VMware vCenter Server 8.0 Update 3 Release Notes

VMware vSphere 8.0

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

https://docs.vmware.com/

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

Contents

1 Introduction 4

2 General Availability 5

3 What's New 6

4 Earlier Releases of vCenter Server 8.0 15

5 Product Support Notices 16

6 Patches Contained in This Release 19

Patch for VMware vCenter Server 8.0 Update 3 19

Download and Installation 19

7 Resolved Issues 21

vSphere Client and vCenter Issues 21

Miscellaneous Issues 22

vSphere Lifecycle Manager Issues 22

High Availability and Fault Tolerance Issues 23

8 Known Issues 24

vCenter Server and vSphere Client Issues 24

vSphere Lifecycle Manager Issues 25

Installation, Upgrade, and Migration Issues 25

vSphere Cluster Service Issues 27

Security Features Issues 27

9 Known Issues from Previous Releases 28

Installation, Upgrade, and Migration Issues 28

Miscellaneous Issues 33

Networking Issues 36

Storage Issues 37

vCenter Server and vSphere Client Issues 38

Virtual Machine Management Issues 40

vSphere Lifecycle Manager Issues 40

VMware Host Client Issues 42

Security Features Issues 42

Introduction

1

vCenter Server 8.0 Update 3 | 25 JUN 2024 | GA ISO Build 24022515

Check for additions and updates to these release notes.

General Availability

2

This vCenter Server 8.0 Update 3 release is a General Availability (GA) designation. For more information on the vSphere 8.0 IA/GA Release Model of vSphere Update releases, see The vSphere 8 Release Model Evolves.

What's New

3

This release resolves CVE-2024-37087. For more information on this vulnerability and its impact on VMware products, see VMSA-2024-0013.

For overview of the new features in vSphere 8.0 Update 3, see What's New in vSphere 8 Update 3 and the following release notes:

- vSphere IaaS Control Plane (former VMware vSphere With Tanzu)
- Tanzu Kubernetes releases
- VMware vSAN
- VMware Host Client
- VMware OVF Tool

The list of new features and enhancements that follows adds some of the key highlights for vSphere 8.0 Update 3:

- **DPU/SmartNIC**
  - **High Availability with VMware vSphere Distributed Services Engine:** Starting with ESXi 8.0 Update 3, vSphere Distributed Services Engine adds support for 2 data processing units (DPUs) to provide high availability or increase offload capacity per ESXi host. Dual-DPU systems can use NVIDIA or Pensando devices. In ESXi 8.0 Update 3, dual-DPU systems are supported by Lenovo server designs. For more information, see High Availability with VMware vSphere Distributed Services Engine.
- **vSphere IaaS control plane**
  - **Support for running vSphere IaaS control plane on vSAN stretched clusters:** vSphere 8.0 Update 3 adds support for vSphere IaaS control plane, formerly known as vSphere with Tanzu, on vSAN stretched clusters. This capability is available only for greenfield installations of the IaaS control plane on stretched vSAN clusters. If you already have IaaS control pane or vSphere with Tanzu plus Content Library running, you need to redeploy it before you use the vSAN Stretched Cluster Storage Policies.

    Before configuring this capability, read Running vSphere IaaS Control Plane on vSAN Stretched Cluster and follow content on https://core.vmware.com.

- Virtual Machine Management
  - **New Virtual Machine Compute Policy for Best Effort Virtual Machine Evacuation**: vSphere 8.0 Update 3 adds a compute policy for best effort evacuation of virtual machines on ESXi hosts that are entering maintenance mode. When the host enters maintenance mode, all VMs are shut down. If shutdown fails, the VMs are powered off. If power off fails, you need to evacuate the VMs. With the new capability, when the VMs are in a powered-off state, vCenter attempts to power them on every few minutes on the best available ESXi host at the time. This policy overrules any DRS overrides set at the VM level, and the hosts on which the VMs are powered on might be different from the original host.
- vSphere Cluster Service
  - **Introducing Embedded vSphere Cluster Service (vCLS)**: vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of these VMs are managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). Earlier versions of vCLS are termed External vCLS. Issues previously encountered with External vCLS are resolved in this release of Embedded vCLS. While existing API compatibility is preserved, minor modifications to customer scripts or automation tools might be necessary. Other products and solutions that have defined business logic around External vCLS might not work with Embedded vCLS. See individual product documentation to understand the interoperability support and impact. For more information, see vSphere Cluster Services and content available from vSphere Technical Marketing on the Broadcom Support Portal.
- vSAN
  - **vSAN add-on licenses based on capacity per tebibyte (TiB):** With vSphere 8.0 Update 3, you can license your use of vSAN storage based on TiB capacity. The new capacity-based license replaces the core and CPU-based licenses. For more information on capacity reporting and licensing in vSAN, see License Requirements and Counting Cores for VMware Cloud Foundation and vSphere Foundation and TiBs for vSAN.
- Security
  - **TLS 1.3 and 1.2 support by using TLS profiles:** Starting with vSphere 8.0 Update 3, you can use TLS profiles to simplify the configuration of TLS parameters and improve supportability in your vSphere system. With vSphere 8.0 Update 3, you get a default TLS profile on ESXi and vCenter Server hosts, COMPATIBLE, which supports TLS 1.3 and some TLS 1.2 connections. For more information, see vSphere TLS Configuration.
  - **PingFederate Identity Provider for vSphere:** With vSphere 8.0 Update 3, you can configure PingFederate as an external identity provider in your vSphere system. For more information, see Configuring vCenter Server Identity Provider for PingFederate.

- Storage/Memory
  - **Memory Tiering:** vSphere 8.0 Update 3 launches in tech preview the Memory Tiering capability, which allows you to use NVMe devices that you can add locally to an ESXi host as tiered memory. Memory tiering over NVMe optimizes memory utilization by directing VM memory allocations to either NVMe devices or faster dynamic random access memory (DRAM) in the host. This allows you to increase your memory footprint and workload capacity, while reducing the total cost of ownership (TCO). For more details on the tech preview, see KB 95944.
  - **Fabric Notification support for SAN clusters:** ESXi 8.0 Update 3 introduces support for Fabric Performance Impact Notifications Link Integrity (FPIN-LI). With FPIN-LI, the vSphere infrastructure layer can manage notifications from SAN switches or targets, identifying degraded SAN links and ensuring only healthy paths are used for storage devices. FPIN can also notify ESXi hosts for storage link congestion and errors.
  - **Support for space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores and Config-vVol:** ESXi 8.0 Update 3 adds support for automatic space reclamation requests from guest operating systems on NVMe-backed vSphere Virtual Volumes datastores. ESXi 8.0 Update 3 also adds support for both command line-based and automatic unmap for vSphere Virtual Volumes objects of type Config-vVol, formatted with VMFS-6. For more information, see Reclaim Space on the vSphere Virtual Volumes Datastores.
  - **Manage the UNMAP load from ESXi hosts at a VMFS datastore level:** Starting with ESXi 8.0 Update 3, you can control the unmap load at datastore level to avoid time delays from space reclamation and reduce overall unmap load on the arrays in your environment. For more information, see Space Reclamation on vSphere VMFS Datastores.
  - **Windows Server Failover Clustering (WSFC) enhancements on vSphere Virtual Volumes:** vSphere 8.0 Update 3 adds support for a WSFC solution for NVMe-backed disks on vSphere Virtual Volumes. This capability allows NVMe reservations support in NVMe/TCP environments apart from Fibre Channel support for WSFC on vSphere Virtual Volumes. Virtual NVMe (vNVME) controllers are supported as the frontend for WSFC with NVMe-backed disks, not with SCSI-backed disks. For more information, see VMware vSphere® Virtual Volumes Support for WSFC.
  - **Support for active-active vSphere Metro Storage Cluster (vMSC) with vSphere Virtual Volumes:** vSphere 8.0 Update 3 introduces a new version of the VMware vSphere Storage APIs for Array Integration (VASA) to add support to active-active stretched storage clusters with vSphere Virtual Volumes, with active-active deployment topologies for SCSI block access between two sites. VASA version 6 includes new architecture and design for VASA Provider High Availability support for both stretched and non-stretched storage clusters. For more information, see Using Stretched Storage Clustering with Virtual Volumes.

