# VMware ESXi 8.0 Update 3 Release Notes

VMware vSphere 8.0

vmware by Broadcom

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

# Introduction

1

VMware ESXi 8.0 Update 3 | 25 JUN 2024 | ISO Build 24022510

Check for additions and updates to these release notes.

# Initial Availability

2

This ESXi 8.0 Update 3 release is an Initial Availability (IA) 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-2024-37085 and CVE-2024-37086. For more information on these vulnerabilities and their 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 storage performance 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.
  - **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.

- For vSphere Quick Boot support in ESXi 8.0 Update 3, see Understanding ESXi Quick Boot Compatibility.

- **ESXi 8.0 Update 3 adds support for vSphere Quick Boot to Marvel OCTEON Fusion CNF105xx drivers.**

For the full list of supported servers and drivers, see the VMware Compatibility Guide.

# Earlier Releases of ESXi 8.0

4

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 2c Release Notes
- VMware ESXi 8.0 Update 1d Release Notes
- VMware ESXi 8.0 Update 2b Release Notes
- VMware ESXi 8.0 Update 2 Release Notes
- VMware ESXi 8.0 Update 1c Release Notes
- VMware ESXi 8.0 Update 1a Release Notes
- 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.

# 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:

    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.

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

- VMware ESXi 8.0 Update 3
- Patch Download and Installation

## VMware ESXi 8.0 Update 3

### Build Details

### VMware vSphere Hypervisor (ESXi ISO) image

| | |
|---|---|
| **Download Filename**: | VMware-VMvisor-Installer-8.0U3-24022510.x86_64.iso |
| **Build**: | 24022510 |
| **Download Size**: | 605.63 MB |
| **SHA256 checksum**: | 05ce214069a3e23265881fbd7a949fb93a99aa13059c6cac31920196e97ba4a1 |
| **Host Reboot Required**: | Yes |
| **Virtual Machine Migration or Shutdown Required**: | Yes |

### VMware vSphere Hypervisor (ESXi) Offline Bundle

| | |
|---|---|
| **Download Filename**: | VMware-ESXi-8.0U3-24022510-depot.zip |
| **Build**: | 24022510 |
| **Download Size**: | 599 MB |
| **SHA256 checksum:** | 82e963b196c9fca6ae2e25de2bfd353e221bf86f041f561b3ef42815e90eb604 |
| **Host Reboot Required**: | Yes |
| **Virtual Machine Migration or Shutdown Required**: | Yes |

### 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 |
|---|---|---|
| ESXi80U3-24022510 | Enhancement | Important |

### 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-8.0U3-24022510-standard |
| ESXi-8.0U3-24022510-no-tools |

### ESXi Image

| Name and Version | Release Date | Category | Detail |
|---|---|---|---|
| ESXi_8.0.3-0.0.24022510 | 06/25/2024 | General | Bugfix image |

## Patch Download and Installation

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

For download instructions, see Download Broadcom products and software.

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.

For more information, see the Upgrading Hosts by Using ESXCLI Commands and the VMware ESXi Upgrade guide.

# Resolved Issues 7

Read the following topics next:



## Virtual Machine Management Issues

- **PR 3237088: A vSphere system with EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001 might become unresponsive after more than 1000 days of continuous uptime**

  Due to AMD erratum 1474, if the core-C6 (CC6) sleep state is active on your EPYC processors of the series 7002, 7Fx2, 7Hx2, and 7001, a core might fail to exit CC6 in about 1044 days after the last system hardware reset, which excludes reboots by using VMware QuickBoot. As a result, your vSphere system might become unresponsive.

  This issue is resolved in this release. The fix automatically deactivates CC6 until the next hardware reset.

- **PR 3035321: 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.

  This issue is resolved in this release.

## Storage Issues

- **PR 2931647: 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.

This issue is resolved in this release.

## Security Issues

- **PR 3373364: The CIMHttpServer firewall ruleset remains active even when you stop the CIM service**

  The `CIMHttpServer` firewall ruleset, which is for internal port 5988, might remain active even when you stop the CIM service.

  This issue is resolved in this release. The `CIMHttpServer` firewall ruleset is permanently deactivated.

## Miscellaneous Issues

- **New - PR 3359237: If a NVMe/TCP driver fails to process an I/O command, an ESXi host might disconnect from a storage array**

  If your ESXi hosts connect to storage arrays and perform read/write operations by using NVMe/TCP connections, when a NVMe/TCP driver fails to process an I/O command, ESXi might cancel the process with an abort command. As a result, an ESXi host might disconnect from a storage array and end previous I/O queues.

  This issue is resolved in this release.

- **PR 3297607: An ESXi host might fail with a purple diagnostic screen while connecting to an NVMe over TCP controller**

  Due to a rare issue in the nvmetcp driver, an ESXi host might fail with a purple diagnostic screen while connecting to an NVMe over TCP controller. In the backtrace, you see a warning such as `WARNING: NVMFDEV:xxxx Failed to connect controller`.

  This issue is resolved in this release.

