



vVols FAQ

VMware Storage

Table of contents

vVols FAQ	3
Introduction and General Information	3
Technical Support	7
vVols and Microsoft WSFC	8
.....	9
Requirements and Capabilities	10

vVols FAQ

Introduction and General Information

Q: What are VMware vSphere Virtual Volumes (vVols)?

vVols are a new model for provisioning, managing, and accessing virtual disks for vSphere VMs. Using Storage Policy-Based Management (SPBM), VMs are provisioned with storage from vVols based arrays or servers through the mediation of a VASA Provider (VP) which provides an out-of-band interface for managing vVols based storage. vVols are designed to work equally over SCSI (FC, iSCSI, or FCoE) or NAS protocols. Storage arrays or servers are designed to manage all aspects of vVols storage and vSphere hosts have no direct access to a vVols storage. Instead, hosts access vVols through an intermediate point in the data path, the so-called "I/O Multiplexer" or "Protocol Endpoint" (PE) which is a SCSI target in the case of block-based vVols storage and an NFS mount point for NAS-based vVols storage.

Q: What is a VASA Provider?

The VASA Provider or VP is the endpoint of the out-of-band management connection. The VP is used to set up access to a vVol, communicate with VC and with ESX hosts as vVols are used, and ultimately release vVols. In addition, VASA commands to the VP can be used to snapshot vVols, revert a vVol to a previously created snapshot, as well as change the storage profile of an existing vVol.

Q: What is a Protocol Endpoint (PE)?

Protocol endpoints are the access points from the hosts to the storage systems, which are created by storage administrators. All paths and policies are administered by protocol endpoints. Protocol Endpoints are compliant with both, iSCSI and NFS. They are intended to replace the concept of LUNs and mount points. For more information, see the Working with Virtual Volumes section of the [VMware vSphere 6.0 Documentation](#).

Q: What is a storage container and how does it relate to a Virtual Datastore?

A vVols Storage Container is a logical abstraction on to which vVols are mapped and stored. Storage containers are setup at the array level and associated with array capabilities. vSphere will map the Storage Container to Virtual Datastore and provide applicable datastore level functionality. The Virtual Datastore is a key element and it allows the vSphere Admin to provision virtual machines without depending on the Storage Admin. Moreover, the Virtual Datastore provides logic abstraction for managing large numbers of vVols. This abstraction can be used for better managing multi-tenancy, various departments within a single organization, etc.

Q: How many Storage Containers can I have per storage array?

It depends on how a given array is configured. There is a limit of 256 storage containers per host. For more information, see the [VMware vSphere 6.0 Configuration Maximums Guide](#) and contact your Storage Array vendor for additional details.

Q: Can a single Virtual Datastore span different physical arrays?

No. However, if your storage vendor presents multiple physical arrays with 1 logical or virtual array then since that is still technically 1 (virtual) array you could potentially have a datastore span different physical arrays.

Q: How does a PE function?

A PE represents the IO access point for a vVol. When a vVol is created, it is not immediately accessible for IO. To access vVols, vSphere needs to issue a "Bind" operation to a VP, which creates an IO access point for a vVol on a PE chosen by a VP. Single PE can be IO access point for multiple vVols. "Unbind" Operation will remove this IO access point for a given vVol.

Q: What is the association of the PE to storage array?

PEs are associated with arrays. One PE is associated with one array only. An array can be associated with multiple PEs. For block arrays, PEs will be special LUN. ESX can identify these special LUNs and make sure the visible list of PEs is reported to VP. For NFS arrays, PEs are regular mount points.

Q: What is the association of a PE to storage containers?

PEs are managed per array. vSphere will assume that all PEs reported for an array are associated with all containers on that array. E.g: If the array has 2 containers and 3 PEs then ESX will assume that vVols on both containers can be bound on all 3 PEs. But internally VPs and storage Arrays can have some specific logic to map Virtual Volumes and storage containers

to PE.

Q: What is the association of a PE to Hosts??

PEs are like LUNs or mount points. They can be mounted or discovered by multiple hosts.

Q: Can I have one PE connect to multiple hosts across clusters?

Yes. VPs can return the same vVols binding information to each host if the PE is visible to multiple hosts.

Q: We have the multi-write VMDK feature today? How will that be represented in Virtual Volumes?

A vVol can be bound by multiple hosts. vSphere provides multi-writer support for vVols.

Q: Does vVols support Array-Based Replication

Yes. Starting with vSphere 6.5 Virtual Volumes offers support for array-based replication. Using supporting VASA providers, VI admins can configure data protection and disaster recovery in VM storage policies.

