VMware ESXi 8.0 Update 1a Release Notes

VMware vSphere 8.0
ESXi 8.0 Update 1

vmware®

You can find the most up-to-date technical documentation on the VMware website at:

https://docs.vmware.com/

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

Introduction

1

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

Check for additions and updates to these release notes.

What's New

2

ESXi 8.0 Update 1a is applicable only for vSAN enabled ESXi hosts, configured with deduplication and compression. For more details, see the Chapter 5 Resolved Issues section.

Earlier Releases of ESXi 8.0

3

New features, resolved, and known issues of ESXi are described in the release notes for each release. Release notes for earlier releases of ESXi 8.0 are:

- VMware ESXi 8.0 Update 1 Release Notes
- VMware ESXi 8.0c Release Notes
- VMware ESXi 8.0b Release Notes
- VMware ESXi 8.0a Release Notes

For internationalization, compatibility, and open source components, see the VMware vSphere 8.0 Release Notes.

Patches Contained in This Release 4

Read the following topics next:

- Patch for VMware ESXi 8.0 Update 1a
- Patch Download and Installation

Patch for VMware ESXi 8.0 Update 1a

This release of ESXi 8.0 Update 1a delivers the following patches:

Build Details

Download Filename:	VMware-ESXi-8.0U1a-21813344-depot.zip
Build:	21813344
Download Size:	583.2 MB
sha256checksum:	52dd56f65367afdc0503419370cee7bf62965940cd9cca3af6b9c2409e615950
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.10.21813344	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.10.21813344	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.10.21813344	Bugfix	Critical

Rollup Bulletin

This rollup bulletin contains the latest VIBs with all the fixes after the initial release of ESXi 8.0.

Bulletin ID	Category	Severity
ESXi80U1a-21813344	Bugfix	Critical

Image Profiles

VMware patch and update releases contain general and critical image profiles. Application of the general release image profile applies to new bug fixes.

Image Profile Name
ESXi-8.0U1a-21813344-standard
ESXi-8.0U1a-21813344-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi-8.0U1a-21813344	06/01/2023	Bugfix	Bugfix image

For information about the individual components and bulletins, see the Product Patches page and the Chapter 5 Resolved Issues section.

Patch Download and Installation

For details on updates and upgrades by using vSphere Lifecycle Manager, see About vSphere Lifecycle Manager and vSphere Lifecycle Manager Baselines and Images. You can also update ESXi hosts without the use of vSphere Lifecycle Manager by using an image profile. To do this, you must manually download the patch offline bundle ZIP file from VMware Customer Connect. From the **Select a Product** drop-down menu, select **ESXi (Embedded and Installable)** and from the **Select a Version** drop-down menu, select **8.0**. For more information, see the Upgrading Hosts by Using ESXCLI Commands and the VMware ESXi Upgrade guide.

Resolved Issues

5

Read the following topics next:

- ESXi_8.0.1-0.10.21813344
- esx-update_8.0.1-0.10.21813344
- esxio-update_8.0.1-0.10.21813344
- ESXi-8.0U1a-21813344-standard
- ESXi-8.0U1a-21813344-no-tools
- ESXi-8.0U1a-21813344

ESXi_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_trx_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal_8.0.1-0.10.21813344 ■ VMware_bootbank_vsanhealth_8.0.1-0.10.21813344 ■ VMware_bootbank_clusterstore_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 ■ VMware_bootbank_crx_8.0.1-0.10.21813344 ■ VMware_bootbank_gc_8.0.1-0.10.21813344 ■ VMware_bootbank_vdfs_8.0.1-0.10.21813344 ■ VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-base_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-base_8.0.1-0.10.21813344 ■ VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 ■ VMware_bootbank_vsan_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 ■ VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio_8.0.1-0.10.21813344
PRs Fixed	3217456
CVE numbers	N/A

This patch updates the `vsan` VIB. Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `trx`, `bmcal-esxio`, `bmcal`, `vsanhealth`, `clusterstore`, `esx-dvfilter-generic-fastpath`, `native-misc-drivers`, `esxio-dvfilter-generic-fastpath`, `esxio-combiner`, `crx`, `gc`, `vdfs`, `gc-esxio`, `esx-base`, `esxio-combiner-esxio`, `native-misc-drivers-esxio`, `esxio-base`, `vds-vsip`, `esx-xserver`, `cpu-microcode`, and `esxio`.

