VMware vCenter Server 8.0 Update 1 Release Notes

VMware vSphere 8.0
Center Server 8.0 Update 1

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

vCenter Server 8.0 Update 1 | 18 APR 2023 | GA ISO Build 21560480

Check for additions and updates to these release notes.

General Availability

2

This vCenter Server 8.0 Update 1 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

- **vSphere Configuration Profiles:** vSphere 8.0 Update 1 officially launches vSphere Configuration Profiles, which allow you to manage ESXi cluster configurations by specifying a desired host configuration at the cluster level, automate the scanning of ESXi hosts for compliance to the specified Desired Configuration and remediate any host that is not compliant. vSphere Configuration Profiles require that you use vSphere Lifecycle Manager images to manage your cluster lifecycle, a vSphere 8.0 Update 1 environment, and Enterprise Plus or vSphere+ license. For more information, see Using vSphere Configuration Profiles to Manage Host Configuration at a Cluster Level.
- **With vSphere 8.0 Update 1, vSphere Distributed Services Engine adds support for**:
  - NVIDIA BlueField-2 DPUs to server designs from Lenovo (Lenovo ThinkSystem SR650 V2).
  - 100G NVIDIA BlueField-2 DPUs to server designs from Dell.
  - UPTv2 for NVIDIA BlueField-2 DPUs.
  - AMD Genoa CPU based server designs from Dell.
- **Recorder for required privileges on vCenter Server workflows:** With vSphere 8.0 Update 1, to help you identify the minimal set of privileges that are required to run a vCenter Server workflow, you can record workflows and store the required privileges. You can use the stored data to create specific roles for users that run the same workflows, similar to templates. Searching through vCenter Server workflow recordings allows you to retrieve lists of privilege checks along with the corresponding sessions, users, managed objects, and operation IDs (opIDs). For more information, see Performing Privilege Checks Operations.
- **Support for heterogenous virtual graphics processing unit (vGPU) profiles on the same GPU hardware:** vSphere 8.0 Update 1 removes the requirement that all vGPUs on a physical GPU must be of the same type and you can set different vGPU profiles, such as compute, graphics, or Virtual Desktop Infrastructure workload, on one GPU to save cost by higher GPU utilization and reduced workload fragmentation.

- **Integration of VMware Skyline™ Health Diagnostics™ with vCenter:** Starting with vSphere 8.0 Update 1, you can detect and remediate issues in your vSphere environment by using the VMware Skyline Health Diagnostics self-service diagnostics platform, which is integrated with the vSphere Client. For more information, see VMware Skyline Health Diagnostics for vSphere Documentation.
- **VM-level power consumption metrics**: Starting with vSphere 8.0 Update 1, you as a vSphere admin can track power consumption at a VM level to support the environmental, social, and governance goals of your organization.
- **NVSwitch support**: vSphere 8.0 Update 1 adds support to up to 8 GPUs with NVSwitch connections between them on a single ESXi host. With the NVSwitch technology, you can run high-performance computing (HPC) and AI applications such as deep learning, scientific simulations, and big data analytics, which require multiple GPUs working together in parallel.
- **Okta Identity Federation for vCenter:** With vSphere 8.0 Update 1, you can configure vCenter Server Identity Provider Federation for Okta as an external identity provider. For more information, see Configure vCenter Server Identity Provider Federation for Okta.
- **Support for Fault Tolerance of virtual machines that use a virtual TPM (vTPM) module**: With vSphere 8.0 Update 1, you can use FT for VMs with a vTPM module to ensure continuous availability and security for your mission critical VMs.
- **Quick Boot support for servers with TPM 2.0 chips:** With vSphere 8.0 Update 1, you no longer need to disable TPM 2.0 for Quick Boot, which allows you to save time and eliminate security gaps.
- **vSphere API for Storage Awareness (VASA) version 5 for vSphere Virtual Volumes**: VASA version 5 for vSphere Virtual Volumes enhances security, certification management for multi vCenter deployments, and usability support. VASA version 5 deprecates self-signed certificates but provides backward compatibility support for earlier VASA versions. Before upgrading to ESXi 8.0 Update 1, see VMware knowledge base article 91387 to ensure continued availability of vSphere Virtual Volumes datastores.
- **Sidecar files become regular files in Config-vVol instead of vSphere Virtual Volumes objects**: To improve scalability and performance of vSphere Virtual Volumes, starting with vSphere 8.0 Update 1, sidecar files are no longer vSphere Virtual Volumes objects but regular files in Config-vVol. When sidecar files are created as vSphere Virtual Volumes objects, for solutions such as First Class Disk (FCD) that create a large number of small sidecar files, they can cause an overhead of VASA operations such as binding and unbinding the volumes to a protocol endpoint. Creating sidecar files under the Config-vVol home directory removes such overhead from VASA and storage.

