



# vVols and VASA Related Fixes by ESXi Release

VMware Storage

## Table of contents

vVols and VASA Related Fixes by ESXi Release .....	4
vSphere Patches and Updates for vVols and VASA .....	4
vVols and VASA on vSphere 8.0 .....	4
VMware ESXi 8.0 Update 2b   29 FEB 2024   Build 23305546 .....	4
VMware ESXi 8.0 Update 2   21 SEP 2023   IA ISO Build 22380479 .....	5
ESXi 8.0 Update 1a   JUNE 01 2023   Build 21813344 .....	5
ESXi 8.0   11 OCT 2022   Build 20513097 .....	5
vVols and VASA on vSphere 7.0 .....	5
VMware ESXi 7.0 Update 3o   28 SEP 2023   Build 22348816 .....	5
.....	6
ESXi 7.0 Update 3l   MAR 30 2023   ISO Build 21424296 .....	6
ESXi 7.0 Update 3i   DEC 08 2022   20842708 .....	6
VMware ESXi 7.0 Update 3f - Build 20036589 .....	7
VMware ESXi 7.0 Update 3c - Build 19193900 .....	7
VMware ESXi 7.0 Update 2 - Build 17630552 .....	7
VMware ESXi 7.0 - Update 1c - Build 17325551 .....	8
vVols and VASA on vSphere 6.7 .....	8
VMware ESXi 6.7 - Patch Release ESXi670-202011002 - Build 17167734 .....	8
VMware ESXi 6.7 - Patch Release ESXi670-202008001 - Build 16713306 .....	9
VMware ESXi 6.7 - Patch Release ESXi670-202004002 - Build 16075168 .....	10
VMware ESXi 6.7 - Patch Release ESXi670-201912001 - Build 15160138 .....	11
VMware ESXi 6.7 Update 3 - Build 14320388 .....	12
VMware ESXi 6.7 Update 2 - Build 13006603 .....	13
VMware ESXi 6.7 Update 1 - Build 10302608 .....	14
vVols and VASA on vSphere 6.5 .....	15
VMware ESXi 6.5 - Patch Release ESXi650-201912002 - Build 15256549 .....	15
VMware ESXi 6.5 Update 3 - Build 13932383 .....	15
VMware ESXi 6.5, Patch Release ESXi650-201811002 - Build 10884925 .....	16
VMware ESXi 6.5 Update 1 - Build 5969303 .....	17
vCenter Patches and Updates for vVols and VASA .....	19
What's New in vSphere 7 Core Storage .....	19
vVols and VASA on vCenter 8.0 .....	19
vCenter Server 8.0 Update 2b   29 FEB 2024   GA ISO Build 23319993 .....	19
vCenter Server 8.0a   15 DEC 2022   ISO Build 20920323 .....	19
vVols and VASA on vCenter 7.0 .....	20

vCenter Server 7.0 Update 3l | 30 MAR 2023 | ISO Build 21477706 ..... 20

vCenter Server 7.0 Update 3d - Build 19480866 ..... 20

VMware vCenter Server 7.0 Update 3d - Build 19480866 ..... 20

..... 20

VMware vCenter Server 7.0.0b - Build 16386292 ..... 20

vVols and VASA on vCenter 6.7 ..... 21

VMware vCenter Server 6.7 Update 3m - Build 17713310 ..... 21

VMware vCenter Server 6.7 Update 3l - Build 17138064 ..... 21

VMware vCenter Server 6.7 Update 3j - Build 16708996 ..... 21

VMware vCenter Server 6.7 Update 2 - Build 13010631 ..... 22

vVols and VASA on vCenter 6.5 ..... 22

VMware vCenter Server 6.5 Update 3k - Build 16613358 ..... 22

vCenter Server 6.5.0e - Build 5705665 ..... 22

vCenter Server 6.5.0b - Build 5178943 ..... 22

VMware vSphere 6.5 - Build 4564106 ..... 23

## vVols and VASA Related Fixes by ESXi Release

### vSphere Patches and Updates for vVols and VASA

With some patches or releases, there may be patches or fixes to vVols and or the VASA provider. This is a collection of those patches and fixes.

For details on ESXi Version, Releases, and build numbers refer to the KB Article: [Build numbers and versions of VMware ESXi/ESX \(2143832\)](#)

For updates and enhancements see the following white papers.

- [What's New with vSphere 8 Core Storage](#)
- [What's New in vSphere 7 Core Storage](#)
- [What's New in vSphere 6.7 Core Storage](#)
- [What's New in vSphere 6.5 Storage](#)

---

## vVols and VASA on vSphere 8.0

VMware ESXi 8.0 Update 2b | 29 FEB 2024 | Build 23305546

### **PR 3293464: The VMX service might fail during migration of virtual machines with vSphere Virtual Volumes if the volume closure on the source takes long**

- In some scenarios, such as virtual machines with vSphere Virtual Volumes and Changed Block Tracking (CBT) or Content Based Read Cache (CBRC) enabled, during storage migration, flushing pending I/Os on volume closure at the source might take long, like 10 sec. If clearing the pending I/Os is not complete by the time the VMX service tries to reach the volumes on the destinations host, the service fails.
- This issue is resolved in this release. The fix makes sure that all vSphere Virtual Volumes and disks are closed on the source host before the destination host tries to open them.

### **PR 3293470: In the vSphere Client, you see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore**

- In the vSphere Client, when you navigate to Monitor >Advanced >Datastore to see the advanced performance metrics of a virtual machine, you might see a flat line in the read/write IO latency stats for a vSphere Virtual Volumes datastore. The issue occurs due to incorrect calculation of the counters for the read/write IO latency as a cumulative average.
- This issue is resolved in this release.

### **PR 3293466: vSphere vMotion operations between ESXi 7.x hosts fail with the error "Destination failed to preopen disks"**

- If on an ESXi 7.x host you use an older Virtual Volumes storage provider, also called a VASA provider, it might not allow binding multiple vSphere Virtual Volumes in batches to optimize the performance of migration and power on. As a result, in the vSphere Client you might see the error Destination failed to preopen disks when running vSphere vMotion operations between 7.x hosts.
- This issue is resolved in this release.

### **If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible**

- Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option Config.HostAgent.ssl.keyStore.allowSelfSigned is set to false by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and

VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

- Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article 91387.