- **vSAN hosts with deduplication and compression might fail after upgrade to ESXi 8.0 Update 1**

  This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:

  a You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.

b You have added new capacity disks to such a disk group.

c You upgrade the vSAN host to ESXi 8.0 Update 1.

The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

```
Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world
2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108
Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/
dedup/dedup_util.h:231 -- VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx-
>devices[devIndex].sb, blkType, blkNumCheck, 1))
```

For more information, see VMware knowledge base article 92458.

This issue is resolved in this release.

esx-update_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	■ VMware_bootbank_esx-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesx_8.0.1-0.10.21813344
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esx-update` and `loadesx`.

esxio-update_8.0.1-0.10.21813344

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	■ VMware_bootbank_esxio-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	N/A
CVE numbers	N/A

Due to their dependency with the `esx-base` VIB, the following VIBs are updated with build number and patch version changes, but deliver no fixes: `esxio-update` and `loadesxio`.

ESXi-8.0U1a-21813344-standard

Profile Name	ESXi-8.0U1a-21813344-standard
Build	For build information, see Chapter 4 Patches Contained in This Release.
Vendor	VMware, Inc.
Release Date	June 1, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_trx_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal_8.0.1-0.10.21813344 ■ VMware_bootbank_vsanhealth_8.0.1-0.10.21813344 ■ VMware_bootbank_clusterstore_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 ■ VMware_bootbank_crx_8.0.1-0.10.21813344 ■ VMware_bootbank_gc_8.0.1-0.10.21813344 ■ VMware_bootbank_vdfs_8.0.1-0.10.21813344 ■ VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-base_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-base_8.0.1-0.10.21813344 ■ VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 ■ VMware_bootbank_vsan_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 ■ VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesx_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

- This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:
  a. You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.
  b. You have added new capacity disks to such a disk group.
  c. You upgrade the vSAN host to ESXi 8.0 Update 1.

  The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

```
Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world

2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108

Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/
dedup/dedup_util.h:231 -- VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx-
>devices[devIndex].sb, blkType, blkNumCheck, 1))
```

For more information, see VMware knowledge base article 92458.

ESXi-8.0U1a-21813344-no-tools

Profile Name	ESXi-8.0U1a-21813344-no-tools
Build	For build information, see Chapter 4 Patches Contained in This Release.
Vendor	VMware, Inc.
Release Date	June 1, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_trx_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_bmcal_8.0.1-0.10.21813344 ■ VMware_bootbank_vsanhealth_8.0.1-0.10.21813344 ■ VMware_bootbank_clusterstore_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner_8.0.1-0.10.21813344 ■ VMware_bootbank_crx_8.0.1-0.10.21813344 ■ VMware_bootbank_gc_8.0.1-0.10.21813344 ■ VMware_bootbank_vdfs_8.0.1-0.10.21813344 ■ VMware_bootbank_gc-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-base_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-base_8.0.1-0.10.21813344 ■ VMware_bootbank_vds-vsip_8.0.1-0.10.21813344 ■ VMware_bootbank_vsan_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-xserver_8.0.1-0.10.21813344 ■ VMware_bootbank_cpu-microcode_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio_8.0.1-0.10.21813344 ■ VMware_bootbank_esx-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesx_8.0.1-0.10.21813344 ■ VMware_bootbank_esxio-update_8.0.1-0.10.21813344 ■ VMware_bootbank_loadesxio_8.0.1-0.10.21813344
PRs Fixed	3154255
Related CVE numbers	N/A

This patch updates the following issue:

- This issue occurs only on vSAN all-flash hosts with deduplication and compression enabled and only when:
  a. You have an existing disk group created on a vSAN version earlier than vSAN 7.0 Update 3.
  b. You have added new capacity disks to such a disk group.
  c. You upgrade the vSAN host to ESXi 8.0 Update 1.