  After updating to ESXi 8.0 Update 1, new virtual machines, or virtual disks in the Config-vVol namespace are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article 90791.

- **Increased default capacity for vSphere Virtual Volumes objects of type Config-vVol**: Starting with vSphere 8.0 Update 1, to allow the use of folders over vSphere Virtual Volumes datastores as content repositories, vSphere Virtual Volumes objects of type Config-vVol are created by default as thin-provisioned 255 GB, up from 4 GB in earlier releases. Also, starting with ESXi 8.0 Update 1, Config-vVol objects use VMFS-6 format instead of VMFS-5.
- **NVMe over TCP support for vSphere Virtual Volumes**: vSphere 8.0 Update 1 adds NVMe over TCP support for vSphere Virtual Volumes.
- **Extended XCOPY support**: With vSphere 8.0 Update 1, you can use `EXTENDED COPY (XCOPY)` commands to optimize the data copy between VMFS datastores that are in different storage arrays, not only between datastores within the same array. `XCOPY` operations across storage arrays are supported only for VMFS 6 datastores and allow you to offload, migrate and clone workloads. While ESXi enables this feature, the actual data migration across the arrays must be implemented on the storage array side. For more information on how to use SCSI T10 commands, see Managing Hardware Acceleration on Block Storage Devices.
- **New file type for OSDATA volumes on SSD devices**: vSphere 8.0 Update 1 adds a new file system type, VMFSOS, specifically for the ESX-OSData system partition on local SSD devices, which allows you to continue using virtual flash resources on other devices. The new file type prevents cases when you format an ESX-OSData volume on a local SSD device, and `fsType` returns a file of type Virtual Flash File System (VFFS). As a result, the disk backing of the ESX-OSData volume is listed under the Virtual Flash resources in vCenter, but such a disk belongs to the ESX-OSData volume and is not a part of the Virtual Flash resource pool.
- **NFS traffic isolation enhancements**: With vSphere 8.0 Update 1, to isolate NFS traffic, you can bind an NFS3 datastore on an NFS share over the network to a VM kernel adaptor on an ESXi host from a cluster. VMKNIC port binding is not supported in this release for vSphere Virtual Volumes datastores backed with NFS3 configurations. For more information, see Configure VMkernel Binding for NFS 3 Datastore.
- **Fourfold increase in NVMe-oF namespace capacity**: Starting with vSphere 8.0 Update 1, you can use 32 paths to a NVMe-oF namespace from 8 in earlier releases.
- **Support for end-to-end NVMe stack without protocol translation**: With vSphere 8.0 Update 1, the Pluggable Storage Architecture (PSA) extends NVMe capabilities to support end-to-end NVMe stack without any protocol translation in any of the layers.

  vSphere 8.0 Update 1 also extends NVMe functionality by adding support for third-party multi-pathing plug-ins to control and manage NVMe arrays.
- **Increased maximum for Windows Server Failover Clusters (WSFC):** vSphere 8.0 Update 1 increases the maximum WSFC clusters per ESXi host to 16 from 3.
- **Scale the number of NVMe over TCP adapters**: vSphere 8.0 Update 1 increases the number of NVMe over TCP storage adapters per ESXi host to 8 from 2 to allow more redundancy and better bandwidth.