Q: Is VASA 2.0 a requirement for vVols support?

vVols does require a minimum of VASA 2.0. The version 2.0 of the VASA protocol introduces a set of APIs specifically for vVols that are used to manage storage containers and vVols. It also provides communication between vCenter, hosts, and storage arrays.

Q: Is VASA 3.0 now a requirement for vVols Support?

vVols supports both VASA 2.0 and VASA 3.0. VASA 3.0 introduces a new set of APIs that support disaster recovery configuration and operation. Using a VASA provider certified for VASA 3.0 users can configure data protection and disaster recovery into VM Storage Policies.

Q: Can Site Recovery Manager manage Virtual Volumes?

SRM can manage vVols that are using vSphere Replication or vVols array-based replication as of vSphere 7. Array-based replication is supported in vVols however configuration and execution of failover activities can be managed with public APIs or PowerCLI 6.5.

Q: How do vVols affect backup software vendors?

vVols are modeled in vSphere exactly as today's virtual disks. The VADP APIs backup vendors use are fully supported on vVols just as they are on vmdk files on a LUN. Backup software using VADP should be unaffected. This blog post provides more information on [backing up Virtual Volumes](#).

Q: Is vSAN using some of the vVols features under the covers?

Although vSAN presents some of the same capabilities (representing virtual disks as objects in storage, for instance) and introduces the ability to manage storage SLAs on a per-object level with SPBM, it does not use the same mechanisms as vVols. vVols uses VASA 2.0 and 3.0 to communicate with an array's VASA Provider to manage vVols on that array but vSAN uses its own APIs to manage virtual disks. SPBM is used by both, and SPBM's ability to present and interpret storage-specific capabilities lets it span vSAN's capabilities and vVols array's capabilities and present a single, uniform way of managing storage profiles and virtual disk requirements.

Q: Can you modify the size of a Storage Container on the fly?

Storage Containers are a logical entity only and are entirely managed by the storage array. You can think of it as a quota as opposed to a pre-allocated disk size of a LUN. In theory, there's nothing to prevent them from growing and shrinking on the fly. That capability is up to the array vendor to implement.

Q: Where are the Protocol Endpoints (PE) setup? In the vCenter with vSphere web client?

PEs are configured on the array side and vCenter is informed about them automatically through the VASA Provider. Hosts discover SCSI PEs as they discover today's LUNs; NFS mount points are automatically configured. For detailed guidance be sure to check with the respective storage vendor.

Q: Where are the array policies (snap, clone, replicate) applied?

Each array will have a certain set of capabilities supported (snapshot, clone, encryption, etc) defined at the storage container level. In vSphere, a VM Storage Policy is a combination of multiple storage capabilities. When a VM is provisioned with a particular VM Storage Policy, recommended datastores matching the policy are presented.

Q: From within vSphere, is a vVols Datastore accessed like a LUN? (storage browser, vm logs, vm.vpx config file, etc)

You can browse a vVols Datastore as you browse any other datastore.

Q: Is there a maximum number of vVols or maximum capacity for an SC/PE?

Those limits are entirely determined by the array vendor's implementation. vVols implementation does not impose any particular limits. You can have up to 64,000 vVols per host and a vVols datastore can be up to a zettabyte.

Q: Can you modify the size of the PE on the fly?

No. The Protocol Endpoint or PE is a protocol-specific access point that's used to provide a path to the vVols data. They're just a conduit for data traffic and as such, there is no reason to modify the size.

Q: Does the virtual disk vVol contain both the .vmdk file and the -flat.vmdk file as a single object?

The .vmdk file (the virtual disk descriptor file) is stored in the config vVol with the other VM description information (vmx file, log files, etc.). The vVol object takes the place of the flat file. The vmdk file has the ID of the vVol.

Q: How many vVols will be created for a VM? Do we have a different vVol for flat.vmdk and different vVol for .vmx files etc?

There is typically a minimum of 3 vVols per VM. The maximum would depend on how many virtual disks and snapshots reside on the VM,

- Config VVol: 1 per VM, holds the information previously in the VM's directory, i.e. vmx file, VM logs, etc
- DATA VVol: 1 for every virtual data disk. Analogous to the VMDK.
- SWAP VVol: 1 for Swap file. Created when the VM is powered on
- SNAP VVol: 1 for each Snapshot
- Other VVol: vSphere specific solution type

As an example the following SQL server has 8 vVols associated with it:

SQL01 vVols

1. Config vVol
2. Data vVol for the OS
3. Data vVol for the Database
4. Data vVol for the Log
5. Swap vVol created when powered on
6. Snapshot vVol for OS
7. SnapShot vVol for Database
8. Snapshot vVol for Log

Q: From a vSphere perspective, are snapshots allowed to be created at individual VMDK (vVol) level?