  The issue does not impact vSAN hosts with compression-only or checksum-only disk groups.

ESXi 8.0 Update 1 includes a background scanner to collect metadata. This background scanner might not recognize new disks added to an existing disk group. This issue can cause the ESXi host to fail with a purple diagnostic screen. Because the host fails and reboots before the background scan completes, repeated failures might occur.

The following details appear in the backtrace:

```
Panic Details: Crash at 2023-05-01T11:31:12.957Z on CPU 51 running world

2019239 - VSAN_0x450080017780_PLOG. VMK Uptime:0:09:54:11.108

Panic Message: @BlueScreen: Failed at bora/modules/vmkernel/plog/
dedup/dedup_util.h:231 -- VMK_ASSERT(blkNum < DDPGetEndBlkNumber(&ddCtx-
>devices[devIndex].sb, blkType, blkNumCheck, 1))
```

For more information, see VMware knowledge base article 92458.

ESXi-8.0U1a-21813344

Name	ESXi
Version	ESXi-8.0U1a-21813344
Release Date	June 01, 2023
Category	Bugfix
Affected Components	■ ESXi Component - core ESXi VIBs ■ ESXi Install/Upgrade Component
PRs Fixed	3217456
Related CVE numbers	N/A

Known Issues from Previous Releases

6

Read the following topics next:

Installation, Upgrade, and Migration Issues

- **If you apply a host profile using a software FCoE configuration to an ESXi 8.0 host, the operation fails with a validation error**

  Starting from vSphere 7.0, software FCoE is deprecated, and in vSphere 8.0 software FCoE profiles are not supported. If you try to apply a host profile from an earlier version to an ESXi 8.0 host, for example to edit the host customization, the operation fails. In the vSphere Client, you see an error such as `Host Customizations validation error`.

  Workaround: Disable the Software FCoE Configuration subprofile in the host profile.

- **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.

- **You cannot use ESXi hosts of version 8.0 as a reference host for existing host profiles of earlier ESXi versions**

  Validation of existing host profiles for ESXi versions 7.x, 6.7.x and 6.5.x fails when only an 8.0 reference host is available in the inventory.

  Workaround: Make sure you have a reference host of the respective version in the inventory. For example, use an ESXi 7.0 Update 2 reference host to update or edit an ESXi 7.0 Update 2 host profile.

- **VMNICs might be down after an upgrade to ESXi 8.0**

  If the peer physical switch of a VMNIC does not support Media Auto Detect, or Media Auto Detect is disabled, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

  Workaround: Use either of these 2 options:

  a Enable the option `media-auto-detect` in the BIOS settings by navigating to System Setup Main Menu, usually by pressing **F2** or opening a virtual console, and then **Device Settings** > *<specific broadcom NIC>* > **Device Configuration Menu** > **Media Auto Detect**. Reboot the host.

  b Alternatively, use an ESXCLI command similar to: `esxcli network nic set -S <your speed> -D full -n <your nic>`. With this option, you also set a fixed speed to the link, and it does not require a reboot.

- **If a vCenter Server Security Token Service (STS) refresh happens during upgrade to ESXi 8.0, the upgrade might fail**

  In vSphere 8.0, vCenter Single Sign-On automatically renews a VMCA-generated STS signing certificate. The auto-renewal occurs before the STS signing certificate expires and before triggering the 90-day expiration alarm. However, in long-running upgrade or remediation tasks by using a vSphere Lifecycle Manager image on multiple ESXi hosts in a cluster, vSphere Lifecycle Manager might create a cache of STS certificates internally. In very rare cases, if an STS certificates refresh task starts in parallel with the long-running upgrade or remediation task, the upgrade task might fail as the STS certificates in the internal cache might be different from the refreshed certificates. After the upgrade task fails, some ESXi hosts might remain in maintenance mode.