  - **VMkernel port binding for NFS v4.1 datastores:** With vSphere 8.0 Update 3, you can bind an NFS 4.1 connection to a specific VM kernel adapter. If you use multipathing, you can provide multiple vmknics for each connection to ensure path isolation and security by directing NFS traffic across a specified subnet/VLAN, so that the NFS traffic does not use other vmknics. For more information, see Configure VMkernel Binding for NFS 4.1 Datastores.
  - **Support for nConnect for NFS v4.1 datastores:** ESXi 8.0 Update 3 adds support to multiple TCP connections for a NFS v4.1 volume, also referred to as nConnect. For NFS v4.1, multiple TCP connections are created for a single session that many datastores can share in parallel. It can be configured by using either the vSphere API or ESXCLI directly on an ESXi host. For more information, see Configure Multiple TCP Connections for NFS.
  - **Reduced time to inflate VMFS disks:** With vSphere 8.0 Update 3, a new VMFS API allows you to inflate a thin-provisioned disk to eagerzeroedthick (EZT) while the disk is in use and up to 10 times faster than previous alternative methods. During the inflation, all blocks are fully allocated and zeroed upfront to allow faster run-time performance. For more information, see Virtual Disk Options.
  - **Improved resiliency against memory corruption on RAM-heavy ESXi hosts**: vSphere 8.0 Update 3 adds proactive measures to prevent memory errors in systems with VMs of more than 1TB that might bring down an entire ESXi host and increase application downtime.
  - **Advanced setting to block deletion and removal of disks for VMs with snapshots:** ESX 8.0 Update 3 adds a per-host advanced option `blockDiskRemoveIfSnapshot` to prevent the removal of disks from a VM that has snapshots, even if you choose to delete files, which might lead to orphaned disks. For more information, see VMware knowledge base article 94545.
  - **Hardware accelerated move (clone operation) support on NVMe devices:** vSphere 8.0 Update 3 supports NVMe copy command for hardware accelerated move (also called clone blocks or copy offload) across or within NVMe namespaces that belong to the same NVMe subsystem.
  - **Granular monitoring for VASA provider accessibility and certification authentication status on ESXi host level:** With vSphere 8.0 Update 3, you can monitor connection status and ESXi authentication with VASA storage providers from the vSphere Client and re-authenticate hosts as necessary. For more information, see Reauthenticate VASA Client in ESXi.
- GuestOS
  - **Guest customization supports RHEL NetworkManager keyfile format:** In vSphere 8.0 Update 3, guest customization adds support for RHEL NetworkManager keyfile format and you can store network configuration in both keyfile and ifcfg format.

- Drivers/Network
  - **Support for Fibre Channel Extended Link Services (FC-ELS):** With vSphere 8.0 Update 3, you can use the command `esxcli storage fpin info set -e= <true/false>` to activate or deactivate the Fabric Performance Impact Notification (FPIN). The command saves the FPIN activation to both ConfigStore and the VMkernel System Interface Shell and persists across ESXi reboots. This is enabled by both Broadcom's lpfc and Marvell's qlnativefc drivers.
  - **Unified Enhanced Networking Stack (UENS) driver for the Elastic Network Adapter (ENA):** vSphere 8.0 Update 3 adds a UENS driver for ENA, which provides connectivity to the AWS underlay Virtual Private Cloud (VPC).
  - **Overlay Filters Supporting Tunnel End Point (TEP):** vSphere 8.0 Update 3 enhances the i40en, qedentv, and sfvmk NIC drivers by adding capabilities that expose overlay filters, supporting the TEP functionality.
  - **Broadcom Driver Updates:**
    - **Broadcom bnxtnet driver:** Adds support for NetQ RSS to facilitate future Unified Enhanced Networking Stack (UENS) support.
  - **Intel Driver Updates:**
    - **Intel i40en driver:** The number of queues per RSS engine is up from 4 to 8, in ENS mode it supports up to 16 queues.
    - **Intel icen driver:** Adds RoCE support on 1000GbE NIC.
    - **Intel irdman driver:** Adds RoCE support on 1000GbE NIC.
  - **Marvell Driver Updates:**
    - **Marvell qedentv driver:** The max number of queues for the Default RSS engine is up from 4 to 16 in ENS mode.
  - **Mellanox Driver Updates:**
    - **Mellanox nmlx5 driver:** Adds support for hardware Large Receive Offload (LRO) in Enhanced Data Path mode to increase inbound throughput of high-bandwidth network connections by reducing CPU overhead.
    - Adds support for dual-DPU servers.
  - **Pensando Driver Updates:**
    - **Pensando ionic_en driver:** Adds support for dual-DPU servers.
  - **Routine Inbox Driver Updates and Bug Fixes**
    - Broadcom bcm_mpi3
    - Broadcom lpfc
    - Cisco nfnic

  - Marvell qlnativefc
  - Microchip smartpqi
- CPU
  - **PCIe hot plug is updated for server platforms utilizing newer generation AMD Genoa and Intel Sapphire Rapid CPUs:** Starting with vSphere 8.0 Update 3, kernel hot plug is supported for newer generation CPUs such as AMD Genoa and Intel Sapphire Rapids.
  - **Support for Intel Xeon Max Series processors with integrated High Bandwidth Memory (HBM):** vSphere 8.0 Update 3 adds support for Intel Xeon Max Series processors (formerly with code name Sapphire Rapids HBM) with 64 GB of integrated HBM, aimed to enhance performance for workloads like high performance computing apps, artificial intelligence (AI), and machine learning (ML).
  - **CPU C-State Power virtualization:** With vSphere 8.0 Update 3, for use cases such as Virtualized Radio Access Network (vRAN) workloads, you can configure and control the C-State power of the physical CPUs dedicated to vRAN VMs from the vSphere Client.
  - **Cluster-wide option to retain virtual NUMA topology:** vSphere 8.0 Update 3 adds a setting to retain preconfigured vNUMA topology even if the VM moves, allowing for better NUMA topology tuning for VMs across all the hosts in the cluster. This is the advanced vCenter Server setting `VPXD_PersistVnuma` to keep the virtual NUMA topology on a cluster level under **Configure > Advanced Settings** in the vSphere Client.
- Analytics and Metrics
  - **vSphere green metrics with Running Average Power Limit (RAPL) technology**: Starting with vSphere 8.0 Update 3, in the **Advanced Performance Charts** in the vSphere Client you can see RAPL data on ESXi hosts that normally do not report individual subsystem power consumption, such as CPU and Memory, but only general host-level power consumption. With individual subsystem power consumption reports, you can plan on a more granular level to match your power and cooling budget.
  - **Set VM log levels without powering off the VMs:** With vSphere 8.0 Update 3, you can set log levels between `VMW_LOG_TRIVIA` and `VMW_LOG_DEBUG_3` to avoid log spew in a healthy running VM without a power cycle by using the `SetLogLevel` service in the vAPI infrastructure.
- GPU
  - **Support for switching between Time Sliced and Multi-Instance GPU (MIG) modes for NVIDIA virtual GPUs:** Starting with vSphere 8.0 Update 3, you do not need to reboot an ESXi host to switch between time sliced and MIG modes for NVIDIA virtual GPUs. vGPU VMs can automatically set the correct device mode according to their vGPU type.
  - **Zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks:** vSphere 8.0 Update 3 adds zero-copy support for vGPUs to enhance vSphere vMotion and vSphere DRS tasks by utilizing throughput of up to 100 Gbps.