The APIs provide for snapshots of individual vVols but note that the UI only provides snapshots on a per-VM basis, which internally translates to requests for (simultaneous) snapshots of all a VM's virtual disks. It's not per LUN, though.

Q: Is there documentation that shows exactly which files belong to each type of vVol object?

There are no "files". vVols reside natively on the storage array and are referenced by the VM as it powers on. There is metadata information linking a vVol to a particular VM, so a management UI could be created for the vVols of a VM.

Q: Is there any NFS or SCSI conversions going on under the PE or array side?

Storage Containers are configured and the PEs set up to use NFS or SCSI for data in advance. The NFS traffic is unaltered using a mount point and a path. SCSI I/O is directed to a particular Virtual Volume using the secondary LUN ID field of the SCSI command.

Q: Is vVols suitable for Business Critical Applications?

Yes. vVols enables storage-level snapshots with VM granularity. Every VMDK is represented by an independent virtual volume and snapshots create a point-in-time copy of the volume that can be used for backup, recovery, and cloning activities. vVols also provides Storage Policy Based Management capability for automating the provisioning of database Virtual Machines which can be used for creating different storage tiers for business-critical databases. vSphere 6.5 introduces support for Oracle RAC 11gR2 and 12cR1

Q: How do you see the vVol ID?

Finding the UUID doesn't provide you with anything you can act upon. The UUID can be different between hosts based on the PE (Primary LUN or T10: Administrative Logic Unit). The UUID is used by the vVols architecture for tracking the actual

vVol (T10: Subsidiary Logical Unit) per PE. That is why you won't find general commands for showing the vVol UUID.

Q: How do we find out which VM is using a LUN?

In PowerCLI you can use the get-harddisk, but this will only show the PE where the vVol is bound.

Q: How do I troubleshoot vVols?

Docs.vmware.com has an article on troubleshooting vVols [here](#).

Q: Does vVols support NFS 4.1?

Yes, vVols supports NFS 4.1 if the storage vendor's implementation of vVols supports NFS 4.1. See the article in docs.vmware.com [NFS Protocols and ESXi](#).

Q: Can you run vCenter and or the VASA appliance on a vVols Datastore?

Running your vCenter or a VASA appliance on your vVols Datastore will work, but isn't a best practice. Because you need your vCenter and VASA online to manage your vVols infrastructure, if you place your vCenter on a vVols Datastore it becomes a circular dependency. The best practice is to place your vCenter and possibly VASA appliance on a traditional VMFS or NFS Datastore. This way, in a catastrophic failure, you can quickly recover your vCenter and/or VASA provider, point to the array and all the vVols objects will easily be available.

You could still get to the vVol object if your vCenter or VASA was on vVols, you would just have to build a new vCenter and VASA, point to the array, then recover your original vCenter and VASA if applicable.

Technical Support

Q: I use multi-pathing policies today. How do I continue to use them with vVols?

All multi-pathing policies today will be applied to PE devices. This means if path failover happens, it is applicable to all Virtual Volumes bound on that PE. Multi-pathing plugins have been modified not to treat internal vVols error conditions as path failure. vSphere makes sure that older MPPs does not claim PE device.

Q: Is there a way to mount a vVol from the esxcli command?

Depends. Config vVols can be “mounted” (bound) by accessing it through osfs (regular ‘ls’ into the path of the config vVol) but other vVol types aren’t accessible that way.

Q: If I clone or snapshot a single vVol in the storage array, is there any way to mount this extra vVol back to the ESXi host/VM as an additional volume?

Yes. A vSphere API exists to import unmanaged snapshots or vVols into vSphere. The imported object can be a data vVol, not necessarily a snapshot. The import of config-vVol (and other non-data vVol types) is unsupported. Storage vendors should ensure that the imported object behaves like a newly created vVol, from the vSphere perspective. The vCenter plugin can call the `ImportUnmanagedSnapshot` method, passing as parameters a datastore path (vdisk), datacenter (optional), and the array-assigned UUID of the previously unmanaged vVol (vvolId). The imported vVol specified by vvolId must be writable because an unmanaged vVol has no associated metadata, vSphere sets its metadata and creates a VMDK descriptor file for the imported vVol at the provided datastore path. The vSphere administrator can then add the resulting VMDK to any virtual machine using the “add disk” workflow. In the following call, datastore path is just an example:

```
importUnmanagedSnapshot("[datastore1]/vm1/imported-vm1.vmdk", Datacenter, vvolId)
```

Q: What's the process for migrating to vVols?