  Workaround: Manually exit any ESXi hosts in maintenance mode and retry the upgrade or remediation. Refreshing or importing and replacing the STS signing certificates happens automatically and does not require a vCenter Server restart, to avoid downtime.

- **After upgrade to ESXi 8.0, you might lose some nmlx5_core driver module settings due to obsolete parameters**

Some module parameters for the `nmlx5_core` driver, such as `device_rss`, `drss` and `rss`, are deprecated in ESXi 8.0 and any custom values, different from the default values, are not kept after an upgrade to ESXi 8.0.

Workaround: Replace the values of the `device_rss`, `drss` and `rss` parameters as follows:

- `device_rss`: Use the `DRSS` parameter.
- `drss`: Use the `DRSS` parameter.
- `rss`: Use the `RSS` parameter.

- **Second stage of vCenter Server restore procedure freezes at 90%**

  When you use the vCenter Server GUI installer or the vCenter Server Appliance Management Interface (VAMI) to restore a vCenter from a file-based backup, the restore workflow might freeze at 90% with an error `401 Unable to authenticate user`, even though the task completes successfully in the backend. The issue occurs if the deployed machine has a different time than the NTP server, which requires a time sync. As a result of the time sync, clock skew might fail the running session of the GUI or VAMI.

  Workaround: If you use the GUI installer, you can get the restore status by using the `restore.job.get` command from the `appliancesh` shell. If you use VAMI, refresh your browser.

Miscellaneous Issues

- **If a PCI passthrough is active on a DPU during the shutdown or restart of an ESXi host, the host fails with a purple diagnostic screen**

  If an active virtual machine has a PCI passthrough to a DPU at the time of shutdown or reboot of an ESXi host, the host fails with a purple diagnostic screen. The issue is specific for systems with DPUs and only in case of VMs that use PCI passthrough to the DPU.

  Workaround: Before shutdown or reboot of an ESXi host, make sure the host is in maintenance mode, or that no VMs that use PCI passthrough to a DPU are running. If you use auto start options for a virtual machine, the Autostart manager stops such VMs before shutdown or reboot of a host.

- **RDMA over Converged Ethernet (RoCE) traffic might fail in Enhanced Networking Stack (ENS) and VLAN environment, and a Broadcom RDMA network interface controller (RNIC)**

  The VMware solution for high bandwidth, ENS, does not support MAC VLAN filters. However, a RDMA application that runs on a Broadcom RNIC in an ENS + VLAN environment, requires a MAC VLAN filter. As a result, you might see some RoCE traffic disconnected. The issue is likely to occur in a NVMe over RDMA + ENS + VLAN environment, or in an ENS+VLAN+RDMA app environment, when an ESXi host reboots or an uplink goes up and down.

  Workaround: None

- **Reset or restore of the ESXi system configuration in a vSphere system with DPUs might cause invalid state of the DPUs**

  If you reset or restore the ESXi system configuration in a vSphere system with DPUs, for example, by selecting **Reset System Configuration** in the direct console, the operation might cause invalid state of the DPUs. In the DCUI, you might see errors such as `Failed to reset system configuration. Note that this operation cannot be performed when a managed DPU is present`. A backend call to the `-f` force reboot option is not supported for ESXi installations with a DPU. Although ESXi 8.0 supports the `-f` force reboot option, if you use `reboot -f` on an ESXi configuration with a DPU, the forceful reboot might cause an invalid state.

  Workaround: Reset System Configuration in the direct console interface is temporarily disabled. Avoid resetting the ESXi system configuration in a vSphere system with DPUs.

- **You cannot mount an IPv6-based NFS 3 datastore with VMkernel port binding by using ESXCLI commands**

  When you try to mount an NFS 3 datastore with an IPv6 server address and VMkernel port binding by using an ESXCLI command, the task fails with an error such as:

  ```
  [:~] esxcli storage nfs add -I fc00:xxx:xxx:xx::xxx:vmk1 -s share1 -v volume1

  Validation of vmknic failed Instance(defaultTcpipStack, xxx:xxx:xx::xxx:vmk1)
  Input(): Not found:
  ```

  The issue is specific for NFS 3 datastores with an IPv6 server address and VMkernel port binding.