- **Support hot-add and hot-remove for VMDirectPath I/O devices**: With vSphere 8.0 Update 1, by using vSphere API you can add or remove a VMDirectPath I/O device without powering off VMs. For more information, see Hot-add Support for VMDirectPath I/O Devices.
- **Increased number of PCI passthrough devices per VM**: With vSphere 8.0 Update 1, you can add up to 64 PCI passthrough devices per VM, up from 32 in previous releases, to allow better support for Virtualized Radio Access Network (vRAN) use cases.
- **Local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts:** Starting with vSphere 8.0 Update 1, to meet demand of certain business use cases, you can manage local depot overrides for Remote Office and Branch Office (ROBO) standalone ESXi hosts in the vSphere Client and by using the vSphere Automation API. For more information, see Manage Depot Overrides for a Cluster or a Standalone Host.
- **Advanced filters in the vSphere Client:** With vSphere 8.0 Update 1, in addition to the existing quick filter option, you can set advanced filters on most inventory objects in the vSphere Client, such as VM lists, Hosts and Clusters, Datastores, and linked vCenters. You can use strings, such as Name, Status, and Guest OS, and enumerations, such as IP address. To refine your filters, you can use various logical operators and 2 conditions.
- **Increased maximum number of NFSv3 datastores that can be mounted with multiple connections:** Starting with vSphere 8.0 Update 1, you can configure up to 8 connections to NFSv3 datastores on an ESXi host, from 1 in earlier releases, and have up to 256 connections used by the NFSv3 datastores on the host, depending on the combination of the number of datastores and the number of connection configured for each datastore. For more information, see VMware knowledge base article 91481.
- **HTTP/JSON-based wire protocol as an alternative to SOAP/XML**: vSphere 8.0 Update 1 adds support for a new HTTP/JSON-based wire protocol as an alternative to SOAP/XML. The new protocol is described using the industry standard OpenAPI specification, version 3.0, and you can use it to access the popular VIM APIs. For more information, see the OpenAPI schema and Web Services Programming Guide, part of the vSphere Management SDK.

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 8.0c Release Notes
- VMware vCenter Server 8.0b Release Notes
- VMware vCenter Server 8.0a Release Notes
- VMware vSphere 8.0 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.

Installation and Upgrades for This Release

5

Before upgrading to VASA version 5, refer to VMware knowledge base article 91387 to ensure continued availability of your vSphere Virtual Volumes datastores.

Compatibility

6

Newly created virtual machines or virtual disks of version ESXi 8.0 Update 1 in the Config-vVol namespace of vSphere Virtual Volumes are not supported on ESXi hosts of earlier versions. For more information and resolution, see VMware knowledge base article 90791.

Product Support Notices

7

- Since the SHA-1 signature algorithm is no longer supported in vSphere 8.0, if you use Enhanced Linked Mode (ELM) to connect vCenter systems of versions 7.0x and 8.0x, you must not add a SHA-1 trusted root certificate to the 7.0x vCenter instance during upgrades, because the SHA-1 certificate might be propagated to the 8.0x vCenter and cause unexpected results, such as disconnecting ESXi hosts.

  If you already have 7.0x vCenter instances connected with ELM, and using a SHA-1 trusted root certificate, do not add an 8.0x vCenter instance to the group, because the SHA-1 certificate might be replicated to the 8.0x instance.

Patches Contained in This Release

8

Read the following topics next:

- Patch for VMware vCenter Server 8.0 Update 1
- Download and Installation

Patch for VMware vCenter Server 8.0 Update 1

Product Patch for vCenter Server containing VMware software fixes, security fixes, and third-party product fixes.

This patch is applicable to vCenter Server.

