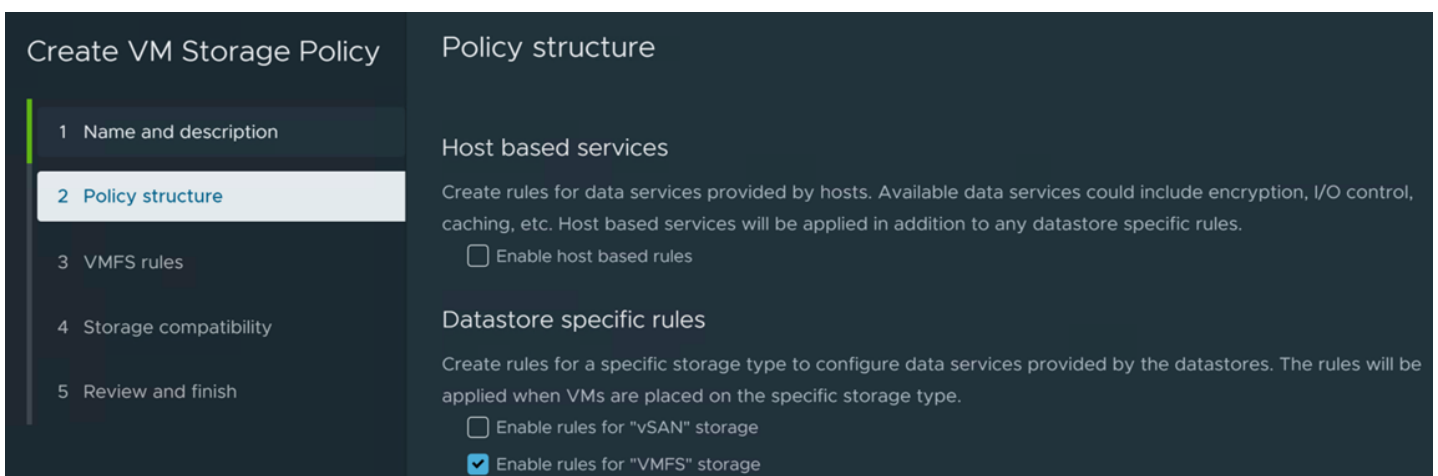
Container Storage CNS/CSI

VMFS and vSANDirect Disk Provisioning Storage Policy.

Choose EZT, LZT, or Thin provisioning via SPBM policy for CNS/Tanzu.

The goal is to add SPBM capability to support the creation/modification of storage policy rules to specify volume allocation options. It will also facilitate compliance checks in SPBM regarding the volume allocation rules in a storage policy.

- Operations supported for virtual disks are: create, reconfigure, clone, and relocate
- Operations supported for FCDs are: create, update storage policy, clone, and relocate.



VMFS rules

Placement Tags

Volume allocation type ⓘ

Conserve space when possible

Conserve space when possible

Fully initialized

Reserve space

Policy

Provision Type

FullyInitialized	All the blocks are allocated and zeroed (EZT)
ReserveSpace	All the blocks are allocated and not all blocks are zeroed (LZT)
ConserveSpaceWhenPossible	Not all the blocks are allocated (THIN)

Utilize SPBM provisioning rules for volume creation and support compliance checks.

NFS Enhancements

Engineering is always working to enhance storage resiliency. In vSphere 8 we have added NFS enhancements increasing the resilience through services, checks and permission validations.

- Retry NFS mounts on failure
- NFS mount validation