  Workaround: Use the vSphere Client as an alternative to mount IPv6-based NFSv3 datastores with VMkernel port binding.

- **In a vCenter Server system with DPUs, if IPv6 is disabled, you cannot manage DPUs**

  Although the vSphere Client allows the operation, if you disable IPv6 on an ESXi host with DPUs, you cannot use the DPUs, because the internal communication between the host and the devices depends on IPv6. The issue affects only ESXi hosts with DPUs.

  Workaround: Make sure IPv6 is enabled on ESXi hosts with DPUs.

- **You might see 10 min delay in rebooting an ESXi host on HPE server with pre-installed Pensando DPU**

  In rare cases, HPE servers with pre-installed Pensando DPU might take more than 10 minutes to reboot in case of a failure of the DPU. As a result, ESXi hosts might fail with a purple diagnostic screen and the default wait time is 10 minutes.

  Workaround: None.

- **If you have an USB interface enabled in a remote management application that you use to install ESXi 8.0, you see an additional standard switch vSwitchBMC with uplink vusb0**

Starting with vSphere 8.0, in both Integrated Dell Remote Access Controller (iDRAC) and HP Integrated Lights Out (ILO), when you have an USB interface enabled, vUSB or vNIC respectively, an additional standard switch `vSwitchBMC` with uplink `vusb0` gets created on the ESXi host. This is expected, in view of the introduction of data processing units (DPUs) on some servers but might cause the VMware Cloud Foundation Bring-Up process to fail.

Workaround: Before vSphere 8.0 installation, disable the USB interface in the remote management application that you use by following vendor documentation.

After vSphere 8.0 installation, use the ESXCLI command `esxcfg-advcfg -s 0 /Net/BMCNetworkEnable` to prevent the creation of a virtual switch `vSwitchBMC` and associated portgroups on the next reboot of host.

See this script as an example:

```
~# esxcfg-advcfg -s 0 /Net/BMCNetworkEnable
```

The value of BMCNetworkEnable is 0 and the service is disabled.

```
~# reboot
```

On host reboot, no virtual switch, PortGroup and VMKNIC are created in the host related to remote management application network.

- **If an NVIDIA BlueField DPU is in hardware offload mode disabled, virtual machines with configured SR-IOV virtual function cannot power on**

  NVIDIA BlueField DPUs must be in hardware offload mode enabled to allow virtual machines with configured SR-IOV virtual function to power on and operate.

  Workaround: Always use the default hardware offload mode enabled for NVIDIA BlueField DPUs when you have VMs with configured SR-IOV virtual function connected to a virtual switch.

- **In the Virtual Appliance Management Interface (VAMI), you see a warning message during the pre-upgrade stage**

  Moving vSphere plug-ins to a remote plug-in architecture, vSphere 8.0 deprecates support for local plug-ins. If your 8.0 vSphere environment has local plug-ins, some breaking changes for such plug-ins might cause the pre-upgrade check by using VAMI to fail.

  In the Pre-Update Check Results screen, you see an error such as:

  ```
  Warning message: The compatibility of plug-in package(s) %s with the new vCenter Server version cannot be validated. They may not function properly after vCenter Server upgrade.

  Resolution: Please contact the plug-in vendor and make sure the package is compatible with the new vCenter Server version.
  ```

Workaround: Refer to the VMware Compatibility Guide and VMware Product Interoperability Matrix or contact the plug-in vendors for recommendations to make sure local plug-ins in your environment are compatible with vCenter Server 8.0 before you continue with the upgrade. For more information, see the blog Deprecating the Local Plugins :- The Next Step in vSphere Client Extensibility Evolution and VMware knowledge base article 87880.