---

### VMware ESXi 8.0 Update 2 | 21 SEP 2023 | IA ISO Build 22380479

#### **VASA API version does not automatically refresh after upgrade to vCenter Server 8.0**

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to **vCenter Server > Configure > Storage Providers** and expand the **General** tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

---

### ESXi 8.0 Update 1a | JUNE 01 2023 | Build 21813344

#### **ESXi 8.0 | 11 OCT 2022 | Build 20513097**

- **VASA API version does not automatically refresh after upgrade to vCenter Server 8.0**

vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:

1. If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to vCenter Server > Configure > Storage Providers and expand the General tab of the registered VASA provider, you still see VASA API version 3.5.
2. If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

## vVols and VASA on vSphere 7.0

### VMware ESXi 7.0 Update 3o | 28 SEP 2023 | Build 22348816

#### **PR 3185560: vSphere vMotion operations for virtual machines with swap files on a vSphere Virtual Volumes datastore intermittently fail after a hot extend of the VM disk**

- vSphere vMotion operations for virtual machines with swap files on a vSphere Virtual Volumes datastore might fail under the following circumstances:
- A VM that runs on ESXi host A is migrated to ESXi host B, then the VM disk is hot extended, and the VM is migrated back to ESXi host A. The issue occurs due to stale statistics about the swap virtual volume.

- This issue is resolved in this release.

### Virtual Machine Management Issues

#### Virtual machine snapshot operations fail in vSphere Virtual Volumes datastores on Purity version 5.3.10

- Virtual machine snapshot operations fail in vSphere Virtual Volumes datastores on Purity version 5.3.10 with an error such as An error occurred while saving the snapshot: The VVol target encountered a vendor specific error. The issue is specific for Purity version 5.3.10.
- Workaround: Upgrade to Purity version 6.1.7 or follow vendor recommendations.

---

### ESXi 7.0 Update 3l | MAR 30 2023 | ISO Build 21424296

#### PR 3060661: Windows Server Failover Cluster (WSFC) applications might lose connectivity to a virtual volume-based disk after initiating a rebind operation on that disk

- Due to reservation conflicts, WSFC applications might lose connectivity to a virtual volume-based disk after a rebind operation on that disk.
- This issue is resolved in this release.

#### This patch updates the following issues:

- Due to reservation conflicts, WSFC applications might lose connectivity to a virtual volume-based disk after a rebind operation on that disk.

---

### ESXi 7.0 Update 3i | DEC 08 2022 | 20842708

#### PR 3006356: ESXi host fail with a purple diagnostic screen due to rebinding of virtual volumes

- In rare cases, vSphere Virtual Volumes might attempt to rebind volumes on ESXi hosts that have SCSI Persistent Reservations. As a result, the ESXi hosts fail with a purple diagnostic screen and an error such as Panic Message: @BlueScreen: PANIC bora/vmkernel/main/dlmalloc.c:4933 - Usage error in dlmalloc in the backtrace.
- This issue is resolved in this release.

#### PR 2981272: vSphere Client displays 0 KB regardless of the actual VMDK size of a virtual machine

- Due to a caching issue, in the vSphere Client you might see a VMDK size of 0 KB regardless of the actual size of virtual machines in a vSphere Virtual Volumes environment.
- This issue is resolved in this release.

#### This patch updates the following issues:

- In rare cases, vSphere Virtual Volumes might attempt to rebind volumes on ESXi hosts that have SCSI Persistent Reservations. As a result, the ESXi hosts fail with a purple diagnostic screen and an error such as Panic Message: @BlueScreen: PANIC bora/vmkernel/main/dlmalloc.c:4933 - Usage error in dlmalloc in the backtrace.
- Due to a caching issue, in the vSphere Client you might see a VMDK size of 0 KB regardless of the actual size of virtual machines in a vSphere Virtual Volumes environment.

### VMware ESXi 7.0 Update 3f - Build 20036589

- **PR 2944919: You see virtual machine disks (VMDK) with stale binding in vSphere Virtual Volumes datastores**

When you reinstall an ESXi host after a failure, since the failed instance never reboots, stale bindings of VMDKs remain intact on the VASA provider and vSphere Virtual Volumes datastores. As a result, when you reinstall the host, you cannot delete the VMDKs due to the existing binding. Over time, many such VMDKs might accrue and consume idle storage space.

This issue is resolved in this release.

- **PR 2882789: While browsing a vSphere Virtual Volumes datastore, you see the UUID of some virtual machines, instead of their name**

In the vSphere Client, when you right-click to browse a datastore, you might see the UUID of some virtual machines in a vSphere Virtual Volumes datastore, such as naa.xxxx, instead of their names. The issue rarely occurs in large-scale environments with a large number of containers and VMs on a vSphere Virtual Volumes datastore. The issue has no functional impact, such as affecting VM operations, backup, or anything than VM display in the vSphere Client.

This issue is resolved in this release.

- **PR 2956080: Read and write operations to a vSphere Virtual Volumes datastore might fail**

If a SCSI command to a protocol endpoint of a vSphere Virtual Volumes datastore fails, the endpoint might get an Unsupported status, which might be cached. As a result, following SCSI commands to that protocol endpoint fail with an error code such as 0x5 0x20, and read and write operations to a vSphere Virtual Volumes datastore fail.

This issue is resolved in this release.

---

### VMware ESXi 7.0 Update 3c - Build 19193900

- **vSphere Virtual Volume snapshot operations might fail on the source volume or the snapshot volume on Pure storage**

Due to an issue that allows the duplication of the unique ID of vSphere Virtual Volumes, virtual machine snapshot operations might fail, or the source volume might get deleted. The issue is specific to Pure storage and affects Purity release lines 5.3.13 and earlier, 6.0.5 and earlier, and 6.1.1 and earlier.

This issue is resolved in this release.

---

### VMware ESXi 7.0 Update 2 - Build 17630552

- **VMware vSphere Virtual Volumes statistics for better debugging**

With ESXi 7.0 Update 2, you can track performance statistics for vSphere Virtual Volumes to quickly identify issues such as latency in third-party VASA provider responses. By using a set of commands, you can get statistics for all VASA providers in your system, or for a specified namespace or entity in the given namespace, or enable statistics tracking for the complete namespace.

For more information, see [Collecting Statistical Information for vVols](#).

---

## VMware ESXi 7.0 - Update 1c - Build 17325551

- **PR 2654686: The vSphere Virtual Volumes algorithm might not pick out the first Config-VVol that an ESXi host requests**

In a vSphere HA environment, the vSphere Virtual Volumes algorithm uses UUID to pick out when multiple ESXi hosts might compete to create and mount a Config-VVol with the same friendly name at the same time. However, the Config-VVol picked by UUID might not be the first one that the ESXi host requests and this might create issues in the vSphere Virtual Volumes datastores.

The vSphere Virtual Volumes algorithm uses a timestamp rather than a UUID to pick out when multiple ESXi hosts might compete to create and mount a Config-VVol with the same friendly name at the same time.