Download Filename	VMware-vCenter-Server-Appliance-8.0.1.00000-21560480-patch-FP.iso
Build	21560480
Download Size	7965.6 MB
sha256checksum	1016f37c466e46957879f4084e176e62d96803fb0d1ad16cb1b8f6fc0e852b24

Download and Installation

Download this patch from Broadcom Support Portal.

For download instructions for earlier releases, see Download Broadcom products and software.

1 Attach the `VMware-vCenter-Server-Appliance-8.0.1.00000-21560480-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

9

Read the following topics next:

- Security Features Issues
- Miscellaneous Issues
- Server Configuration Issues

Security Features Issues

- **Enabling TLS 1.0 on ESXi 8.0 hosts causes connection drops**

  ESXi 8.0 supports OpenSSL 3.0 and the only TLS protocol enabled by default is TLS 1.2. If you try to enable TLS 1.0 in `/UserVars/ESXiVPsDisabledProtocols` by using ESXCLI commands, this leads to connection drops.

  This issue is resolved in this release.

Miscellaneous Issues

- **You might see wrong metadata for delete operation of a vSphere Lifecycle Manager depot that you use to create an image to manage a standalone ESXi host**

  Starting with vSphere 8.0, you can create an image based on a vSphere Lifecycle Manager depot, online or offline, to manage the lifecycle of any standalone ESXi host that is part of your vCenter Server inventory. In rare cases, when you delete such a depot, the metadata for the delete task might not be correct. For example, you see the hostid and hostIP details populated under clusterName and clusterID, such as:

  ```
  clusterName = 10.161.153.136,

  clusterId = host-54,

  entityName = <null>,

  entityId = <null>,
  ```

The issue also occurs when you delete a depot downloaded by using VMware vSphere Update Manager Download Service (UMDS), which is part of the desired state of a standalone host. The issue has no impact on any vSphere Lifecycle Manager operations and affects only task metadata for standalone hosts.

This issue is resolved in this release.

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

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

  This issue is resolved in this release.

Server Configuration Issues

- **Changing an Input/Output Operations Per Second (IOPS) limit might cause a significant drop in the I/O throughput of virtual machines**

  When you change an IOPS limit based on the Storage I/O Control (SIOC) by using a Storage Policy Based Management (SPBM) policy, you might see significantly slower VM performance. Normally, when you set an SPBM policy, IOPS limits are handled by an I/O filter, while mClock, the default I/O scheduler, handles reservations and shares. Due to a logic fault, when you change an existing IOPS limit, I/Os might throttle at the mClock scheduler instead at the I/O filter. As a result, I/Os get with a significant delay to the I/O filter which causes a drop in the I/O throughput of virtual machines.

  This issue is resolved in this release. The fix makes sure that IOPS limits are handled by the I/O filter, while mClock handles the reservations and shares. For more information, see VMware knowledge base article 89951.

Known Issues

10

Read the following topics next:

- Installation, Upgrade, and Migration Issues
- Miscellaneous Issues
- Networking Issues
- vSphere Lifecycle Manager Issues

Installation, Upgrade, and Migration Issues

- **You cannot specify IP properties on virtual machines deployed by using a customized OVA on vCenter 8.0 Update 1**

  While deploying virtual machines by using a customized OVA file on a vCenter 8.0 Update 1 system, the field to provide IP properties of the virtual machine might not appear and you must manually add IP properties after the deployment.

  Workaround: If the VM deployment succeeds, you can manually set the IP properties from the **vApp Options** under the **Configure** tab of the newly deployed VM. If the VM deployment fails, you must first add the property `ovf:required="false"` to the properties with attribute `vmw:qualifiers='Ip'` in the OVF file before applying the workaround. For more information, see VMware knowledge base article 93677.

- **Upgrades to vCenter Server 8.0 Update 1 fail with an error Exception occurred in postInstallHook**

  If a vCenter Single Sign-On domain contains a capitalization such as `administrator@vSphere.Local` or `administrator@VSPHERE.LOCAL`, upgrades to vCenter Server 8.0 Update 1 might fail with the error `Exception occurred in postInstallHook`.