- **You cannot remove a PCI passthrough device assigned to a virtual Non-Uniform Memory Access (NUMA) node from a virtual machine with CPU Hot Add enabled**

  Although by default when you enable CPU Hot Add to allow the addition of vCPUs to a running virtual machine, virtual NUMA topology is deactivated, if you have a PCI passthrough device assigned to a NUMA node, attempts to remove the device end with an error. In the vSphere Client, you see messages such as `Invalid virtual machine configuration. Virtual NUMA cannot be configured when CPU hotadd is enabled`.

  Workaround: See VMware knowledge base article 89638.

- **If you deploy a virtual machine from an OVF file or from the Content Library, the number of cores per socket for the VM is set to 1**

  If you deploy a virtual machine from an OVF file or from the Content Library, instead of ESXi automatically selecting the number of cores per socket, the number is pre-set to 1.

  Workaround: You can manually set the number of cores per socket by using the vSphere Client. For more information, see VMware knowledge base article 89639.

- **If you configure a VM at HW version earlier than 20 with a Vendor Device Group, such VMs might not work as expected**

  Vendor Device Groups, which enable binding of high-speed networking devices and the GPU, are supported only on VMs with HW version 20 and later, but you are not prevented to configure a VM at HW version earlier than 20 with a Vendor Device Group. Such VMs might not work as expected: for example, fail to power-on.

  Workaround: Ensure that VM HW version is of version 20 before you configure a Vendor Device Group in that VM.

Networking Issues

- **You cannot set the Maximum Transmission Unit (MTU) on a VMware vSphere Distributed Switch to a value larger than 9174 on a Pensando DPU**

  If you have the vSphere Distributed Services Engine feature with a Pensando DPU enabled on your ESXi 8.0 system, you cannot set the Maximum Transmission Unit (MTU) on a vSphere Distributed Switch to a value larger than 9174.

  Workaround: None.

- **ESXi reboot takes long due to NFS server mount timeout**

When you have multiple mounts on an NFS server that is not accessible, ESXi retries connection to each mount for 30 seconds, which might add up to minutes of ESXi reboot delay, depending on the number of mounts.

Workaround: ESXi Update 8.0 Update 1 adds a configurable option to override the default mount timeout: `esxcfg-advcfg -s <timeout val> /NFS/MountTimeout`. For example, if you want to reconfigure mount timeout to 10 seconds, you can run the following command: `- esxcfg-advcfg -s 10 /NFS/MountTimeout`. Use the command `esxcfg-advcfg -g /NFS/MountTimeout` to verify the current configured mount timeout.

- **You see link flapping on NICs that use the ntg3 driver of version 4.1.3 and later**

  When two NICs that use the `ntg3` driver of versions 4.1.3 and later are connected directly, not to a physical switch port, link flapping might occur. The issue does not occur on `ntg3` drivers of versions earlier than 4.1.3 or the `tg3` driver. This issue is not related to the occasional Energy Efficient Ethernet (EEE) link flapping on such NICs. The fix for the EEE issue is to use a `ntg3` driver of version 4.1.7 or later, or disable EEE on physical switch ports.

  Workaround: Upgrade the `ntg3` driver to version 4.1.8 and set the new module parameter `noPhyStateSet` to `1`. The `noPhyStateSet` parameter defaults to `0` and is not required in most environments, except they face the issue.

- **VMware NSX installation or upgrade in a vSphere environment with DPUs might fail with a connectivity error**

  An intermittent timing issue on the ESXi host side might cause NSX installation or upgrade in a vSphere environment with DPUs to fail. In the `nsxapi.log` file you see logs such as `Failed to get SFHC response. MessageType MT_SOFTWARE_STATUS`.

  Workaround: Wait for 10 min and retry the NSX install or upgrade.

- **If you do not reboot an ESXi host after you enable or disable SR-IOV with the icen driver, when you configure a transport node in ENS Interrupt mode on that host, some virtual machines might not get DHCP addresses**

  If you enable or disable SR-IOV with the `icen` driver on an ESXi host and configure a transport node in ENS Interrupt mode, some Rx (receive) queues might not work if you do not reboot the host. As a result, some virtual machines might not get DHCP addresses.

  Workaround: Either add a transport node profile directly, without enabling SR-IOV, or reboot the ESXi host after you enable or disable SR-IOV.