### This patch updates the following issues:

- In a vSphere HA environment, the vSphere Virtual Volumes algorithm uses UUID to pick out when multiple ESXi hosts might compete to create and mount a Config-VVol with the same friendly name at the same time. However, the Config-VVol picked by UUID might not be the first one that the ESXi host requests and this might create issues in the vSphere Virtual Volumes datastores.

---

## vVols and VASA on vSphere 6.7

### VMware ESXi 6.7 - Patch Release ESXi670-202011002 - Build 17167734

- **PR 2649677: You cannot access or power on virtual machines on a vSphere Virtual Volumes datastore**

In rare cases, an ESXi host is unable to report protocol endpoint LUNs to the vSphere API for Storage Awareness (VASA) provider while a vSphere Virtual Volumes datastore is being provisioned. As a result, you cannot access or power on virtual machines on the vSphere Virtual Volumes datastore. This issue occurs only when a networking error or a timeout of the VASA provider happens exactly at the time when the ESXi host attempts to report the protocol endpoint LUNs to the VASA provider.

This issue is resolved in this release.

- **PR 2656196: You cannot use a larger batch size than the default for vSphere API for Storage Awareness calls**

If a vendor provider does not publish or define a max batch size, the default max batch size for vSphere API for Storage Awareness calls is 16. This fix increases the default batch size to 1024.

This issue is resolved in this release.

- **PR 2630045: The vSphere Virtual Volumes algorithm might not pick out the first Config-VVol that an ESXi host requests**

In a vSphere HA environment, the vSphere Virtual Volumes algorithm uses UUID to pick out when multiple ESXi hosts compete to create and mount a Config-VVol with the same friendly name at the same time. However, the Config-VVol picked by UUID might not be the first one that the ESXi hosts in the cluster requests and this might create issues in the vSphere Virtual Volumes datastores.

This issue is resolved in this release.

The vSphere Virtual Volumes algorithm uses a timestamp rather than a UUID to pick out when multiple ESXi hosts



compete to create and mount a Config-VVol with the same friendly name at the same time.

#### This patch updates the following issues:

- In rare cases, an ESXi host is unable to report protocol endpoint LUNs to the vSphere API for Storage Awareness (VASA) provider while a vSphere Virtual Volumes datastore is being provisioned. As a result, you cannot access or power on virtual machines on the vSphere Virtual Volumes datastore. This issue occurs only when a networking error or a timeout of the VASA provider happens exactly at the time when the ESXi host attempts to report the protocol endpoint LUNs to the VASA provider.
- In a vSphere HA environment, the vSphere Virtual Volumes algorithm uses UUID to pick out when multiple ESXi hosts compete to create and mount a Config-VVol with the same friendly name at the same time. However, the Config-VVol picked by UUID might not be the first one that the ESXi hosts in the cluster requests and this might create issues in the vSphere Virtual Volumes datastores.

### VMware ESXi 6.7 - Patch Release ESXi670-202008001 - Build 16713306

- **PR 2601778: When migrating virtual machines between vSphere Virtual Volume datastores, the source VM disks remain undeleted**

In certain cases, such as when a VASA provider for a vSphere Virtual Volume datastore is not reachable but does not return an error, for instance a transport error or a provider timeout, the source VM disks remain undeleted after migrating virtual machines between vSphere Virtual Volume datastores. As result, the source datastore capacity is not correct.

This issue is resolved in this release.

- **PR 2586088: A virtual machine cloned to a different ESXi host might be unresponsive for a minute**

A virtual machine clone operation involves a snapshot of the source VM followed by creating a clone from that snapshot. The snapshot of the source virtual machine is deleted after the clone operation is complete. If the source virtual machine is on a vSphere Virtual Volumes datastore in one ESXi host and the clone virtual machine is created on another ESXi host, deleting the snapshot of the source VM might take some time. As a result, the cloned VM stays unresponsive for 50 to 60 seconds and might cause disruption of applications running on the source VM.

This issue is resolved in this release.

- **PR 2337784: Virtual machines on a VMware vSphere High Availability-enabled cluster display as unprotected when power on**

If an ESXi host in a vSphere HA-enabled cluster using a vSphere Virtual Volumes datastore fails to create the .vSphere-HA folder, vSphere HA configuration fails for the entire cluster. This issue occurs due to a possible race condition between ESXi hosts to create the .vSphere-HA folder in the shared vSphere Virtual Volumes datastore.

This issue is resolved in this release.

- **PR 2583029: Some vSphere vMotion operations fail every time when an ESXi host goes into maintenance mode**

If you put an ESXi host into maintenance mode and migrate virtual machines by using vSphere vMotion, some operations might fail with an error such as `A system error occurred:` in the vSphere Client or the vSphere Web Client.

In the `hostd.log`, you can see the following error:

```
2020-01-10T16:55:51.655Z warning hostd[2099896] [Originator@6876 sub=Vcsvc.VMotionDst.5724640822704079343
opID=k3l22s8p-5779332-auto-3fvd3-h5:70160850-b-01-b4-3bbc user=vpxuser:<user>] TimeoutCb: Expired
```

The issue occurs if vSphere vMotion fails to get all required resources before the defined waiting time of vSphere Virtual

Volumes due to slow storage or VASA provider.

This issue is resolved in this release. The fix makes sure vSphere vMotion operations are not interrupted by vSphere Virtual Volumes timeouts.

### This patch updates the following issues:

- In certain cases, such as when a VASA provider for a vSphere Virtual Volume datastore is not reachable but does not return an error, for instance a transport error or a provider timeout, the source VM disks remain undeleted after migrating virtual machines between vSphere Virtual Volume datastores. As result, the source datastore capacity is not correct.
- A virtual machine clone operation involves a snapshot of the source VM followed by creating a clone from that snapshot. The snapshot of the source virtual machine is deleted after the clone operation is complete. If the source virtual machine is on a vSphere Virtual Volumes datastore in one ESXi host and the clone virtual machine is created on another ESXi host, deleting the snapshot of the source VM might take some time. As a result, the cloned VM stays unresponsive for 50 to 60 seconds and might cause disruption of applications running on the source VM.
- If an ESXi host in a vSphere HA-enabled cluster using a vSphere Virtual Volumes datastore fails to create the `.vsphere-HA` folder, vSphere HA configuration fails for the entire cluster. This issue occurs due to a possible race condition between ESXi hosts to create the `.vsphere-HA` folder in the shared vSphere Virtual Volumes datastore.
- If you put an ESXi host into maintenance mode and migrate virtual machines by using vSphere vMotion, some operations might fail with an error such as `A system error occurred:` in the vSphere Client or the vSphere Web Client. In the `hostd.log`, you can see the following error:  

```
2020-01-10T16:55:51.655Z warning hostd[2099896] [Originator@6876 sub=Vcsvc.VMotionDst.5724640822704079343 opID=k3l22s8p-5779332-auto-3fvd3-h5:70160850-b-01-b4-3bbc user=vpuser:<user>] TimeoutCb: Expired
```

The issue occurs if vSphere vMotion fails to get all required resources before the defined waiting time of vSphere Virtual Volumes due to slow storage or VASA provider.