  - **Support for heterogeneous vGPU profiles on physical GPUs:** Starting with vSphere 8.0 Update 3, you can set vGPU profiles with different types or sizes on a single physical GPU to achieve greater flexibility with vGPU workloads and better utilization of GPU devices. For more information, see Configuring vGPU Size.

- vSphere Lifecycle Management

  - **Support for parallel hardware and firmware upgrade with vSphere Lifecycle Manager:** With vCenter Server 8.0 Update 3, you can run parallel hardware and firmware remediation by using an integration between the vSphere Lifecycle Manager and the Hardware Support Manager.

  - **Non-disruptive certificates for ESXi:** With vCenter Server 8.0 Update 3, you can renew or replace ESXi certificates non-disruptively without the need for rebooting ESXi or setting up the maintenance window.

  - **VMware Photon™ 5.0 support for Update Manager Download Service (UMDS):** vSphere 8.0 Update 3 adds VMware Photon™ 5.0 to the supported Linux-based operating systems for installing UMDS. For more information, see Installing UMDS.

  - **Additional lifecycle management capabilities for standalone ESXi hosts:** Starting with vSphere 8.0 Update 3, you can add a standalone host to a data center or folder by importing an image from another ESXi host in the vCenter Server inventory or by using the current image on the host. For more information, see Managing Standalone ESXi Hosts with vSphere Lifecycle Manager Images. When you move a host out of a cluster that you manage with a vSphere Lifecycle Manager image to a data center or a folder, the host becomes standalone and can retain the image from the cluster. For more information, see Specifics of the Transitioning Workflow.

  - **Lifecycle management of standalone ESXi hosts with VMware NSX:** Starting with vSphere 8.0 Update 3, NSX Manager and vSphere Lifecycle Manager work together to coordinate remediation of standalone ESXi hosts with VMware NSX. For more information, see Using vSphere Lifecycle Manager Images for Standalone Hosts with NSX 4.2 and later.

  - **Convert baseline-managed clusters to clusters managed by a single vSphere Lifecycle Manager image:** Starting with vSphere 8.0 Update 3, to start managing a cluster with a single vSphere Lifecycle Manager image, you can use the image installed on one of the ESXi hosts inside the cluster managed with baselines. For more information, see Use an Image from a Host in the Cluster.

  - **Patch VMX-related security vulnerabilities without any disruption to your workloads:** Starting with vSphere 8.0 Update 3, you can use the **Live Patch** capability to apply VMX-related security patches and bug fixes to ESXi hosts in a cluster managed with a vSphere Lifecycle Manager image. A pre-check prior to remediation lists all suitable hosts in a cluster. After you activate **LivePatch** in the remediation settings, qualified hosts in a cluster do not require maintenance mode, or reboot, or VM migration during the update procedure.

  - **Customize vSphere Lifecycle Manager desired state images:** Starting with vSphere 8.0 Update 3, you can remove the Host Client and VMware Tools components from the base image, remove unnecessary drivers from vendor add-ons and components, and override existing drivers in a desired image. For more information, see Edit a vSphere Lifecycle Manager Image.
  - **Support for dual DPUs with vSphere Lifecycle Manager:** Starting with vSphere 8.0 Update 3, you can install dual DPUs on ESXi hosts and use the vSphere Lifecycle Manager workflow to upgrade the dual DPU system.
  - **Extended support for vSphere Configuration Profiles (VCP):** With vSphere 8.0 Update 3, you can use VCP with the following new capabilities:
    - **Support to baseline-managed clusters (formerly referred to as VUM clusters):** Having an image-managed cluster is no longer a prerequisite for using VCP. You can use VCP to configure either baseline-managed clusters or image-managed clusters.
    - **Support for vSphere Distributed Switch (VDS):** VCP is fully integrated with VDS and supports drift detection and remediation of VDS configurations at a cluster level.
    - **Firewall ruleset management:** You can manage custom firewall rules at a cluster level by using VCP.
    - **ESXi Lockdown Mode:** vSphere admins can use the VCP desired configuration document to enforce Lockdown Mode on all hosts in a cluster.
    - **Support for SNMP and PCI device configurations:** You can manage SNMP and PCI devices at a cluster level by using VCP.

    For more information, see Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level.

- **vSphere Client and vCenter**
  - **Warning on the maximum number of remote https connections for vCenter**: Starting with vSphere 8.0 Update 3, to prevent a vCenter system to become unresponsive due to exceeding the limit of 2048 https connections, you will see HTTP error code **503 Service Unavailable** and **x-envoy-local-overloaded: true** in the response headers.
  - **Merging the vSAN Management SDK with the Python SDK for the VMware vSphere API:** Starting with vSphere 8.0 Update 3, the vSAN Management SDK for Python is integrated into the Python SDK for the VMware vSphere API (pyVmomi). From the Python Package Index (PyPI), you can download a single package to manage vSAN, vCenter, and ESXi. This integration streamlines the discovery and installation process and enables automated pipelines instead of the series of manual steps previously.

- **vCenter Universally Unique Identifier (VC_UUID) field in the vSphere Client:** With vSphere 8.0 Update 3, the property table under **VMs > Virtual Machines** in the vSphere Client includes a new column that contains the VC_UUID. This identifier is automatically assigned to every virtual machine within a vCenter instance. The VC_UUID field helps clarify the correlation between the VM's UUID displayed on the switch's fabric and the VM name in vCenter.
- **Default HTTP response compression:** vSphere 8.0 Update 3 adds support for HTTP response compression by default on the edge proxies of vCenter and ESXi hosts, management traffic between vCenter and ESXi hosts, and for outgoing HTTP requests on sidecar proxies. HTTP response compression reduces the HTTP traffic in a vCenter Server system, shrinks page load time and speeds up API operations. You can deactivate HTTP response compression if required, but the capability does not require any changes in your environment and is backward compatible.
- **Remove restrictions on virtual machine operations from the vSphere Client:** Starting with vCenter Server 8.0 Update 3, in the vSphere Client under **Virtual Machine** > **Configure** > **Disabled Methods** you can remove restrictions on virtual machine operations such as migration. For more information, see VMware knowledge base article 2044369.

  Unified Management and Automation API Sessions: Starting with vCenter 8.0 Update 3, you can combine vSphere Management API (VIM) and vSphere Automation API (vAPI) sessions, effectively unifying authentication across SOAP and REST vCenter endpoints. For more information see Unified Management and Automation Session.
- **Migration of SPBM, SMS, EAM, and VLSM APIs to HTTP/JSON-based wire protocol:** Starting with vCenter Server 8.0 Update 3, the following 4 SOAP/XML service interfaces migrate to a HTTP/JSON-based wire protocol, providing OpenAPI 3.0 specifications for each, and integration with the REST API documentation:
  - The Storage Policy (SPBM) API, which simplifies the task of matching available storage to virtual machines.
  - VMware vCenter Storage Monitoring Service (SMS), which facilitates access to all vCenter storage information associated with VMware vCenter servers.
  - vSphere ESX Agent Manager API (EAM), which gives access to the objects that manage, monitor, and control lifecycle operations in vSphere.
  - Virtual Storage Lifecycle Management (VSLM) API that you use to manage first class disks (FCD).

  For more information, see What's New in vSphere Automation: vSphere APIs, JSON, OpenAPI.