VMs using today's SAN or NAS datastores can be upgraded to vVol-based storage using storage vMotion to a vVol datastore.

Q: Is it possible to Storage vMotion from Containers on vVols capable Storage Arrays to LUNs on older arrays and back?

Yes, you can use storage vMotion to migrate between different datastore types.

Q: Can I copy policies from one VM to another?

Yes. Dynamically.

Q: If a vVol VM is created fresh on ESXi 8.0U1 using new VASA spec, will it be possible to vMotion the VM to older ESXi hosts (8.0 GA or 7.0U3)?

No, with the new VASA spec it is not backwards compatible. You should upgrade all hosts before using the new vVols VASA Spec in vSphere 8 U1.

Q: What is the best way to get the config vVol increased in size from 4GB to 256GB?

VM's created after upgrading to vSphere 8 U1 will automatically have the new 255GB thin provisioned config-vvol. Previous VM's will retain the 4GB config-vvol unless rebuilt.

vVols and Microsoft WSFC

Does vVols support Microsoft WSFC?

With the release of vSphere 6.7 we added SCSI3-PR support with vVols. With that, we also added support for Microsoft WSFC.

Is there a doc on the details for vVols and WSFC?

Yes, docs.vmware.com has an article on vVols and WSFC [here](#).

Is it possible to storage vMotion VMs with RDMS

Within vSphere, when there is a shared disk between VM storage vMotion is not supported.

What backup methods are supported with vVols and WSFC?

With vVols and WSFC using shared disks, you must use an in-guest backup solution.

Can you migrate a WSFC from pRDMs to vVols?

Yes, although there are some caveats. The WSFC and all nodes must be powered off. (Cold migration) Details on the process can be found [here](#).

When setting up a WSFC, what controller mode is required?

SCSI bus sharing mode must be set to physical. Make sure to follow the [vSphere MSCS Setup Checklist](#) when migrating or setting up your WSFC. More information on VMware and WSFC can be found at [About Setup for Failover Clustering and Microsoft Cluster Service](#).

Is vMotion supported with vVols and WSFC?

Yes, you can vMotion VMs within the same vSphere cluster.

Can you resize disks with the WSFC online?

Currently, resizing shared VMDKs with vVols is not supported.

What are the vendor requirements to support WSFC on vVols?

The storage array vendor must support SCSI3-PR type WEAR with vVols

What disk provisioning types are supported?

With vVols, both thick and thin are supported. Which type is supported is specific to the storage array vendor.

What transport protocols are supported with vVols?

WSFC on vVols supports FC, iSCSI, and FCoE.

Is Cluster-in-a-Box (CiB) supported?

- Placing all VMs, nodes of a WSFC on the same ESXi host (i.e. Cluster-in-a-Box (CiB) is not supported).
- VMs, nodes of a WSFC, must be placed on different ESXi hosts (i.e. Cluster Across Boxes (CAB)). The placement must be enforced with DRS MUST Anti-Affinity Rules.

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

What VMware vSphere features are NOT supported for WSFC?

- Live Storage vMotion.
- Fault Tolerance (FT).
- N-Port ID Virtualization (NPIV).
- Mixed versions of ESXi hosts in a vSphere cluster in production use.

What are the WSFC limits with vVols?

- 5 VM in a WSFC
- 3 WSFCs per host
- 128 clustered VMDKs per host

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

Requirements and Capabilities

Q: What are the software and storage hardware requirements for vVols?

You need the VMware vSphere 6.0 bits and your equivalent array vendor Virtual Volumes bits. For more information, see the [VMware Compatibility Guide](#).

Q: Is there an extra license in vSphere for vVols?

No. Standard, Enterprise, and Enterprise Plus all support vVols.

Q: Where can I get the storage array vendor vVols bits?

Storage vendors are providing Virtual Volumes integration in different ways. Contact your storage vendor for more details or visit the website of your vendor for more information on Virtual Volumes integration.

Q: Which VMware Products are interoperable with Virtual Volumes (vVols)?

VMware Products that are interoperable with Virtual Volumes (vVols) are:

- VMware vSphere 6.0.x
- VMware vRealize Automation 6.2.x (formerly known as VMware vCloud Automation Center)
- VMware Horizon 6.1.x
- VMware vSphere Replication 6.0.x

Q: Which VMware Products are currently NOT interoperable with Virtual Volumes (vVols)?