- **PR 3355070: ESXi hosts might become unresponsive after VMFS heap exhaustion due to the pointer block cache size**

  The VMFS heap copy of the pointer block cache might fill up the heap memory and cause memory allocation failures for other operations that require memory from the VMFS heap. As a consequence, you might see some ESXi hosts becoming unresponsive. The issue is more likely to occur if the `maxAddressableSpaceTB` parameter is set to a higher value than the default of 32TB, as the time for cache refresh increases.

This issue is resolved in this release.

## Guest OS Issues

- **PR 3091866: You see I/O errors in the logs of Linux virtual machines**

  In rare cases, unmap requests from a Linux Guest OS might compete with other workflows running on a different host such as block allocation on the same resource cluster. As a result, you see I/O errors in the logs of Linux VMs.

  This issue is resolved in this release.

# Known Issues

8

Read the following topics next:

- Storage Issues
- Miscellaneous Issues
- Networking Issues
- Installation, Upgrade, and Migration Issues

## Storage Issues

- **vSphere vMotion operations of virtual machines residing on Pure-backed vSphere Virtual Volumes storage might time out**

  vSphere vMotion operations for VMs residing on vSphere Virtual Volumes datastores depend on the vSphere API for Storage Awareness (VASA) provider and the timing of VASA operations to complete. In rare cases, and under specific conditions when the VASA provider is under heavy load, response time from a Pure VASA provider might cause ESXi to exceed the timeout limit of 120 sec for each phase of vSphere vMotion tasks. In environments with multiple stretched storage containers you might see further delays in the Pure VASA provider response. As a result, running vSphere vMotion tasks time out and cannot complete.

  Workaround: Reduce parallel workflows, especially on Pure storage on vSphere Virtual Volumes datastores exposed from the same VASA provider, and retry the vSphere vMotion task.

- **In a vSphere Virtual Volumes stretched storage cluster environment, some VMs might fail to power on after recovering from a cluster-wide APD**

  In high scale Virtual Volumes stretched storage cluster environments, after recovering from a cluster-wide APD, due to the high load during the recovery some VMs might fail to power on even though the datastores and protocol endpoints are online and accessible.

  Workaround: Migrate the affected VMs to a different ESXi host and power on the VMs.

- **You see "Object or item referred not found" error for tasks on a First Class Disk (FCD)**

Due to a rare storage issue, during the creation of a snapshot of an attached FCD, the disk might be deleted from the Managed Virtual Disk Catalog. If you do not reconcile the Managed Virtual Disk Catalog, all consecutive operations on such a FCD fail with the `Object or item referred not found` error.

Workaround: See Reconciling Discrepancies in the Managed Virtual Disk Catalog.

## Miscellaneous Issues

- **The irdman driver might fail when you use Unreliable Datagram (UD) transport mode ULP for RDMA over Converged Ethernet (RoCE) traffic**

  If for some reason you choose to use the UD transport mode upper layer protocol (ULP) for RoCE traffic, the irdman driver might fail. This issue is unlikely to occur, as the irdman driver only supports iSCSI Extensions for RDMA (iSER), which uses ULPs in Reliable Connection (RC) mode.

  Workaround: Use ULPs with RC transport mode.

## Networking Issues

- **Connection-intensive RDMA workload might lead to loss of traffic on Intel Ethernet E810 Series devices with inbox driver irdman-1.4.0.1**

  The inbox irdman driver version 1.4.0.1 does not officially support vSAN over RDMA. Tests running 10,000 RDMA connections, usual for vSAN environments, might occasionally lose all traffic on Intel Ethernet E810 Series devices with NVM version 4.2 and irdman driver version 1.4.0.1.

  Workaround: None.

- **Capture of network packets by using the PacketCapture tool on ESXi does not work**

  Due to tightening of the rhttpproxy security policy, you can no longer use the PacketCapture tool as described in Collecting network packets using the lightweight PacketCapture on ESXi.

  Workaround: Use the pktcap-uw tool. For more information, see Capture and Trace Network Packets by Using the pktcap-uw Utility.

## Installation, Upgrade, and Migration Issues

- **New - Fresh installation or creating VMFS partitions on a Micron 7500 or Intel D5-P5336 NVMe drives might fail with a purple diagnostic screen**

  UNMAP commands enable ESXi hosts to release storage space that is mapped to data deleted from the host. In NVMe, the equivalent of UNMAP commands is a deallocate DSM request. Micron 7500 and Intel D5-P5336 devices advertise a very large value in one of the deallocate limit attributes, DMSRL, which is the maximum number of logical blocks in a single

range for a Dataset Management command. This leads to an integer overflow when the ESXi unmap split code converts number of blocks to number of bytes, which in turn might cause a failure of either installation or VMFS creation. You see a purple diagnostics screen with an error such as `Exception 14 or corruption in dlmalloc`. The issue affects ESXi 8.0 Update 2 and later.

Workaround: For a fresh ESXi installation, first install ESXi 8.0, deactivate UNMAP for the affected disk by using the command `esxcli storage core device vaai status set -D 0 -d <device-id>` and then upgrade to 8.0 Update 3. To create VMFS partitions, deactivate UNMAP for the affected disk by using the command `esxcli storage core device vaai status set -D 0 -d <device-id>` and then create VMFS as usual. You can reactivate UNMAP after you create a VMFS datastore, but if you delete or create a partition, UNMAP must remain deactivated.