Earlier Releases of vCenter Server 8.0

4

Features, resolved and known issues of vCenter Server are described in the release notes for each release. Release notes for earlier releases of vCenter Server 8.0 are:

- VMware vCenter Server 8.0 Update 2d Release Notes
- VMware vCenter Server 8.0 Update 1e Release Notes
- VMware vCenter Server 8.0 Update 2c Release Notes
- VMware vCenter Server 8.0 Update 2b Release Notes
- VMware vCenter Server 8.0 Update 2a Release Notes
- VMware vCenter Server 8.0 Update 1d Release Notes
- VMware vCenter Server 8.0 Update 2 Release Notes
- VMware vCenter Server 8.0 Update 1c Release Notes
- VMware vCenter Server 8.0 Update 1b Release Notes
- VMware vCenter Server 8.0 Update 1a Release Notes
- VMware vCenter Server 8.0 Update 1 Release Notes
- VMware vCenter Server 8.0c Release Notes
- VMware vCenter Server 8.0b Release Notes
- VMware vCenter Server 8.0a Release Notes

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the VMware vSphere 8.0 Release Notes.

For more information on vCenter Server supported upgrade and migration paths, please refer to VMware knowledge base article 67077.

Product Support Notices

5

- **Deprecation of APIs for vCenter for Windows to Linux migration:**

  Starting with vSphere 8.0 Update 3, APIs from vSphere 6.7.x for migrating vCenter for Windows to Linux during upgrade are deprecated and will be removed in a future vSphere release.

- **Deprecation of VMkernel API (vmkapi) version 2.5:**

  Starting with vSphere 8.0 Update 3, vmkapi version 2.5 is deprecated and will be removed in a future major release. This requires an update of any third-party component released for vSphere 6.7.

- **Removal of Integrated Windows Authentication (IWA)**

  vSphere 8.0 Update 3 is the final release to support Integrated Windows Authentication. IWA was deprecated in vSphere 7.0 and will be removed in the next major release. To ensure continued secure access, migrate from IWA to Active Directory over LDAPS or to Identity Federation with Multi-Factor Authentication. For more information, see vSphere Authentication with vCenter Single Sign-On and Deprecation of Integrated Windows Authentication.

- **Removal of VMware Enhanced Authentication Plug-in (EAP):**

  Starting with vSphere 8.0 Update 3, you cannot use EAP to log in to a vCenter system by using the vSphere Client. For more information, see Removing the deprecated VMware Enhanced Authentication Plugin (EAP) to address CVE-2024-22245 and CVE-2024-22250.

- **Deprecation of the vCenter Service Lifecycle Management API:**

  vSphere 8.0 Update 3 is deprecating the use of vCenter Service Lifecycle Management (vmonapi service) API and the service is not active by default, you must manually activate it. The service will be removed in a future release. For more information, see VMware knowledge base article 80775.

- **Removal of the internal runtime option execInstalledOnly:**

Starting with ESXi 8.0 Update 3, the internal runtime option that deactivates the security option execInstalledOnly is deprecated and will be removed in the next major release. The boot option execInstalledOnly, which helps protect hosts against ransomware attacks, will be activated on ESXi hosts by default in the next major release.

- **Deprecation of Storage DRS Load Balancer and Storage I/O Control (SIOC):**

  The Storage DRS (SDRS) I/O Load Balancer, SDRS I/O Reservations-based load balancer, and vSphere Storage I/O Control Components will be deprecated in a future vSphere release. Existing 8.x and 7.x releases will continue to support this functionality.

  The deprecation affects I/O latency-based load balancing and I/O reservations-based load balancing among datastores within a Storage DRS datastore cluster. In addition, enabling of SIOC on a datastore and setting of Reservations and Shares by using SPBM Storage policies are also being deprecated.

  Storage DRS Initial placement and load balancing based on space constraints and SPBM Storage Policy settings for limits are not affected by the deprecation.

- **Deprecation of vSphere Trust Authority (vTA):**

  Starting with vSphere 8.0 Update 3, vSphere Trust Authority is deprecated. vSphere continues to offer advanced workload attestation in its baseline functionality.

- **Removal of Patch Manager APIs:**

  Patch Manager APIs are supported in vSphere 8.x, but support will discontinue in a future release of vSphere. Instead of Patch Manager APIs, you can use the latest vSphere APIs, documented in the vSphere API automation reference guide.

- **Deprecation of locales:**

  Beginning with the next major release, we will be reducing the number of supported localization languages. The three supported languages will be:

  - Japanese
  - Spanish
  - French

  The following languages will no longer be supported:

  - Italian, German, Brazilian Portuguese, Traditional Chinese, Korean, Simplified Chinese

  Impact:

  - Users who have been using the deprecated languages will no longer receive updates or support in these languages.
  - All user interfaces, help documentation, and customer support will be available only in English or in the three supported languages mentioned above.

- **Removal of vSphere Lifecycle Manager baselines:**

Support for managing clusters with vSphere Lifecycle Manager baselines and baseline groups (legacy vSphere Update Manager workflows) will drop in a future vSphere release. Instead of baselines and baseline groups, you can use vSphere Lifecycle Manager images to perform tasks on a standalone host or on a cluster level such as install a desired ESXi version on all hosts in a cluster, install and update third-party software, update, and upgrade ESXi or firmware, generate recommendations, and use a recommended image for your cluster.

Patches Contained in This Release

6

Read the following topics next:

- Patch for VMware vCenter Server 8.0 Update 3
- Download and Installation

Patch for VMware vCenter Server 8.0 Update 3

Product Patch for vCenter Server containing VMware software fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.3.00000-24022515-patch-FP.iso
Build	24022515
Download Size	8507.3 MB
sha256checksum	f611bba1fca57bfc81a021b0de2433a1df284b5283e0750f49eb2272fdd908ed

Download and Installation

Log in to the Broadcom Support Portal to download this patch.

For download instructions, see Download Broadcom products and software.

1 Attach the `VMware-vCenter-Server-Appliance-8.0.3.00000-24022515-patch-FP.iso` file to the vCenter Server CD or DVD drive.

2 Log in to the appliance shell as a user with super administrative privileges (for example, **root**) and run the following commands:

   - To stage the ISO:

     ```
     software-packages stage --iso
     ```

   - To see the staged content:

     ```
     software-packages list --staged
     ```

- To install the staged rpms:

  ```
  software-packages install --staged
  ```

For more information on using the vCenter Server shells, see VMware knowledge base article 2100508.

For more information on patching vCenter Server, see Patching and Updating vCenter Server 8.0 Deployments.

For more information on staging patches, see Upgrading the vCenter Server Appliance.

Resolved Issues

7

Read the following topics next:

- vSphere Client and vCenter Issues
- Miscellaneous Issues
- vSphere Lifecycle Manager Issues
- High Availability and Fault Tolerance Issues

vSphere Client and vCenter Issues

- **PR 3308624: ESXi hosts fail to boot with an error such as Could not boot image: Invalid argument (http://ipxe.org/err/1c0de8)**

  After a change in the root certificate authority in a vCenter instance, when you use vSphere Auto Deploy to boot ESXi hosts, the TLS handshake with the vCenter might not succeed. As a result, booting of the hosts fails with a message such as `Could not boot image: Invalid argument (http://ipxe.org/err/1c0de8).`

  This issue is resolved in this release. For more information, see KB 96827.

- **PR 3210868: vCenter support bundles fill up the /storage/log partition**

  Due to a change in the vCenter permissions model, the vpxd service might not be able to delete the vCenter support bundles it temporarily stores in the `/storage/log` partition. As a result, you might see vCenter support bundles fill up the`/storage/log` partition.

  This issue is resolved in this release. If you already face the issue, you can safely delete such files manually.

