VMware vCenter Server 7.0 Update 3r Release Notes

VMware vSphere 7.0

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 vCenter Server 7.0 Update 3r | 17 JUN 2024 | ISO Build 24026615

Check for additions and updates to these release notes.

What's New

2

- This release resolves CVE-2024-37079, CVE-2024-37080, and CVE-2024-37081. For more information on these vulnerabilities and their impact on VMware products, see VMSA-2024-0012.

Earlier Releases of vCenter Server 7.0

3

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 7.0 are:

- VMware vCenter Server 7.0 Update 3q Release Notes
- VMware vCenter Server 7.0 Update 3p Release Notes
- VMware vCenter Server 7.0 Update 3o Release Notes
- VMware vCenter Server 7.0 Update 3n Release Notes
- VMware vCenter Server 7.0 Update 3m Release Notes
- VMware vCenter Server 7.0 Update 3l Release Notes
- VMware vCenter Server 7.0 Update 3k Release Notes
- VMware vCenter Server 7.0 Update 3j Release Notes
- VMware vCenter Server 7.0 Update 3i Release Notes
- VMware vCenter Server 7.0 Update 3h Release Notes
- VMware vCenter Server 7.0 Update 3g Release Notes
- VMware vCenter Server 7.0 Update 3f Release Notes
- VMware vCenter Server 7.0 Update 3e Release Notes
- VMware vCenter Server 7.0 Update 3d Release Notes
- VMware vCenter Server 7.0 Update 3c Release Notes
- VMware vCenter Server 7.0 Update 3a Release Notes
- VMware vCenter Server 7.0 Update 3 Release Notes
- VMware vCenter Server 7.0 Update 2d Release Notes
- VMware vCenter Server 7.0 Update 2c Release Notes
- VMware vCenter Server 7.0 Update 2b Release Notes
- VMware vCenter Server 7.0 Update 2a Release Notes
- VMware vCenter Server 7.0 Update 2 Release Notes
- VMware vCenter Server 7.0 Update 1c Release Notes

- VMware vCenter Server 7.0 Update 1a Release Notes
- VMware vCenter Server 7.0 Update 1 Release Notes
- VMware vCenter Server 7.0.0d Release Notes
- VMware vCenter Server 7.0.0c Release Notes
- VMware vCenter Server 7.0.0b Release Notes
- VMware vCenter Server 7.0.0a Release Notes

For internationalization, compatibility, installation, upgrade, open source components and product support notices, see the VMware vSphere 7.0 Release Notes.

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

Patches Contained in This Release 4

Read the following topics next:

- Patch for VMware vCenter Server Appliance 7.0 Update 3r
- Download and Installation

Patch for VMware vCenter Server Appliance 7.0 Update 3r

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-7.0.3.02000-24026615-patch-FP.iso
Build	24026615
Download Size	6789.2 MB
md5sum	fa90aee2a9b6c525552e6762a45ee0f4
sha256checksum	aa490abbe2bf6c45041aea93bdbdcd37d147bbe188ceb85e18ed060b50b85c88
PRs fixed	NA
CVEs	CVE-2024-37079, CVE-2024-37080, CVE-2024-37081

Download and Installation

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

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

1 Attach the `VMware-vCenter-Server-Appliance-7.0.3.02000-24026615-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 the vCenter Server Appliance.

For more information on staging patches, see Stage Patches to vCenter Server Appliance.

For more information on installing patches, see Install vCenter Server Appliance Patches.

For more information on patching using the Appliance Management Interface, see Patching the vCenter Server by Using the Appliance Management Interface.

Resolved Issues

5

Read the following topics next:

- Security Issues

Security Issues

- This release resolves CVE-2024-37079, CVE-2024-37080, and CVE-2024-37081. For more information on these vulnerabilities and their impact on VMware products, see VMSA-2024-0012.

Known Issues

6

Read the following topics next:

- Installation, Upgrade, and Migration Issues

Installation, Upgrade, and Migration Issues

- **Patching vCenter to version 7.0 Update 3q might fail due to unsupported TLS ciphers in the Envoy proxy configuration**

  During the update of a vCenter instance to version 7.0 Update 3q, some services might fail to start due to custom configuration of the TLS cipher suites in the file `/etc/vmware-rhttpproxy/config.xml`. As a result, in the vSphere Client you see a warning such as `An error occurred while starting service 'vpxd-svcs'` and the process fails. If you use the Virtual Appliance Management Interface, you see an error such as `Installation Failed - Exception occurred in postInstallHook`.

  Workaround: See KB 369485 for steps how to reset the TLS ciphers to their defaults and retry the vCenter update.

Known Issues from Prior Releases 7

Read the following topics next:

vSphere Cluster Services Issues

- **If all vSphere Cluster Service agent virtual machines in a cluster are down, vSphere DRS does not function in the cluster**

  If vSphere Cluster Service agent virtual machines fail to deploy or power on in a cluster, services such as vSphere DRS might be impacted.

  Workaround: For more information on the issue and workarounds, see VMware knowledge base article 79892.

- **System virtual machines that support vSphere Cluster Services might impact cluster and datastore maintenance workflows**

In vCenter Server 7.0 Update 1, vSphere Cluster Services adds a set of system virtual machines in every vSphere cluster to ensure the healthy operation of vSphere DRS. The system virtual machines deploy automatically with an implicit datastore selection logic. Depending on your cluster configuration, the system virtual machines might impact some of the cluster and datastore maintenance workflows.

Workaround: For more information on the impacted workflows and possible workarounds, see VMware knowledge base articles 79892 and 80483.