---

## VMware ESXi 6.7 - Patch Release ESXi670-202004002 - Build 16075168

- **PR 2458201: Some vSphere Virtual Volumes snapshot objects might not get a virtual machine UUID metadata tag**

During snapshot operations, especially a fast sequence of creating and deleting snapshots, a refresh of the virtual machine configuration might start prematurely. This might cause incorrect updates of the vSphere Virtual Volumes metadata. As a result, some vSphere Virtual Volumes objects, which are part of a newly created snapshot, might remain untagged or get tagged and then untagged with the virtual machine UUID.

This issue is resolved in this release.

- **PR 2424969: If the first attempt of an ESXi host to contact a VASA provider fails, vSphere Virtual Volumes datastores might remain inaccessible**

If a VASA provider is not reachable or not responding at the time an ESXi host boots up and tries to mount vSphere Virtual Volumes datastores, the mount operation fails. However, if after some time a VASA provider is available, the ESXi host does not attempt to reconnect to a provider and datastores remain inaccessible.

This issue is resolved in this release.

- **PR 2424363: During rebind operations, I/Os might fail with NOT\_BOUND error**

During rebind operations, the source protocol endpoint of a virtual volume might start failing I/Os with a `NOT_BOUND` error even when the target protocol endpoint is busy. If the target protocol endpoint is in `WAIT_RBZ` state and returns a status `PE_NOT_READY`, the source protocol endpoint must retry the I/Os instead of failing them.

This issue is resolved in this release. With the fix, the upstream relays a `BUSY` status to the virtual SCSI disk (vSCSI) and the ESXi host operating system to ensure a retry of the I/O.

- **PR 2449462: You might not be able to mount a Virtual Volumes storage container due to a stale mount point**

If the mount point was busy and a previous unmount operation has failed silently, attempts to mount a Virtual Volumes storage container might fail with an error that the container already exists.

This issue is resolved in this release.

- **PR 2467765: Upon failure to bind volumes to protocol endpoint LUNs on an ESXi host, virtual machines on vSphere Virtual Volumes might become inaccessible**

If a VASA provider fails to register protocol endpoint IDs discovered on an ESXi host, virtual machines on vSphere Virtual Volumes datastores on this host might become inaccessible. You might see an error similar to `vim.fault.CannotCreateFile`. A possible reason for failing to register protocol endpoint IDs from an ESXi host is that the `SetPEContext()` request to the VASA provider fails for some reason. This results in failing any subsequent request for binding virtual volumes, and losing accessibility to data and virtual machines on vSphere Virtual Volumes datastores.

This issue is resolved in this release. The fix is to reschedule `SetPEContext` calls to the VASA provider if a `SetPEContext()` request on a VASA provider fails. This fix allows the ESXi host eventually to register discovered protocol endpoint IDs and ensures that volumes on vSphere Virtual Volumes datastores remain accessible.

- **PR 2429068: Virtual machines might become inaccessible due to wrongly assigned second level LUN IDs (SLLID)**

The `nfnc` driver might intermittently assign wrong SLLID to virtual machines and as a result, Windows and Linux virtual machines might become inaccessible.

This issue is resolved in this release. Make sure that you upgrade the `nfnc` driver to version 4.0.0.44.

#### This patch updates the following issues:

- During snapshot operations, especially a fast sequence of creating and deleting snapshots, a refresh of the virtual machine configuration might start prematurely. This might cause incorrect updates of the vSphere Virtual Volumes metadata. As a result, some vSphere Virtual Volumes objects, which are part of a newly created snapshot, might remain untagged or get tagged and then untagged with the virtual machine UUID.
- If a VASA provider is not reachable or not responding at the time an ESXi host boots up and tries to mount vSphere Virtual Volumes datastores, the mount operation fails. However, if after some time a VASA provider is available, the ESXi host does not attempt to reconnect to a provider and datastores remain inaccessible.
- During rebind operations, the source protocol endpoint of a virtual volume might start failing I/Os with a `NOT_BOUND` error even when the target protocol endpoint is busy. If the target protocol endpoint is in `WAIT_RBZ` state and returns a status `PE_NOT_READY`, the source protocol endpoint must retry the I/Os instead of failing them.
- If the mount point was busy and a previous unmount operation has failed silently, attempts to mount a Virtual Volumes storage container might fail with an error that the container already exists.
- If a VASA provider fails to register protocol endpoint IDs discovered on an ESXi host, virtual machines on vSphere Virtual Volumes datastores on this host might become inaccessible. You might see an error similar to `vim.fault.CannotCreateFile`. A possible reason for failing to register protocol endpoint IDs from an ESXi host is that the `SetPEContext()` request to the VASA provider fails for some reason. This results in failing any subsequent request for binding virtual volumes, and losing accessibility to data and virtual machines on vSphere Virtual Volumes datastores.

---

#### VMware ESXi 6.7 - Patch Release ESXi670-201912001 - Build 15160138

- **PR 2419339: ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore.**

- ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore. You see a message indicating a PF Exception 14 on the screen. This issue might affect multiple hosts.

This issue is resolved in this release.

- **PR 2419339: ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore**

ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore. You see a message indicating a PF Exception 14 on the screen. This issue might affect multiple hosts.

This issue is resolved in this release.

- **PR 2432530: You cannot use a batch mode to unbind VMware vSphere Virtual Volumes**

ESXi670-201912001 implements the `UnbindVirtualVolumes ()` method in batch mode to unbind VMware vSphere Virtual Volumes. Previously, unbinding took one connection per vSphere Virtual Volume. This sometimes led to consuming all available connections to a vStorage APIs for Storage Awareness (VASA) provider and delayed response from or completely failed other API calls.

This issue is resolved in this release.

### This patch updates the following issues:

- ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore. You see a message indicating a PF Exception 14 on the screen. This issue might affect multiple hosts.
- ESXi670-201912001 implements the `UnbindVirtualVolumes ()` method in batch mode to unbind VMware vSphere Virtual Volumes. Previously, unbinding took one connection per vSphere Virtual Volume. This sometimes led to consuming all available connections to a vStorage APIs for Storage Awareness (VASA) provider and delayed response from or completely failed other API calls.