- **You cannot use Mellanox ConnectX-5, ConnectX-6 cards Model 1 Level 2 and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0**

  Due to hardware limitations, Model 1 Level 2, and Model 2 for Enhanced Network Stack (ENS) mode in vSphere 8.0 is not supported in ConnectX-5 and ConnectX-6 adapter cards.

  Workaround: Use Mellanox ConnectX-6 Lx and ConnectX-6 Dx or later cards that support ENS Model 1 Level 2, and Model 2A.

- **Pensando DPUs do not support Link Layer Discovery Protocol (LLDP) on physical switch ports of ESXi hosts**

  When you enable LLDP on an ESXi host with a DPU, the host cannot receive LLDP packets.

  Workaround: None.

Storage Issues

- **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:

  a 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.

  b 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.

- **vSphere Storage vMotion operations might fail in a vSAN environment due to an unauthenticated session of the Network File Copy (NFC) manager**

  Migrations to a vSAN datastore by using vSphere Storage vMotion of virtual machines that have at least one snapshot and more than one virtual disk with different storage policy might fail. The issue occurs due to an unauthenticated session of the NFC manager because the Simple Object Access Protocol (SOAP) body exceeds the allowed size.

  Workaround: First migrate the VM home namespace and just one of the virtual disks. After the operation completes, perform a disk only migration of the remaining 2 disks.

- **You cannot create snapshots of virtual machines due to an error in the Content Based Read Cache (CBRC) that a digest operation has failed**

  A rare race condition when assigning a content ID during the update of the CBRC digest file might cause a discrepancy between the content ID in the data disk and the digest disk. As a result, you cannot create virtual machine snapshots. You see an error such as `An error occurred while saving the snapshot: A digest operation has failed` in the backtrace. The snapshot creation task completes upon retry.

Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

- **The Utilization view of resource pools and clusters might not automatically refresh when you change the object**

  When you have already opened the **Utilization** view under the **Monitor** tab for a resource pool or a cluster and then you change the resource pool or cluster, the view might not automatically refresh. For example, when you open the **Utilization** view of one cluster and then select a different cluster, you might still see the statistics of the first cluster.

  Workaround: Click the refresh icon.

- **If you load the vSphere virtual infrastructure to more than 90%, ESXi hosts might intermittently disconnect from vCenter Server**

  In rare occasions, if the vSphere virtual infrastructure is continuously using more than 90% of its hardware capacity, some ESXi hosts might intermittently disconnect from the vCenter Server. Connection typically restores within a few seconds.

  Workaround: If connection to vCenter Server accidentally does not restore in a few seconds, reconnect ESXi hosts manually by using vSphere Client.

- **In the vSphere Client, you do not see banner notifications for historical data imports**

  Due to a backend issue, you do not see banner notifications for background migration of historical data in the vSphere Client.

  Workaround: Use the vCenter Server Management Interface as an alternative to the vSphere Client. For more information, see Monitor and Manage Historical Data Migration.

- **You see an error for Cloud Native Storage (CNS) block volumes created by using API in a mixed vCenter environment**

  If your environment has vCenter Server systems of version 8.0 and 7.x, creating Cloud Native Storage (CNS) block volume by using API is successful, but you might see an error in the vSphere Client, when you navigate to see the CNS volume details. You see an error such as `Failed to extract the requested data. Check vSphere Client logs for details. + TypeError: Cannot read properties of null (reading 'cluster')`. The issue occurs only if you review volumes managed by the 7.x vCenter Server by using the vSphere Client of an 8.0 vCenter Server.

  Workaround: Log in to vSphere Client on a vCenter Server system of version 7.x to review the volume properties.

- **ESXi hosts might become unresponsive, and you see a vpxa dump file due to a rare condition of insufficient file descriptors for the request queue on vpxa**

In rare cases, when requests to the vpxa service take long, for example waiting for access to a slow datastore, the request queue on vpxa might exceed the limit of file descriptors. As a result, ESXi hosts might briefly become unresponsive, and you see a `vpxa-zdump.00*` file in the `/var/core` directory. The vpxa logs contain the line `Too many open files`.