- **PR 2993755: 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.x 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.x vCenter Server.

This issue is resolved in this release.

Miscellaneous Issues

- **New - PR 3330963: The PCI slot of a virtual device might change after upgrade from hardware version 19**

  When you upgrade a virtual machine with hardware version 19 and configured with EFI firmware to a newer virtual hardware version, the PCI device addresses of the virtual devices in the VM might change. As a result, the guest operating system of the VM might not apply configurations on a PCI device at a particular address.

  This issue is resolved in this release.

- **PR 3020259: 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`.

  This issue is resolved in this release.

vSphere Lifecycle Manager Issues

- **PR 3080596: During the generation of a new image recommendation by the vSphere Lifecycle Manager, the vpxd service fails due to memory exhaustion**

  If the desired state image on a vSphere Lifecycle Manager cluster has several additional Components, and especially if some of them are not up to date, during the generation of a new image recommendation, the memory required to calculate all possible combinations of Components that have new versions might grow to 85% of the overall memory of the vpxd service. As a result, the service might fail with `coreDump.#####` files. In the`/var/log/vmware/vmware-updatemgr/vum-server/vmware-vum-server.log`, you see errors such as:

```
2022-12-08T11:00:12.208 info vmware-vum-server[50005] [Originator@6876
sub=RecommendationEngine::ReUtil] [reUtil 108] GenerateImageUnitsCombination:
Generating ImageUnit combinations

2022-12-08T11:05:34.821 error vmware-vum-server[50005] [Originator@6876
sub=Default] Unable to allocate memory
```

This issue is resolved in this release.

High Availability and Fault Tolerance Issues

- **PR 3354058: vSphere HA advanced options das.respectVmVmAntiAffinityRules and das.respectVmVmAffinityRules do not work as expected**

  In some cases, the vSphere HA advanced options `das.respectVmVmAntiAffinityRules` and `das.respectVmVmAffinityRules` might not work as expected. For example, when you set `das.respectVmVmAntiAffinityRules` to `FALSE`, expecting that the vm-vm anti-affinity rules to be ignored during a vSphere HA failover, some VMs fail to failover with a `RuleViolation` fault. In other cases, you might expect a VM failover to respect the vm-vm anti-affinity rule as by default `das.respectVmVmAntiAffinityRules` is `TRUE`, but some VMs fail over and the rule is violated.

  This issue is resolved in this release.

Known Issues

8

Read the following topics next:

- vCenter Server and vSphere Client Issues
- vSphere Lifecycle Manager Issues
- Installation, Upgrade, and Migration Issues
- vSphere Cluster Service Issues
- Security Features Issues

vCenter Server and vSphere Client Issues

- **When you move a standalone ESXi host back to a cluster that you manage with vSphere Lifecycle Manager images, you might see an error in the vSphere Client**

  If you move out an ESXi host from a cluster that you manage with vSphere Lifecycle Manager images, it becomes a standalone host that is connected to a vCenter Server instance but is not part of any cluster. If you move such a host back to a cluster, in the **Updates** tab of the vSphere Client you might see an error such as `The host [IP] is not a vLCM managed standalone host`. This message does not indicate a functional issue.

  Workaround: Refresh the vSphere Client session or change the tab and return to the **Updates** tab.

- **You see the same license product name displayed multiple times in the vSphere Client**

  Starting with vSphere 8.0 Update 2b, you can apply a single Solution License to the components of VMware vSphere Foundation and VMware Cloud Foundation. In different screens in the vSphere Client, you might see the license product name multiplied by the number of components for which you use the license. For example, if you use a vSphere 8 Enterprise Plus for VMware Cloud Foundation license for 3 components, such as ESXi, vCenter, and vSphere with Tanzu, you see the name listed 3 times.

  Workaround: None

- **An external gateway firewall might block vSphere Pod traffic to clusterIPs**

When you deploy Supervisor Services on vSphere Pods with Supervisors configured with the VDS stack, traffic from the vSphere Pod to ClusterIPs goes through the external gateway and a custom firewall can block it.

Workaround: For more information, see vSphere Pod Traffic to ClusterIP Time-outs.

vSphere Lifecycle Manager Issues

- **In the vSphere Client, you see a different count of components in a vSphere Lifecycle Manager image**

  When you prepare a vSphere Lifecycle Manager image for a cluster or standalone ESXi host, and manually add or remove components, you might see a different count of the components in **Updates > Hosts > Image > Components** and in the list of components when you click **Show details**. The discrepancy occurs when the vSphere Lifecycle Manager considers not to use some of the additional components, for example if a driver is being deprecated, or if you remove a component, such as VMware Tools or Host Client.

  Workaround: None. You can safely ignore the difference in the count, because it does not impact the actual list of components that vSphere Lifecycle Manager uses for the image installation.

Installation, Upgrade, and Migration Issues

- **New - If you do not open port 9087 in your firewall between ESXi hosts and vCenter, compliance checks and vSphere HA might fail after vCenter update to 8.0 Update 3**

  Starting from 8.0 Update 3, the vSphere Lifecycle Manager downloads updates for ESXi hosts by a HTTPS connection to the vCenter instance on port 9087. If you do not open port 9087 in your firewall between ESXi hosts and vCenter, you might see compliance check errors. For example, in the `lifecycle.log` you see messages such as:

  ```
  <Timestamp> In(14) lifecycle[2112988]: Downloader:373
  Opening https://<VC-FQDN>:9087/vum/repository/hostupdate/__micro-depot__vendor-
  DEL__DEL-ESXi-8.0-Addon-cumulative_metadata__index__.xml for download
  ```

  ```
  <Timestamp> Wa(12) lifecycle[2112988]: Downloader:210 Download failed: <urlopen
  error timed out>, 9 retry left...
  ```

  In addition, vSphere HA might fail to start.

  Workaround: Open port 9087 in your firewall between ESXi hosts and vCenter.

- **vSphere Lifecycle Manager baseline check compliance fails on ESXi hosts of version 7.0 GA in a vCenter 8.0 Update 3 system**

  Due to a limitation of the `esxupdate` memory resource pool on 7.0 GA ESXi hosts, a vSphere Lifecycle Manager baseline check compliance scan might fail on 7.0 GA ESXi host in a vCenter 8.0 Update 3 system with error such as `The host returns esxupdate error codes: -1`.

The host `esxupdate.log` contains an error such as:

```
<Timestamp> esxupdate: <PID>: esxupdate: ERROR: vmware.runcommand.RunCommandError:

Error running command '['/sbin/smbiosDump']': [Errno 12] Cannot allocate memory
```

The `vmkernel.log` contains error such as:

```
<Timestamp> cpu0:<PID>)Admission failure in path: host/vim/vmvisor/esxupdate/
python.<PID>:python.<PID>:uw.<PID>
```

This issue is specific for 7.0 GA ESXi hosts, as `esxupdate` memory allocation is increased in 7.0 Update 1 and later.

Workaround: Use an ISO upgrade baseline to upgrade the 7.0 GA ESXi hosts to a 8.x version or revert the vCenter system back to version 7.0 and upgrade the hosts to a version of 7.0 Update 1 and later.

- **If you click NEXT before the progress bar reaches 100% in a vCenter Server Lifecycle Manager plug-in upgrade, the Server Lifecycle Manager service fails**

  In the **Upgrade Plug-in** step of the vCenter Server Lifecycle Manager plug-in upgrade wizard, if you click **NEXT** before the progress bar reaches 100%, the Server Lifecycle Manager service fails.

  Workaround: Wait for the progress bar to reach 100% before you click **NEXT**.