---

## VMware ESXi 6.7 Update 3 - Build 14320388

- **PR 2312215: A virtual machine with VMDK files backed by vSphere Virtual Volumes might fail to power on when you revert it to a snapshot**

This issue is specific to vSphere Virtual Volumes datastores when a VMDK file is assigned to different SCSI targets across snapshots. The lock file of the VMDK is reassigned across different snapshots and might be incorrectly deleted when you revert the virtual machine to a snapshot. Due to the missing lock file, the disk does not open, and the virtual machine fails to power on.

This issue is resolved in this release.

- **PR 2363202: The monitoring services show that the virtual machines on a vSphere Virtual Volumes datastore are in a critical state**

In the vSphere Web Client, incorrect Read or Write latency is displayed for the performance graphs of the vSphere Virtual Volumes datastores at a virtual machine level. As a result, the monitoring service shows that the virtual machines are in a critical state.

This issue is resolved in this release.

- **PR 2402409: Virtual machines with enabled Changed Block Tracking (CBT) might fail while a snapshot is created due to lack of allocated memory for the CBT bit map**

While a snapshot is being created, a virtual machine might power off and fail with an error similar to:

```
2019-01-01T01:23:40.047Z| vcpu-0| I125: DISKLIB-CTK : Failed to mmap change bitmap of size 167936: Cannot allocate memory.
2019-01-01T01:23:40.217Z| vcpu-0| I125: DISKLIB-LIB_BLOCKTRACK : Could not open change tracker /vmfs/volumes/DATASTORE_UUID/VM_NAME/VM_NAME_1-ctk.vmdk: Not enough memory for change tracking.
The error is a result of lack of allocated memory for the CBT bit map.
```

This issue is resolved in this release.

### This patch updates the following issues:

- This issue is specific to vSphere Virtual Volumes datastores when a VMDK file is assigned to different SCSI targets across snapshots. The lock file of the VMDK is reassigned across different snapshots and might be incorrectly deleted when you revert the virtual machine to a snapshot. Due to the missing lock file, the disk does not open, and the virtual machine fails to power on.
- In the vSphere Web Client, incorrect Read or Write latency is displayed for the performance graphs of the vSphere Virtual Volumes datastores at a virtual machine level. As a result, the monitoring service shows that the virtual machines are in a critical state.

---

## VMware ESXi 6.7 Update 2 - Build 13006603

- **PR 2250697: Windows Server Failover Cluster validation might fail if you configure Virtual Volumes with a Round Robin path policy**

If during the Windows Server Failover Cluster setup you change the default path policy from Fixed or Most Recently Used to Round Robin, the I/O of the cluster might fail and the cluster might stop responding.

This issue is resolved in this release.

- **PR 2279897: Creating a snapshot of a virtual machine might fail due to a null VvolId parameter**

If a vSphere API for Storage Awareness provider modifies the vSphere Virtual Volumes policy unattended, a null `VvolId` parameter might update the vSphere Virtual Volumes metadata. This results in a VASA call with a null `VvolId` parameter and a failure when creating a virtual machine snapshot.

This issue is resolved in this release. The fix handles the policy modification failure and prevents the null `VvolId` parameter.

- **PR 2227623: Parallel cloning of multiple virtual machines on a vSphere Virtual Volumes datastore might fail with an error message for failed file creation**

If a call from a vSphere API for Storage Awareness provider fails due to all connections to the virtual provider being busy, operations for parallel cloning of multiple virtual machines on a vSphere Virtual Volumes datastore might become unresponsive or fail with an error message similar to `Cannot complete file creation operation`.

This issue is resolved in this release.

- **PR 2268826: An ESXi host might fail with a purple diagnostic screen when the VMware APIs for Storage Awareness (VASA) provider sends a rebind request to switch the protocol endpoint for a vSphere Virtual Volume**

When the VASA provider sends a rebind request to an ESXi host to switch the binding for a particular vSphere Virtual Volume driver, the ESXi host might switch the protocol endpoint and other resources to change the binding without any I/O disturbance. As a result, the ESXi host might fail with a purple diagnostic screen.

This issue is resolved in this release.

- **PR 2246891: An I/O error `TASK_SET_FULL` on a secondary LUN might slow down the I/O rate on all secondary LUNs behind the protocol endpoint of Virtual Volumes on HPE 3PAR storage if I/O throttling is enabled**

When I/O throttling is enabled on a protocol endpoint of Virtual Volumes on HPE 3PAR storage and if an I/O on a secondary LUN fails with an error `TASK_SET_FULL`, the I/O rate on all secondary LUNs that are associated with the protocol endpoint slows down.

This issue is resolved in this release. With this fix, you can enable I/O throttling on individual Virtual Volumes to avoid the slowdown of all secondary LUNs behind the protocol endpoint if the `TASK_SET_FULL` error occurs.

### This patch updates the following issues:

- If a vSphere API for Storage Awareness provider modifies the vSphere Virtual Volumes policy unattended, a null `vvolID` parameter might update the vSphere Virtual Volumes metadata. This results in a VASA call with a null `vvolId` parameter and a failure when creating a virtual machine snapshot.
- If a call from a vSphere API for Storage Awareness provider fails due to all connections to the virtual provider being busy, operations for parallel cloning of multiple virtual machines on a vSphere Virtual Volumes datastore might become unresponsive or fail with an error message similar to `Cannot complete file creation operation`.
- When the VASA provider sends a rebind request to an ESXi host to switch the binding for a particular vSphere Virtual Volume driver, the ESXi host might switch the protocol endpoint and other resources to change the binding without any I/O disturbance. As a result, the ESXi host might fail with a purple diagnostic screen.
- When I/O throttling is enabled on a protocol endpoint of Virtual Volumes on HPE 3PAR storage and if an I/O on a secondary LUN fails with an error `TASK_SET_FULL`, the I/O rate on all secondary LUNs that are associated with the protocol endpoint slows down.

---

## VMware ESXi 6.7 Update 1 - Build 10302608

- **PR 2039186: VMware vSphere Virtual Volumes metadata might not be updated with associated virtual machines and make virtual disk containers untraceable**

vSphere Virtual Volumes set with `VMW_VVOLType` metadata key `Other` and `VMW_VVOLTypeHint` metadata key `Sidecar` might not get `VMW_VmID` metadata key to the associated virtual machines and cannot be tracked by using IDs.

This issue is resolved in this release.