Workaround: None. The vpxa service automatically restarts and corrects the issue.

- **If you use custom update repository with untrusted certificates, vCenter Server upgrade or update by using vCenter Lifecycle Manager workflows to vSphere 8.0 might fail**

  If you use a custom update repository with self-signed certificates that the VMware Certificate Authority (VMCA) does not trust, vCenter Lifecycle Manager fails to download files from such a repository. As a result, vCenter Server upgrade or update operations by using vCenter Lifecycle Manager workflows fail with the error `Failed to load the repository manifest data for the configured upgrade`.

  Workaround: Use CLI, the GUI installer, or the Virtual Appliance Management Interface (VAMI) to perform the upgrade. For more information, see VMware knowledge base article 89493.

Virtual Machine Management Issues

- **When you add an existing virtual hard disk to a new virtual machine, you might see an error that the VM configuration is rejected**

  When you add an existing virtual hard disk to a new virtual machine by using the VMware Host Client, the operation might fail with an error such as `The VM configuration was rejected. Please see browser Console`. The issue occurs because the VMware Host Client might fail to get some properties, such as the hard disk controller.

  Workaround: After you select a hard disk and go to the **Ready to complete** page, do not click **Finish**. Instead, return one step back, wait for the page to load, and then click **Next** > **Finish**.

vSphere Lifecycle Manager Issues

- **If you use an ESXi host deployed from a host profile with enabled stateful install as an image to deploy other ESXi hosts in a cluster, the operation fails**

  If you extract an image of an ESXi host deployed from a host profile with enabled stateful install to deploy other ESXi hosts in a vSphere Lifecycle Manager cluster, the operation fails. In the vSphere Client, you see an error such as `A general system error occurred: Failed to extract image from the host: no stored copy available for inactive VIB VMW_bootbank_xxx. Extraction of image from host xxx.eng.vmware.com failed`.

  Workaround: Use a different host from the cluster to extract an image.

- **You see error messages when try to stage vSphere Lifecycle Manager Images on ESXi hosts of version earlier than 8.0**

ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

Workaround: None. After you confirm that the affected ESXi hosts are of version earlier than 8.0, ignore the errors.

- **A remediation task by using vSphere Lifecycle Manager might intermittently fail on ESXi hosts with DPUs**

  When you start a vSphere Lifecycle Manager remediation on an ESXi hosts with DPUs, the host upgrades and reboots as expected, but after the reboot, before completing the remediation task, you might see an error such as:

  `A general system error occurred: After host … remediation completed, compliance check reported host as 'non-compliant'. The image on the host does not match the image set for the cluster. Retry the cluster remediation operation.`

  This is a rare issue, caused by an intermittent timeout of the post-remediation scan on the DPU.

  Workaround: Reboot the ESXi host and re-run the vSphere Lifecycle Manager compliance check operation, which includes the post-remediation scan.

VMware Host Client Issues

- **VMware Host Client might display incorrect descriptions for severity event states**

  When you look in the VMware Host Client to see the descriptions of the severity event states of an ESXi host, they might differ from the descriptions you see by using Intelligent Platform Management Interface (IPMI) or Lenovo XClarity Controller (XCC). For example, in the VMware Host Client, the description of the severity event state for the PSU Sensors might be `Transition to Non-critical from OK`, while in the XCC and IPMI, the description is `Transition to OK`.

  Workaround: Verify the descriptions for severity event states by using the ESXCLI command `esxcli hardware ipmi sdr list` and Lenovo XCC.

Security Features Issues

- **If you use an RSA key size smaller than 2048 bits, RSA signature generation fails**

Starting from vSphere 8.0, ESXi uses the OpenSSL 3.0 FIPS provider. As part of the FIPS 186-4 requirement, the RSA key size must be at least 2048 bits for any signature generation, and signature generation with SHA1 is not supported.

Workaround: Use RSA key size larger than 2048.