VMware Products that are currently not interoperable with Virtual Volumes (vVols) are:

- VMware vRealize Operations Manager 6.0.x to 6.1.0 (formerly known as VMware vCenter Operations Manager) **vROps 8.1 now supports vVols** [vROps 8.1 announcement blog](#)
- VMware Site Recovery Manager 5.x to 6.1.0 **SRM 8.3 now support vVols** [Announcement blog for SRM 8.3](#)
- VMware vSphere Data Protection 5.x to 6.1.0
- VMware vCloud Director 5.x

Q: Which VMware Products have deployment considerations to be aware of with vVols?

VMware NSX for vSphere 6.x - The deployment of virtual machine workloads attached to the NSX networks and stored on vVol datastores is supported. The deployment of NSX infrastructure components (the NSX Manager and Controller instances) on vVol storage is currently not supported by VMware.

Q: Which VMware vSphere 6.0.x features are interoperable with vVols?

VMware vSphere 6.0.x features that are interoperable with Virtual Volumes (vVols) are:

- High Availability (HA)
- Linked Clones
- Native Snapshots
- NFS version 3.x
- Storage Policy-Based Management (SPBM)
- Storage vMotion
- Thin Provisioning
- View Storage Accelerator/Content-Based Read Cache (CBRC)
- vSAN (VSAN)
- vSphere Auto Deploy
- vSphere Flash Read Cache
- vSphere Software Development Kit (SDK)
- vSphere API for I/O Filtering (VAIO)
- vMotion

- xvMotion

VMware vSphere 6.7 vVols features supported:

- IPv6 support for management access to the VASA provider.
- SCSI-3 Persistent Group Reservations (PGRs) support for supported arrays.
- TLS 1.2 default VP security.

Q: Which VMware vSphere 6.0.x features are currently NOT interoperable with vVols?

VMware vSphere 6.0.x features that are not interoperable with Virtual Volumes (vVols) are:

- Fault Tolerance (FT)
- NFS version 4.1
- IPv6 (Now supported with vSphere 6.7)
- Microsoft Failover Clustering (Now supported with vSphere 6.7)
- Raw Device Mapping (RDM)
- Storage Distributed Resource Scheduler (SDRS)
- Storage I/O Control

Q: Can vSAN and vVols co-exist and if yes, is data migration possible between the two?

Yes, and yes. vVols and vSAN work very complementary together because they use the same Storage Policy Based Management.

Q: Do vVols conform to requirements of applications such as exchange and clusters which require block-based storage and as such is this a block-based storage type?

vVols are just the VM objects and derivatives which are controlled through the storage policy framework. Whatever the array can expose as a feature of the array or capability by grouping application-focused capabilities that are presented to vSphere in that fashion, then yes, you can have specific application policy offerings for things like exchange, SQL, etc. This all happens regardless of whether its block or IP based storage.

Q: Is VMware APIs for Storage Awareness (VASA) 2.0 a requirement for vVols support?

Yes, vVols does require VASA 2.0. The version 2.0 of the VASA protocol introduces a new set of APIs specifically for Virtual Volumes that are used to manage storage containers and vVols. It also provides communication between the vCenter Server, ESXi hosts, and the storage arrays. For more information on the list of certified Storage Arrays for vVols support, see the [VMware Certified Compatibility Guides](#) .

Q: Can I use SDRS/SIOC to provision vVols enabled arrays?

No. SDRS is not supported. SIOC for vVols is not currently supported, but many storage vendors have QoS or storage I/O control via vVols and their VASA.

See the following blog post for detailed information on why [Storage DRS is less compelling with vVols](#)

Q: Can I use VAAI enabled storage arrays along with vVols enabled arrays?

Yes, VMware vSphere will use VAAI support whenever possible. VMware mandates ATS support for configuring vVols on SCSI.

Q: Can I use legacy datastores along with vVols?

Yes, vVols can co-exist with NFS and VMFS on the same array and NFS, VMFS, and vSAN in the same vSphere cluster.

Q: Can I replace RDMs with vVols?

With vSphere 6.7, vVols now supports SCSI-3 Persistent Group Reservations (PGRs) and with that RDM in most cases can be replaced with vVols. See Page on [vSphere 6.7 and vVols](#). Whenever an application requires direct access to the physical storage device, a pass-through RDM is required to be configured. Virtual Volumes are not a replacement for pass-thru RDM (ptRDM). Virtual Volumes are superior to non-pass-thru RDM (nptRDM) in a majority of virtual disks use cases.

Q: Will PowerCLI provide support for native vVols cmdlets?

Yes, many of the PowerCLI cmdlets work on vVols and there are some that are specific to vVols. Here's a blog on VirtualBlocks with some examples. [Automating vVols DR with Power CLI 6.5](#)