- **PR 2119610: Migration of a virtual machine with a Filesystem Device Switch (FDS) on a vSphere Virtual Volumes datastore by using VMware vSphere vMotion might cause multiple issues**

If you use vSphere vMotion to migrate a virtual machine with file device filters from a vSphere Virtual Volumes datastore to another host, and the virtual machine has either of the Changed Block Tracking (CBT), VMware vSphere Flash Read Cache (VFRC) or I/O filters enabled, the migration might cause issues with any of the features. During the migration, the file device filters might not be correctly transferred to the host. As a result, you might see corrupted incremental backups in CBT, performance degradation of VFRC and cache I/O filters, corrupted replication I/O filters, and disk corruption, when cache I/O filters are configured in write-back mode. You might also see issues with the virtual machine encryption.

This issue is resolved in this release.

- **PR 2145089: vSphere Virtual Volumes might become unresponsive if an API for Storage Awareness (VASA) provider loses binding information from the database**

vSphere Virtual Volumes might become unresponsive due to an infinite loop that loads CPUs at 100% if a VASA provider loses binding information from the database. Hostd might also stop responding. You might see a fatal error message.

This issue is resolved in this release. This fix prevents infinite loops in case of database binding failures.

- **PR 2146206: vSphere Virtual Volumes metadata might not be available to storage array vendor software**

vSphere Virtual Volumes metadata might be available only when a virtual machine starts running. As result, storage array vendor software might fail to apply policies that impact the optimal layout of volumes during regular use and after a

failover.

This issue is resolved in this release. This fix makes vSphere Virtual Volumes metadata available at the time vSphere Virtual Volumes are configured, not when a virtual machine starts running.

### This patch updates the following issues:

- vSphere Virtual Volumes set with `VMW_VVolType` metadata key `Other` and `VMW_VVolTypeHint` metadata key `Sidecar` might not get `VMW_VmID` metadata key to the associated virtual machines and cannot be tracked by using IDs.
- If you use vSphere vMotion to migrate a virtual machine with file device filters from a vSphere Virtual Volumes datastore to another host, and the virtual machine has either of the Changed Block Tracking (CBT), VMware vSphere Flash Read Cache (VFRC) or I/O filters enabled, the migration might cause issues with any of the features. During the migration, the file device filters might not be correctly transferred to the host. As a result, you might see corrupted incremental backups in CBT, performance degradation of VFRC and cache I/O filters, corrupted replication I/O filters, and disk corruption, when cache I/O filters are configured in write-back mode. You might also see issues with the virtual machine encryption.
- vSphere Virtual Volumes might become unresponsive due to an infinite loop that loads CPUs at 100% if a VASA provider loses binding information from the database. Hostd might also stop responding. You might see a fatal error message.
- vSphere Virtual Volumes metadata might be available only when a virtual machine starts running. As a result, storage array vendor software might fail to apply policies that impact the optimal layout of volumes during regular use and after a failover.

---

## vVols and VASA on vSphere 6.5

### VMware ESXi 6.5 - Patch Release ESXi650-201912002 - Build 15256549

- **PR 2271176: ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore**

ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore. You see a message indicating a PF Exception 14 on the screen. This issue might affect multiple hosts.

This issue is resolved in this release.

### This patch updates the following issues:

- ESXi hosts might fail with a purple diagnostic screen during power off or deletion of multiple virtual machines on a vSphere Virtual Volumes datastore. You see a message indicating a PF Exception 14 on the screen. This issue might affect multiple hosts.

---

## VMware ESXi 6.5 Update 3 - Build 13932383

- **PR 2282080: Creating a snapshot of a virtual machine from a virtual volume datastore might fail due to a null `VVolId` parameter**

If a vSphere API for Storage Awareness provider modifies the vSphere Virtual Volumes policy unattended, a null `VVolId` parameter might update the vSphere Virtual Volumes metadata. This results in a VASA call with a null `VVolId` parameter and a failure when creating a virtual machine snapshot.

This issue is resolved in this release. The fix handles the policy modification failure and prevents the null `VvolId` parameter.

- **PR 2265828: A virtual machine with VMDK files backed by vSphere Virtual Volumes might fail to power on when you revert it to a snapshot**

This issue is specific to vSphere Virtual Volumes datastores when a VMDK file is assigned to different SCSI targets across snapshots. The lock file of the VMDK is reassigned across different snapshots and might be incorrectly deleted when you revert the virtual machine to a snapshot. Due to the missing lock file, the disk does not open, and the virtual machine fails to power on.

This issue is resolved in this release.

- **PR 2113782: vSphere Virtual Volumes datastore might become inaccessible if you change the vCenter Server instance or refresh the CA certificate**

vSphere Virtual Volume datastore uses VMware CA signed certificate to communicate with VASA providers. When the vCenter Server instance or the CA certificate changes, vCenter Server imports the new vCenter Server CA signed certificate and the vSphere Virtual Volume datastore gets SSL reset signal which might not be triggered. As a result, the communication between vSphere Virtual Volume datastore and VASA providers might fail and the vSphere Virtual Volume datastore might become inaccessible.

This issue is resolved in this release.

- **PR 2278591: Cloning multiple virtual machines simultaneously on vSphere Virtual Volumes might stop responding**

When you clone multiple virtual machines simultaneously from vSphere on a vSphere Virtual Volumes datastore, a `setPEContext` VASA API call is issued. If all connections to the VASA Providers are busy or unavailable at the time of issuing the `setPEContext` API call, the call might fail and the cloning process stops responding.

This issue is resolved in this release.

### This patch updates the following issues:

- If a vSphere API for Storage Awareness provider modifies the vSphere Virtual Volumes policy unattended, a null `VvolId` parameter might update the vSphere Virtual Volumes metadata. This results in a VASA call with a null `VvolId` parameter and a failure when creating a virtual machine snapshot.
- This issue is specific to vSphere Virtual Volumes datastores when a VMDK file is assigned to different SCSI targets across snapshots. The lock file of the VMDK is reassigned across different snapshots and might be incorrectly deleted when you revert the virtual machine to a snapshot. Due to the missing lock file, the disk does not open, and the virtual machine fails to power on.
- When you clone multiple virtual machines simultaneously from vSphere on a vSphere Virtual Volumes datastore, a `setPEContext` VASA API call is issued. If all connections to the VASA Providers are busy or unavailable at the time of issuing the `setPEContext` API call, the call might fail and the cloning process stops responding.