- **You see a "vc.health.error.dbjob3" warning after a vCenter upgrade**

  In the vSphere Client or the in Virtual Appliance Management Interface, you might see the warning `vc.health.error.dbjob3` after a vCenter upgrade although the vCenter overall health status is green. This issue does not affect vCenter operations, it can only affect historical statistics when some data does not roll up for more than 72 hours.

  Workaround: See Delete old tasks, events and statistics data in vCenter Server 5.x, 6.x, 7.x and 8.x how to clear stats data, if not relevant anymore.

- **Update to vCenter 8.0 Update 3 might fail with an error "Destination path '/storage/analytics/stage/...' already exists"**

  Starting with vCenter 8.0 Update 3, the analytics service stores telemetry log files in a new directory.

  In rare cases, if previous update attempts had failed and reverted, when patching to vCenter 8.0 Update 3, the attempt to copy the telemetry log files to the new location might fail because these files already exist from a prior patching attempt. In the `patchrunner.log` file, you see an error such as `Destination path '/storage/analytics/stage/XXXXXX_processed_logs' already exists`.

  Workaround: Delete all files under `/storage/log/vmware/analytics/stage`,`/storage/log/vmware/analytics/prod`, `/storage/analytics/stage`, and `/storage/analytics/prod`, and retry the update.

vSphere Cluster Service Issues

- **In the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected**

  vSphere 8.0 Update 3 introduces a redesign of vCLS to Embedded vCLS, which utilizes vSphere Pod technology. Deployment and lifecycle of such VMs are now managed within ESXi and are no longer managed by the vSphere ESX Agent Manager (EAM). On the **Summary** tab in the vSphere Client, you see the vSphere HA status of Embedded vSphere Cluster Service VMs as Unprotected, but this is expected, because vSphere HA does not manage Embedded vSphere Cluster Service VMs.

  Workaround: Ignore the information in the vSphere HA card on the **Summary** tab in the vSphere Client. For more information, see vSphere Cluster Services.

Security Features Issues

- **New - Adding a Standard Key Provider to your vCenter system from the vSphere Client fails with Vim.fault.DatabaseError**

  Starting with vCenter 8.0 Update 1c, you must use a RSA certificate when adding a Standard Key Provider to vCenter otherwise in the vSphere Client you see a `Vim.fault.DatabaseError` error. In the `vmafdd.log` file, you see the following line:

  ```
  <Timestamp> [vmafdd][ERROR] Certificate uses an unsupported signature algorithm (NID=ecdsa-with-SHA256). Only SHA-2 RSA algorithms are supported on the vCenter Server.
  ```

  Workaround: Use a RSA certificate with a SHA-2 digital signature algorithm when trusting a KMS.

Known Issues from Previous Releases

9

Read the following topics next:

- Installation, Upgrade, and Migration Issues
- Miscellaneous Issues
- Networking Issues
- Storage Issues
- vCenter Server and vSphere Client Issues
- Virtual Machine Management Issues
- vSphere Lifecycle Manager Issues
- VMware Host Client Issues
- Security Features Issues

Installation, Upgrade, and Migration Issues

- **You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade**

  During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

  ```
  Source ESX Agent Manager Configuration contains URLs that are not trusted by the
  System! Please refer to https://kb.vmware.com/s/article/93526 to trust the URLs:
  <LIST_OF_URLs>
  ```

  Or

  ```
  Source vSphere ESX Agent Manager (EAM) upgrade failed to obtain EAM URLs to check
  against trusted certificates by the System! Verify that the ESX Agent Manager
  extension is running properly on the source vCenter Server instance and https://
  VC_IP/eam/mob presents correct data. If log in to the MOB is not successful, try
  resolving the issue with https://kb.vmware.com/s/article/94934.
  ```

  Workaround: For more information, see VMware knowledge base articles 93526 and 94934.

- **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.

- **You see an error Failed to get ceip status in the Virtual Appliance Management Interface (VAMI) during update to vCenter Server 8.0 Update 1**

  During an update, vCenter stops and restarts the VMDir service and within this interval, if you try to log in to the VAMI, you might see an error such as **Failed to get ceip status**. This is expected and does not indicate an actual issue with the vCenter system.

  Workaround: Wait for the VMDir service to restart and refresh the Virtual Appliance Management Interface.

- **vCenter upgrade or update to 8.0 Update 2a or later fails during precheck with the error "VMCA root certificate validation failed"**

  If your vCenter system has a legacy VMCA root certificate dating back to version 5.x which does not have the Subject Key Identifier (SKID) extension, upgrades and updates to vCenter 8.0 Update 2 and later fail because the OpenSSL version 3.0 in 8.0 Update 2 is not compatible with legacy root certificates. vCenter Server 8.0 Update 2a adds a precheck to detect this issue and shows the error message **VMCA root certificate validation failed** if the source vCenter has VMCA root certificate without SKID.

  Workaround: Regenerate a VMCA root certificate by following the steps in VMware knowledge base article 94840.

- **A reduced downtime upgrade (RDU) on a vCenter system might fail when you use Update Planner**

  During RDU, if you use Update Planner, in the vSphere Client you might see an error such as: `Update 8.0.2.00000 for component vlcm is not found`.

  Workaround: For more information, see VMware knowledge base articles 94779 and 92659.

- **Failed parallel remediation by using vSphere Lifecycle Manager on one ESXi host might cause other hosts to remain in a pending reboot state**

  An accidental loss of network connectivity during a parallel remediation by using vSphere Lifecycle Manager might cause the operation to fail on one of the ESXi hosts. Remediation on other hosts continues, but the hosts cannot reboot to complete the task.

  Workaround: If an ESXi host consistently fails remediation attempts, manually trigger a reboot. For more information, see VMware knowledge base article 91260**.**

- **You see a security warning for the ESX Agent Manager configuration during the pre-check phase of a vCenter upgrade**

  During the pre-check phase of a vCenter update or upgrade, in the vSphere Client and logs, you might see an error such as:

  ```
  Source ESX Agent Manager Configuration contains URLs that are not trusted by the
  System! Verify following URLs and their respective statuses and follow KB 93526.
  ```

  Workaround: For more information, see VMware knowledge base article 93526.

- **Firmware compliance details are missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host**

  Firmware compliance details might be missing from a vSphere Lifecycle Manager image compliance report for an ESXi standalone host in two cases:

  a You run a compliance report against a standalone host managed with a vSphere Lifecycle Manager image from vSphere Client and then navigate away before the compliance report gets generated.

  b You trigger a page refresh after the image compliance reports are generated.

  In such cases, even when you have the firmware package available in the Desired State, the firmware compliance section remains empty when you revisit or refresh the vSphere Client browsing session. If you use GET image compliance API, then firmware compliance details are missing from the response.

  Workaround: Invoke the image compliance scan for a standalone host managed with a vSphere Lifecycle Manager image by using the vSphere Client and do not navigate away or refresh the browser. For API, use the Check image compliance API for fetching the firmware details as apposed to GET image compliance.

- **You see vCenter update status as failed although it completes successfully**

  A rare race condition might cause vCenter to report a successful update as failed. The issue occurs if during vCenter reboot `/storage/core` unmounts before the system acknowledges the **Installation complete** status. As a result, the update fails with an error such as `No such file or directory: '/storage/core/software-update/install_operation'`. In the `software-packages.logs`, you see errors such as:

  ```
  2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
  State._get using state file /etc/applmgmt/appliance/software_update_state.conf

  2023-08-17T10:57:59.229 [15033]INFO:vmware.appliance.update.update_state:Found
  operation in progress /storage/core/software-update/install_operation

  2023-08-17T10:57:59.229 [15033]ERROR:vmware.appliance.update.update_functions:Can't
  read JSON file /storage/core/software-update/install_operation [Errno 2] No such
  file or directory: '/storage/core/software-update/install_operation'
  ```

```
2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation
in progress is orphaned

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Operation
in progress is finished

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:Writing
to state file from State._get

2023-08-17T10:57:59.229 [15033]DEBUG:vmware.appliance.update.update_state:In
State._writeInfo writing to state file /etc/applmgmt/appliance/
software_update_state.conf

2023-08-17T10:57:59.229 [15033]INFO:vmware.vherd.base.software_update:Installation
failed. Please collect the VC support bundle.
```

Workaround: Check if vCenter restarts successfully and the vCenter health status is green, and ignore the failure report.