- **The Cancel option in an interactive ESXi installation might not work as expected**

  Due to an update of the Python library, the **Cancel** option by pressing the **ESC** button in an interactive ESXi installation might not work as expected. The issue occurs only in interactive installations, not in scripted or upgrade scenarios.

  Workaround: Press the **ESC** key twice and then press any other key to activate the **Cancel** option.

- **Standard image profiles for ESXi 8.0 Update 3 show last modified date as release date**

  The **Release Date** field of the standard image profile for ESXi 8.0 Update 3 shows the **Last Modified Date** value. The issue is only applicable to the image profiles used in Auto Deploy or ESXCLI. Base images used in vSphere Lifecyce Manager workflows display the release date correctly. This issue has no functional impact. The side effect is that if you search for profiles by release date, the profile does not show with the actual release date.

  Workaround: Search by release version, such as 8.0.3.

# Known Issues from Previous Releases

9

Read the following topics next:



## Installation, Upgrade, and Migration Issues

- **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 update to ESXi 8.0 Update 2b by using esxcli software vib commands**

  Starting with ESXi 8.0 Update 2, upgrade or update of ESXi by using the commands `esxcli software vib update` or `esxcli software vib install` is not supported. If you use `esxcli software vib update` or `esxcli software vib install` to update your ESXi 8.0 Update 2 hosts to 8.0 Update 2b or later, the task fails. In the logs, you see an error such as:

```
ESXi version change is not allowed using esxcli software vib commands.

Please use a supported method to upgrade ESXi.

 vib = VMware_bootbank_esx-base_8.0.2-0.20.22481015 Please refer to the log file
for more details.
```

Workaround: If you are upgrading or updating ESXi from a depot zip bundle downloaded from the VMware website, VMware supports only the update command `esxcli software profile update --depot=<depot_location> --profile=<profile_name>`. For more information, see Upgrade or Update a Host with Image Profiles.

- **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.

- **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.

- **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

- **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

- **You might see compliance errors during upgrade to ESXi 8.0 Update 2b on servers with active Trusted Platform Module (TPM) encryption and vSphere Quick Boot**

  If you use the vSphere Lifecycle Manager to upgrade your clusters to ESXi 8.0 Update 2b, in the vSphere Client you might see compliance errors for hosts with active TPM encryption and vSphere Quick Boot.

  Workaround: Ignore the compliance errors and proceed with the upgrade.

- **If IPv6 is deactivated, you might see 'Jumpstart plugin restore-networking activation failed' error during ESXi host boot**

In the ESXi console, during the boot up sequence of a host, you might see the error banner `Jumpstart plugin restore-networking activation failed`. The banner displays only when IPv6 is deactivated and does not indicate an actual error.

Workaround: Activate IPv6 on the ESXi host or ignore the message.

- **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.

- **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.

- **TCP connections intermittently drop on an ESXi host with Enhanced Networking Stack**

  If the sender VM is on an ESXi host with Enhanced Networking Stack, TCP checksum interoperability issues when the value of the TCP checksum in a packet is calculated as `0xFFFF` might cause the end system to drop or delay the TCP packet.

  Workaround: Disable TCP checksum offloading on the sender VM on ESXi hosts with Enhanced Networking Stack. In Linux, you can use the command `sudo ethtool -K <interface> tx off`.

- **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

- **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.

- **Transfer speed in IPv6 environments with active TCP segmentation offload is slow**

  In environments with active IPv6 TCP segmentation offload (TSO), transfer speed for Windows virtual machines with an e1000e virtual NIC might be slow. The issue does not affect IPv4 environments.

  Workaround: Deactivate TSO or use a vmxnet3 adapter instead of e1000e.

- **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.

- **When you migrate a VM from an ESXi host with a DPU device operating in SmartNIC (ECPF) Mode to an ESXi host with a DPU device operating in traditional NIC Mode, overlay traffic might drop**

  When you use vSphere vMotion to migrate a VM attached to an overlay-backed segment from an ESXi host with a vSphere Distributed Switch operating in offloading mode (where traffic forwarding logic is offloaded to the DPU) to an ESXi host with a VDS operating in a non-offloading mode (where DPUs are used as a traditional NIC), the overlay traffic might drop after the migration.

  Workaround: Deactivate and activate the virtual NIC on the destination ESXi host.

- **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.

# vCenter Server and vSphere Client Issues

- **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.

- **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

# vSphere Lifecycle Manager Issues

- **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.

- **Even though you deactivate Lockdown Mode on an ESXi host, the lockdown is still reported as active after a host reboot**

  Even though you deactivate Lockdown Mode on an ESXi host, you might still see it as active after a reboot of the host.

  Workaround: Add users `dcui` and `vpxuser` to the list of lockdown mode exception users and deactivate Lockdown Mode after the reboot. For more information, see Specify Lockdown Mode Exception Users and Specify Lockdown Mode Exception Users in the VMware Host Client.