---

## VMware ESXi 6.5, Patch Release ESXi650-201811002 - Build 10884925

- **PR 2119609: Migration of a virtual machine with a Filesystem Device Switch (FDS) on a VMware vSphere Virtual Volumes datastore by using VMware vSphere vMotion might cause multiple issues**

If you use vSphere vMotion to migrate a virtual machine with file device filters from a vSphere Virtual Volumes datastore to another host, and the virtual machine has either of the Changed Block Tracking (CBT), VMware vSphere Flash Read Cache (VFRC), or IO filters enabled, the migration might cause issues with any of the features. During the migration, the file device filters might not be correctly transferred to the host. As a result, you might see corrupted incremental backups in CBT, performance degradation of VFRC and cache IO filters, corrupted replication IO filters and disk corruption, when cache IO filters are configured in write-back mode. You might also see issues with the virtual machine encryption.



This issue is resolved in this release.

- **PR 2142767: VMware vSphere Virtual Volumes might become unresponsive if a vSphere API for Storage Awareness provider loses binding information from its database**

vSphere Virtual Volumes might become unresponsive due to an infinite loop that loads CPUs at 100% if a vSphere API for Storage Awareness provider loses binding information from its database. Hostd might also stop responding. You might see a fatal error message. This fix prevents infinite loops in case of database binding failures.

This issue is resolved in this release.

- **PR 2133634: An ESXi host might fail with a purple diagnostic screen during migration of clustered virtual machines by using vSphere vMotion**

An ESXi host might fail with a purple diagnostic screen during migration of clustered virtual machines by using vSphere vMotion. The issue affects virtual machines in clusters containing shared non-RDM disk, such as VMDK or vSphere Virtual Volumes, in a physical bus sharing mode.

This issue is resolved in this release.

#### This patch updates the following issues:

- If you use vSphere vMotion to migrate a virtual machine with file device filters from a vSphere Virtual Volumes datastore to another host, and the virtual machine has either of the Changed Block Tracking (CBT), VMware vSphere Flash Read Cache (VFRC), or IO filters enabled, the migration might cause issues with any of the features. During the migration, the file device filters might not be correctly transferred to the host. As a result, you might see corrupted incremental backups in CBT, performance degradation of VFRC and cache IO filters, corrupted replication IO filters and disk corruption, when cache IO filters are configured in write-back mode. You might also see issues with the virtual machine encryption.
- vSphere Virtual Volumes might become unresponsive due to an infinite loop that loads CPUs at 100% if a vSphere API for Storage Awareness provider loses binding information from its database. Hostd might also stop responding. You might see a fatal error message. This fix prevents infinite loops in case of database binding failures.

---

#### VMware ESXi 6.5 Update 1 - Build 5969303

- **When you use vSphere Storage vMotion, the UUID of a virtual disk might change**

When you use vSphere Storage vMotion on vSphere Virtual Volumes storage, the UUID of a virtual disk might change. The UUID identifies the virtual disk and a changed UUID makes the virtual disk appear as a new and different disk. The UUID is also visible to the guest OS and might cause drives to be misidentified.

This issue is resolved in this release.

- **Disabled frequent lookup to an internal vSAN metadata directory (.upit) on virtual volume datastores. This metadata folder is not applicable to virtual volumes**

The frequent lookup to a vSAN metadata directory (.upit) on virtual volume datastores can impact its performance. The .upit directory is not applicable to virtual volume datastores. The change disables the lookup to the .upit directory.

This issue is resolved in this release.

- **Non-Latin characters might be displayed incorrectly in VM storage profile names**

UTF-8 characters are not handled properly before being passed on to a VVol Vasa Provider. As a result, the VM storage profiles which are using international characters are either not recognized by the Vasa Provider or are treated, or displayed incorrectly by the Vasa Provider.

This issue is resolved in this release.

---

---

## vCenter Patches and Updates for vVols and VASA

vVols and or VASA patches and fixes for vCenter

For updates and enhancements see the following white papers.

- [What's New in vSphere 7 Core Storage](#)
- [What's New in vSphere 6.7 Core Storage](#)
- [What's New in vSphere 6.5 Storage](#)

---

### vVols and VASA on vCenter 8.0

vCenter Server 8.0 Update 2b | 29 FEB 2024 | GA ISO Build 23319993

#### **In a mixed vSphere 7.x and 8.x environment with vSphere DRS, an incompatible virtual machine might prevent an ESXi host to enter maintenance mode**

- In a mixed vSphere 7.x and 8.x environment with DRS, a VM of 7.x version that is not compatible with ESXi 8.0 Update 1 and later might prevent an ESXi 8.0 Update 1 host to enter maintenance mode. The issue is specific for virtual machines with VMDK on a vSphere Virtual Volumes datastore. In the vSphere Client, you see an error such as **Waiting for all VMs to be powered off or suspended or migrated. In a DRS cluster check the Faults page on the DRS tab for troubleshooting.**
- Workaround: Power-off the incompatible virtual machine.

#### **VASA API version does not automatically refresh after upgrade to vCenter Server 8.0**

- vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:
- If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the upgrade, when you navigate to vCenter Server > Configure > Storage Providers and expand the General tab of the registered VASA provider, you still see VASA API version 3.5.
- If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.
- Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

---

### vCenter Server 8.0a | 15 DEC 2022 | ISO Build 20920323

#### **VASA API version does not automatically refresh after upgrade to vCenter Server 8.0**

- vCenter Server 8.0 supports VASA API version 4.0. However, after you upgrade your vCenter Server system to version 8.0, the VASA API version might not automatically change to 4.0. You see the issue in 2 cases:
- If a VASA provider that supports VASA API version 4.0 is registered with a previous version of VMware vCenter, the VASA API version remains unchanged after you upgrade to VMware vCenter 8.0. For example, if you upgrade a VMware vCenter system of version 7.x with a registered VASA provider that supports both VASA API versions 3.5 and 4.0, the VASA API version does not automatically change to 4.0, even though the VASA provider supports VASA API version 4.0. After the

upgrade, when you navigate to vCenter Server > Configure > Storage Providers and expand the General tab of the registered VASA provider, you still see VASA API version 3.5.

- If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.
- Workaround: Unregister and re-register the VASA provider on the VMware vCenter 8.0 system.

### vVols and VASA on vCenter 7.0

vCenter Server 7.0 Update 3l | 30 MAR 2023 | ISO Build 21477706

#### You cannot add Raw Device Mapping (RDM) as an existing disk to a virtual machine

- If the VMX file of a VM resides on an NFS or vSphere Virtual Volumes datastore, when you create or edit that VM, you cannot add an RDM as an existing disk. In the vSphere Client, you see a validation error for the newly added disk. If you change the location for the new disk to a VMFS datastore, the validation error disappears, but the operation fails with error such as: Invalid virtual machine configuration. Storage policy change failure: @&!\*@\*(mgs.disklin.INVALIDDISK) The file specified is not a virtual disk.
- This issue is resolved in this release.