  In the vSphere Client, you see an error such as `wcp service failed to start`.

  In the `PatchRunner.log`, you see an error such as `An error occurred while starting service 'wcp'`.

In the `wcpsvc.log`, you see the following errors: `Failed to parse VC JWKS: invalid character '<' looking for beginning of value` or `Unable to get VC public key configuration : invalid character '<' looking for beginning of value`.

After such failure, you must restore vCenter from a backup or revert from a snapshot.

Workaround: See VMware knowledge base article 92436.

- **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.

- **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 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.

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.

Networking Issues

- **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.

- **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.

vSphere Lifecycle Manager Issues

- **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.

Known Issues from Previous Releases

11

Read the following topics next:

Installation, Upgrade, and Migration Issues

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

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

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

- **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 the auto negotiate option, or the option is deactivated, and the VMNIC link is set down and then up, the link remains down after upgrade to or installation of ESXi 8.0.

  Workaround: Use either of these 2 options:

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

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

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

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

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

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

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

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

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

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

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

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

Miscellaneous Issues

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

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

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

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

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

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

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

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

  Workaround: None.

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

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

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

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

See this script as an example:

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

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

```
~# reboot
```

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

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

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

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

- **Some ionic_en driver uplinks might work with just a single receive queue and you see slower performance in native mode**

  Pensando Distributed Services Platform (DSC) adapters have 2 high speed ethernet controllers (for example `vmnic6` and `vmnic7`) and one management controller (for example `vmnic8`):

  ```
  :~] esxcfg-nics -l
  ```

  ```
  vmnic6 0000:39:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:48 1500 Pensando Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
  ```

  ```
  vmnic7 0000:3a:00.0 ionic_en_unstable Up 25000Mbps Full 00:ae:cd:09:c9:49 1500 Pensando Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Ethernet Controller
  ```

  ```
  :~] esxcfg-nics -lS
  ```

  ```
  vmnic8 0000:3b:00.0 ionic_en_unstable Up 1000Mbps Full 00:ae:cd:09:c9:4a 1500 Pensando Systems DSC-25 10/25G 2-port 4G RAM 8G eMMC G1 Services Card, Management Controller
  ```

  The high-speed ethernet controllers `vmnic6` and `vmnic7` register first and operate with RSS set to 16 receive queues.

`:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6…Num of RSS-Q=16, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }`

However, in rare cases, if the management controller `vmnic8` registers first with the vSphere Distributed Switch, the high-speed ethernet controllers `vmnic6` or `vmnic7` uplink might end up operating with RSS set to 1 receive queue.`:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int networkinternal nic privstats get -n vmnic6…Num of RSS-Q=1, ntxq_descs=2048, nrxq_descs=1024, log_level=3, vlan_tx_insert=1, vlan_rx_strip=1, geneve_offload=1 }`

As a result, you might see slower performance in native mode.

Workaround: Reload the `ionic_en driver` on ESXi by using the following commands:`:~] esxcfg-module -u ionic_en:~] esxcfg-module ionic_en:~] localcli --plugin-dir /usr/lib/vmware/esxcli/int/ deviceInternal bind.`

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

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

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

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

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

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

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

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

  Workaround: See VMware knowledge base article 89638.

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

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

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

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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

  Workaround: None.

Storage Issues

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

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

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

  b If you register a VASA provider that supports VASA API version 3.5 with a VMware vCenter 8.0 system and upgrade the VASA API version to 4.0, even after the upgrade, you still see VASA API version 3.5.

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

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

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

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

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

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

  Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

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

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

  Workaround: Click the refresh icon.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Virtual Machine Management Issues

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

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

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

vSphere Lifecycle Manager Issues

- **If 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.

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

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

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

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

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

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

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

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

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

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

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

VMware Host Client Issues

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

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

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

Security Features Issues

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

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

  Workaround: Use RSA key size larger than 2048.