- **Patching to vCenter 8.0 Update 2 fails in IPv6 environments with no DNS server and hostname**

  When you update your vCenter system to 8.0 Update 2 from an earlier version of 8.x, if your system uses an IPv6 network without a hostname, such as PNID, and a DNS server, in the VMware Appliance Management Interface you might see an error such as `Data conversion/ Post install hook failed`.

  Workaround: Manually update the `/etc/hosts` file with the missing IPv6 loopback entry: `::1 localhost ipv6-localhost ipv6-loopback` and reboot the system.

  See this example:

  ```
  root@localhost []# cat /etc/hosts

  # Begin /etc/hosts (network card version)

  127.0.0.1 localhost.localdomain

  127.0.0.1 localhost

  ::1 localhost ipv6-localhost ipv6-loopback
  ```

- **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.

- **During a reduced downtime upgrade (RDU), when configuring Target VM network settings, you see no network portgroups**

In very rare cases, during a reduced downtime upgrade of a single self-managed vCenter instance that uses a migration-based method, when a source vCenter VM has thin disk provisioning and the target vCenter cluster does not have enough storage to accommodate the required space for the default thick disk mode selected by the validation process, you might see no network portgroups in the **Target VM deployment** wizard. In the vSphere Client, if you select **Same Configuration** in the **Deployment type** step of the **Target VM deployment** wizard, you see an empty error message in the **Network Settings** screen and no portgroups available.

Workaround: In the **Deployment type** step of the **Target VM deployment** wizard, select **Detailed Configuration**.

- **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.

- **URL-based patching or file-based backup of vCenter 8.0 Update 2 might fail due to OpenSSL noncompliance to Federal Information Processing Standards (FIPS)**

  With vCenter 8.0 Update 2, OpenSSL works only with Diffie-Hellman parameters compliant to NIST SP 800-56A and FIPS 140-2. For URL-based patching or file-based backup of vCenter 8.0 Update 2 systems, FTPS servers in your environment must support the following ciphers:

OpenSSL Cipher Suite	Name AT-TLS Cipher Suite Name	Hexadecimal Value
DHE-RSA-AES256-SHA	TLS_DHE_RSA_WITH_AES_256_CBC_SHA	39
DHE-DSS-AES256-SHA	TLS_DHE_DSS_WITH_AES_256_CBC_SHA	38
AES256-SHA	TLS_RSA_WITH_AES_256_CBC_SHA	35
EDH-RSA-DES-CBC3-SHA	TLS_DHE_RSA_WITH_3DES_EDE_CBC_SHA	16
EDH-DSS-DES-CBC3-SHA	TLS_DHE_DSS_WITH_3DES_EDE_CBC_SHA	13
DES-CBC3-SHA	TLS_RSA_WITH_3DES_EDE_CBC_SHA	0A
DHE-RSA-AES128-SHA	TLS_DHE_RSA_WITH_AES_128_CBC_SHA	33
DHE-DSS-AES128-SHA	TLS_DHE_DSS_WITH_AES_128_CBC_SHA	32
AES128-SHA	TLS_RSA_WITH_AES_128_CBC_SHA	2F

  Workaround: Make sure your file servers are FIPS compliant.

- **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.

- **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

- **In Hybrid Linked Mode, the cloud vCenter is not able to discover plug-ins deployed on an on-prem vCenter**

Hybrid Linked Mode allows you to link your cloud vCenter Server instance with an on-premises vCenter Single Sign-On domain, but the cloud vCenter might not be able to discover plug-ins deployed on the on-prem instance because it does not have the necessary permissions.

Workaround: Install the vCenter Cloud Gateway in your on-premises environment and either browse the plug-ins deployed on the on-prem instance from the VMware Cloud Console or directly from the vSphere Client on the on-prem vCenter.

- **You see incorrect Maximum Size value for thin-provisioned virtual disks**

  In the vSphere Client, when you look at the settings of a VM with a thin-provisioned virtual disk, you might see the **Maximum Size** value for the disk larger than the capacity of the datastore where the disk is located. For example, if a datastore capacity is 100GB of which 90GB are available, and a thin-provisioned virtual disk has 50GB capacity, of which only 10GB are utilized, you might see a **Maximum Size** value of 140 GB, adding the available datastore capacity to the overall disk capacity, not its actual utilization.

  Workaround: None

- **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 cannot customize firewall rule configuration with the option 'Only allow connections from the following networks' on ESXi hosts**

  Starting with vSphere 8.0 Update 2, you cannot customize firewall rule configuration with the option **Only allow connections from the following networks** on ESXi hosts. For example, in the VMware Host Client, when you navigate to **Networking > Firewall rules**, select **DHCP client**, provide an IP and check **Only allow connections from the following networks**, the operation fails with an error such as `Operation failed, diagnostics report: Invalid operation requested: Can not change allowed ip list this ruleset, it is owned by system service.`. This is expected behavior.

  Workaround: None

- **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.

Networking Issues

- **Overlapping hot-add and hot-remove operations for DirectPath I/O devices might fail**

  With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. However, if you run several operations at the same time, some of the overlapping tasks might fail.

  Workaround: Plan for 20 seconds processing time between each hot-add or hot-remove operation for DirectPath I/O devices.

- **Hot adding and removing of DirectPath I/O devices is not automatically enabled on virtual machines**

  With vSphere 8.0 Update 1, by using vSphere API you can add or remove a DirectPath I/O device without powering off VMs. When you enable the hotplug functionality that allows you to hot add and remove DirectPath I/O devices to a VM, if you use such a VM to create an OVF and deploy a new VM, the new VM might not have the hotplug functionality automatically enabled.

  Workaround: Enable the hotplug functionality as described in Hot-add and Hot-remove support for VMDirectPath I/O Devices.

- **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.

- **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.

- **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.

- **Migration of a First Class Disk (FCD) might fail and the FCD remains in a tentative state**

  In certain scenarios, when you migrate a FCD to another datastore by invoking the `RelocateVStorageObject` API, the operation might intermittently fail and the FCD remains in a tentative state. As a result, you cannot complete any other operation on the FCD. For example, if you try another migration, in the backlog you see the error `com.vmware.vim.fcd.error.fcdAlreadyInTentativeState`.

  Workaround: Reconcile the source datastore of the FCD by following the steps described in VMware knowledge base article 2147750.

- **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.

vCenter Server and vSphere Client Issues

- **You see overlapped labels for parameters in the Edit VM Startup/Shutdown Configuration dialog box**

  In the vSphere Client, when you select an ESXi host and click **Configure** > **Virtual Machines** > **VM Startup/Shutdown** > **Edit**, you see overlapped labels for some parameters in the **Edit VM Startup/Shutdown Configuration** dialog box that opens. The overlapped labels are as follows:

  - System influence: labels the checkbox **Automatically start and stop the virtual machines with the system**.
  - Startup delay: numeric value that specifies the delay time that a host waits before powering on the next virtual machine in automatic startup configuration.
  - Shutdown delay: numeric value that defines the maximum time the ESXi host waits for a shutdown command to complete, and the option **Continue if VMware Tools is started**.
  - Shutdown action: such as Guest Shutdown, Power Off, Suspend, and None.