---

### vCenter Server 7.0 Update 3d - Build 19480866

- **vCenter Server and vSphere Client Issues**
- **You cannot refresh storage provider certificates from the vSphere Client**

In the vSphere Client, when you navigate to **Configure > Storage Providers**, you see the option **Refresh certificate** dimmed.

This issue is resolved in this release.

---

### VMware vCenter Server 7.0 Update 3d - Build 19480866

- **You cannot refresh storage provider certificates from the vSphere Client**

In the vSphere Client, when you navigate to **Configure > Storage Providers**, you see the option **Refresh certificate** dimmed.

This issue is resolved in this release.

---

### VMware vCenter Server 7.0.0b - Build 16386292

- **vCenter Server does not report compatible datastores when datastores are mounted on multiple data centers**

If you have ESXi hosts across different data centers that use the same VMware vSphere Virtual Volumes datastore, you do not see compatible datastores when you run the `Get-SpbmCompatibleStorage` PowerCLI command or use the storage policy

editor in the vSphere Client.

This issue is resolved in this release.

---

---

## vVols and VASA on vCenter 6.7

### VMware vCenter Server 6.7 Update 3m - Build 17713310

- **After an upgrade to vCenter Server 6.7 Update 3l and later, you might see existing or new HPE VASA providers in disconnected state**

If the inventory of your vCenter Server system has vSphere Virtual Volumes supported by either of HPE 3PAR StoreServ or HPE Primera VASA providers, you might see the providers get into a disconnected state after an upgrade to vCenter Server 6.7 Update 3l or later. The issue affects 3PAR 3.3.1 MU5 storage, but not 3PAR 3.3.1 MU3 storage.

Workaround: Upgrade to vCenter Server 7.0 Update 1c and later. For upgrade compatibility, see VMware knowledge base article [67077](#).

Alternatively, you can restore your system to a backup prior to vCenter Server 6.7 Update 3l.

If you are not already using HPE 3PAR 3.3.1 MU5 VASA provider, postpone the VASA provider upgrade to HPE 3PAR 3.3.1 MU5 until HPE resolves the issue. For more information, see VMware knowledge base article [83038](#).

---

### VMware vCenter Server 6.7 Update 3l - Build 17138064

- **NEW: After an upgrade to vCenter Server 6.7 Update 3l and later, you might see existing or new HPE VASA providers in disconnected state**

If the inventory of your vCenter Server system has vSphere Virtual Volumes supported by either of HPE 3PAR StoreServ or HPE Primera VASA providers, you might see the providers get into a disconnected state after an upgrade to vCenter Server 6.7 Update 3l or later. The issue affects 3PAR 3.3.1 MU5 storage, but not 3PAR 3.3.1 MU3 storage.

Workaround: Upgrade to vCenter Server 7.0 Update 1c and later. For upgrade compatibility, see VMware knowledge base article [67077](#). Alternatively, you can restore your system to a backup prior to vCenter Server 6.7 Update 3l.

If you are not already using HPE 3PAR 3.3.1 MU5 VASA provider, postpone the VASA provider upgrade to HPE 3PAR 3.3.1 MU5 until HPE resolves the issue. For more information, see VMware knowledge base article [83038](#).

---

### VMware vCenter Server 6.7 Update 3j - Build 16708996

- **vCenter Server does not report compatible datastores when one datastore is mounted on multiple datacenters**

If you have ESXi hosts across different datacenters that use the same VMware vSphere Virtual Volumes datastore, you do not see compatible datastores when you run the `Get-SpbmCompatibleStorage` PowerCLI command or use the storage policy editor in the vSphere Client.

This issue is resolved in this release.

### VMware vCenter Server 6.7 Update 2 - Build 13010631

- **Posting of VMware vSphere Virtual Volumes compliance alarms for a StorageObject type to a vCenter Server system might fail**

If you use an API for Storage Awareness (VASA) provider, posting of vSphere Virtual Volumes compliance alarms for a StorageObject type to a vCenter Server system might fail due to a mapping mismatch.

This issue is resolved in this release.

---

### vVols and VASA on vCenter 6.5

#### VMware vCenter Server 6.5 Update 3k - Build 16613358

- **vCenter Server does not report compatible datastores when datastores are mounted on multiple data centers**

If you have ESXi hosts across different data centers that use the same VMware vSphere Virtual Volumes datastore, you do not see compatible datastores when you run the `Get-SpbmCompatibleStorage` PowerCLI command or use the storage policy editor in the vSphere Client.

This issue is resolved in this release.

---

#### vCenter Server 6.5.0e - Build 5705665

- **After upgrade from vSphere 6.0 to vSphere 6.5, the Virtual Volumes storage policy might disappear from the VM Storage Policies list**

After you upgrade your environment to vSphere 6.5, the Virtual Volumes storage policy that you created in vSphere 6.0 might no longer be visible in the list of VM storage policies.

Workaround: Log out of the vSphere Web Client, and then log in again.

- **The vSphere Web Client fails to display information about the default profile of a Virtual Volumes datastore**

Typically, you can check information about the default profile associated with the Virtual Volumes datastore. In the vSphere Web Client, you do it by browsing to the datastore, and then clicking **Configure > Settings > Default Profiles**.

However, the vSphere Web Client is unable to report the default profiles when their IDs, configured at the storage side, are not unique across all the datastores reported by the same Virtual Volumes provider.

Workaround: None.

---

#### vCenter Server 6.5.0b - Build 5178943

- **Virtual Volumes (VVOL) replication groups in INTEST or FAILEDOVER replication state cannot be selected in the vSphere Web Client**

After a failover that is based on VVol replication, the newly created replication groups at the target site which are in INTEST or FAILEDOVER replication state, cannot be viewed or selected in the vSphere Web Client.

This issue is resolved in this release.

---

### VMware vSphere 6.5 - Build 4564106

- **After upgrade from vSphere 6.0 to vSphere 6.5, the Virtual Volumes storage policy might disappear from the VM Storage Policies list**

After you upgrade your environment to vSphere 6.5, the Virtual Volumes storage policy that you created in vSphere 6.0 might no longer be visible in the list of VM storage policies.

Workaround: Log out of the vSphere Web Client, and then log in again.

- **The vSphere Web Client fails to display information about the default profile of a Virtual Volumes datastore**

Typically, you can check information about the default profile associated with the Virtual Volumes datastore. In the vSphere Web Client, you do it by browsing to the datastore, and then clicking **Configure > Settings > Default Profiles**.

However, the vSphere Web Client is unable to report the default profiles when their IDs, configured at the storage side, are not unique across all the datastores reported by the same Virtual Volumes provider.

Workaround: None.

---

---