  Workaround: None. See the screenshot to figure out the sequence of labels:

  Figure 9-1.

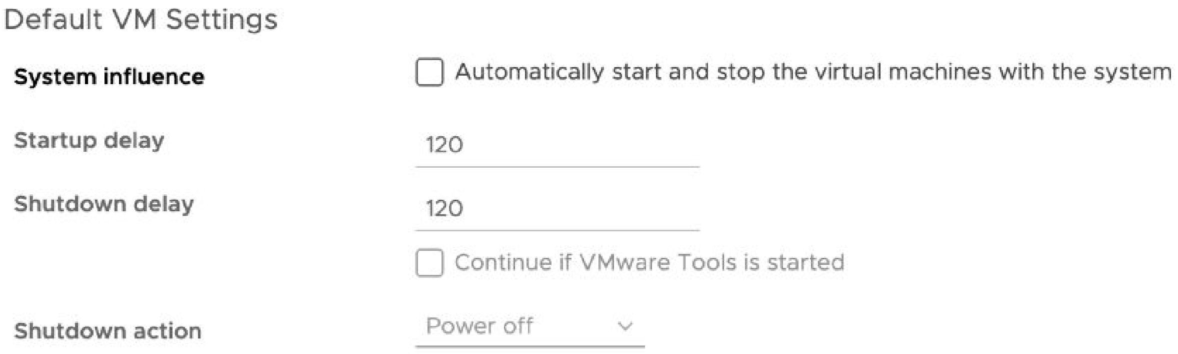

Default VM Settings

System influence ☐ Automatically start and stop the virtual machines with the system

Startup delay 120

Shutdown delay 120

☐ Continue if VMware Tools is started

Shutdown action Power off ⌄

- **VMware vCenter Lifecycle Manager might fail to load latest certificates and cannot complete a range of tasks**

VMware vCenter Lifecycle Manager might fail to load the latest certificates when you opt for a non-disruptive certificate renewal in vCenter 8.0 Update 2. As a result, any functionality relying on vCenter Lifecycle Manager, which provides the underlying VMware vCenter Orchestration platform, such as the Update Planner, vSphere+ vCenter Lifecycle Management Service, and reduced downtime upgrade for vCenter, might fail.

Workaround: Restart the vCenter Lifecycle Manager to get the latest certificates. For more information, see VMware knowledge base article 2109887.

- **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.

- **You do not see the option to push root certificates to vCenter hosts**

  In the **Add Trusted Root Certificate** screen under the **Certificate Management** tab in the vSphere Client, you do not see the option **Start Root certificate push to vCenter Hosts**.

  Workaround: This change in the **Add Trusted Root Certificate** screen is related to the non-disruptive certificate management capability introduced with vCenter 8.0 Update 2 and is expected. For more information, see Non-disruptive Certificate Management.

- **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.

- **A scheduled task fails and doesn't schedule further runs**

  With vSphere 8.0 Update 2, if a vCenter user is unauthorized or unauthenticated, all scheduled tasks they own fail and cannot be scheduled until the user privileges are restored or a different vSphere user takes over the scheduled tasks. In the vSphere Client, you see messages for failed tasks and reasons for the failure.

  Workaround: Any vSphere user with sufficient privileges to edit scheduled tasks, including the current task owner with restored privileges, can click **Edit** and submit the scheduled task, without actually changing the scheduled task. For more information, see Scheduling vSphere Tasks.

Virtual Machine Management Issues

- **In a mixed vCenter environment, when you clone a VM with First Class Disk (FCD) attached and delete it, the attached FCD in the cloned VM is also deleted**

  In a mixed vCenter environment, where vCenter is on version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, when you clone a VM with FCD, the parameter `KeepAfterDeleteVM` is set to `FALSE` by default. As a result, if the cloned VM is deleted, the attached cloned FCD is also deleted.

  Workaround: In a mixed vCenter environment, where vCenter is of version 8.0 Update 2 or later and ESXi is on version 7.0 Update 3 or earlier, set the `KeepAfterDeleteVM` parameter to `TRUE` by using the FCD API : `setVStorageObjectControlFlags`. You can invoke the FCD API at `https://<VC_IP>/mob/?moid=VStorageObjectManager&method=setVStorageObjectControlFlags` and pass the control flag : `KeepAfterDeleteVM`.

vSphere Lifecycle Manager Issues

- **You do not see a warning or error when entering non-numeric values for a desired cluster configuration setting in the vSphere Client that requires numbers**

  When you edit the host settings of the draft configuration for a cluster that uses vSphere Configuration Profiles, you can enter non-numeric values in a field that expects only numbers and you see no error or warning. For example, if you set non-numeric characters in the setting for syslog rotations, `esx/syslog/global_settings/rotations`, which expects a number, the **Edit** dialog box closes without an error and seems to save the value, but the setting actually keeps the previous valid value.

  Workaround: Use numeric values in fields that expect numbers. Use numbers in text inputs that expect numbers.

- **You cannot edit the VMware vSphere Lifecycle Manager Update Download scheduled task**

  In the vSphere Client, when you navigate to a vCenter Server instance and select **Scheduled Tasks** under the **Configure** tab, if you select the **VMware vSphere Lifecycle Manager Update Download** task and click **Edit**, you cannot modify the existing settings.

  Workaround: You can edit the **VMware vSphere Lifecycle Manager Update Download** task by following the steps in the topic Configure the vSphere Lifecycle Manager Automatic Download Task.

- **Manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations**

  Some of all manually applied security advanced options on a vCenter system might not persist across vSphere Lifecycle Manager operations, such as upgrade, update, backup, or restore.

Workaround: Re-apply the manual settings after the vSphere Lifecycle Manager task completes.

- **If a parallel remediation task fails, you do not see the correct number of ESXi hosts that passed or skipped the operation**

  With vSphere 8.0, you can enable vSphere Lifecycle Manager to remediate all hosts that are in maintenance mode in parallel instead of in sequence. However, if a parallel remediation task fails, in the vSphere Client you might not see the correct number of hosts that passed, failed, or skipped the operation, or even not see such counts at all. The issue does not affect the vSphere Lifecycle Manager functionality, but only the reporting in the vSphere Client.

  Workaround: None.

- **You see an authentication error on the vSphere Lifecycle Manager home view in one of several linked vCenter instances**

  After an update of a linked vCenter system, access to the vSphere Lifecycle Manager home page in the vSphere Client from one of the linked vCenter instances might fail. When you select **Menu > Lifecycle Manager**, you see the error `Authentication failed, Lifecycle Manager server could not be contacted`. The issue also affects vSphere Lifecycle Manager baseline pages and workflows. Workflows with vSphere Lifecycle Manager images are not affected.

  Workaround: Log in to the vSphere Client from another vCenter instance in the linked environment or restart the **vsphere-ui** service to resolve the issue.

- **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.

- **You cannot encrypt virtual machines when connected to a vCenter version earlier than 8.0 Update 1**

  When you use the vSphere Client to connect to a vCenter system of version 7.x or earlier than 8.0 Update 1, and try to encrypt a VM either in the **New Virtual Machine** wizard or in the **Edit Settings** dialog of an existing VM, you see errors such as `Operation failed! RuntimeFault.Summary and A general runtime error occurred. Key /default KMS cluster not found`. The task completes successfully when you use the vSphere Client to log in to a vCenter system of version 8.0 Update 1 or later.

  Workaround: Use the vSphere Client from the vCenter instance of version earlier than 8.0 Update 1 to encrypt the VM. Alternatively, you can enable VM encryption on another vCenter instance of version 8.0 Update 1 and later, and migrate the already encrypted VM to the vCenter instance of earlier version.