Installation, Upgrade, and Migration Issues

- **Upgrade to vSphere 7.0 Update 3c might require additional steps to force a full host-sync**

  The supported upgrade sequence for vSphere systems is first to upgrade vCenter Server and then ESXi. However, in certain environments with ESXi hosts of version 7.0 Update 2d and later, you need to update ESXi first to 7.0 Update 3c and then vCenter Server. Such an upgrade sequence requires additional steps to force a full host-sync.

  Workaround: Log in to the appliance shell as a user with super administrative privileges (for example, `root`) and follow these steps:

  a Stop the vpxd service.

  b Run the command `/opt/vmware/vpostgres/current/bin/psql -U postgres -d VCDB -c "update VPX_HOST_SYNC_GEN set master_gen=0 where host_id in (select id from VPX_HOST where product_version like '7.0.3%')"`.

  c Start the vpxd service.

- **Hot-patched ESXi hosts with both i40en and i40enu Intel network drivers installed might fail to configure vSphere High Availability after upgrade to vCenter Server 7.0 Update 3c**

  Due to the name change in the Intel i40en driver to i40enu and back to i40en, vCenter Server 7.0 Update 3c adds an upgrade precheck to make sure that ESXi hosts affected from the change are properly upgraded. However, if you apply an ESXi hot patch that is released after vCenter Server 7.0 Update 3c and then upgrade your system to vCenter Server 7.0 Update 3c, the hot patch might not be listed in the precheck. As a result, you might not follow the proper steps to the upgrade and vSphere HA might fail to configure on such hosts.

  Workaround: Upgrade the hot-patched ESXi hosts to version 7.0 Update 3c.

- **The vCenter Upgrade/Migration pre-checks fail with "Unexpected error 87"**

  The vCenter Server Upgrade/Migration pre-checks fail when the Security Token Service (STS) certificate does not contain a Subject Alternative Name (SAN) field. This situation occurs when you have replaced the vCenter 5.5 Single Sign-On certificate with a custom certificate that has no SAN field, and you attempt to upgrade to vCenter Server 7.0. The upgrade considers the STS certificate invalid and the pre-checks prevent the upgrade process from continuing.

Workaround: Replace the STS certificate with a valid certificate that contains a SAN field then proceed with the vCenter Server 7.0 Upgrade/Migration.

- **Problems upgrading to vSphere 7.0 with pre-existing CIM providers**

  After upgrade, previously installed 32-bit CIM providers stop working because ESXi requires 64-bit CIM providers. Customers may lose management API functions related to CIMPDK, NDDK (native DDK), HEXDK, VAIODK (IO filters), and see errors related to **uwglibc** dependency.

  The syslog reports module missing, "32 bit shared libraries not loaded."

  Workaround: There is no workaround. The fix is to download new 64-bit CIM providers from your vendor.

- **Patching to vCenter Server 7.0 Update 1 from earlier versions of vCenter Server 7.x is blocked when vCenter Server High Availability is enabled**

  Patching to vCenter Server 7.0 Update 1 from earlier versions of vCenter Server 7.x is blocked when vCenter Server High Availability is active.

  Workaround: To patch your system to vCenter Server 7.0 Update 1 from earlier versions of vCenter Server 7.x, you must remove vCenter Server High Availability and delete the passive and witness nodes. After the upgrade, you must re-create your vCenter Server High Availability clusters.

- **Migration of a 6.7.x vCenter Server system to vCenter Server 7.x fails with an UnicodeEncodeError**

  If you select the option to import all data for configuration, inventory, tasks, events, and performance metrics, the migration of a 6.7.x vCenter Server system to vCenter Server 7.x might fail for any vCenter Server system that uses a non-English locale. At step 1 of stage 2 of the migration, in the vSphere Client, you see an error such as:

  ```
  Error while exporting events and tasks data: …ERROR UnicodeEncodeError: Traceback (most recent call last):
  ```

  Workaround: You can complete the migration operation by doing either:

  - Select the default option **Configuration and Inventory** at the end of stage 1 of the migration.

    This option does not include tasks and events data.

  - Clean the data in the events tables and run the migration again.

- **If a Windows vCenter Server system has a database password containing non-ASCII characters, pre-checks of the VMware Migration Assistant fail**

If you try to migrate a 6.x vCenter Server system to vCenter Server 7.x by using the VMware Migration Assistant, and your system has a Windows OS, and uses an external database with a password containing non-ASCII characters, the operation fails. For example, Admin!23 迁移. In the Migration Assistant console, you see the following error:

```
Error:Component com.vmware.vcdb failed with internal errorResolution:File Bugzilla PR to VPX/VPX/vcdb-upgrade
```

Workaround: None

- **During an update from vCenter Server 7.x to vCenter Server 7.0 Update 1, you get prompts to provide the vCenter Single Sign-On password**

  During an update from vCenter Server 7.x to vCenter Server 7.0 Update 1, you get prompts to provide vCenter Single Sign-On administrator password.

  Workaround: If you run the update by using the vCenter Server Management Interface, you must provide the vCenter Single Sign-On administrator password.

  If you run the update by using software-packages or CLI in an interactive manner, you must interactively provide the vCenter Single Sign-On administrator password.

  If you run the update by using software-packages or CLI in a non-interactive manner, you must provide the vCenter Single Sign-On administrator password by an answer file in the format

  ```
  { "vmdir.password": "SSO Password of Administrator@<SSO-DOMAIN> user" }
  ```

- **Smart Card and RSA SecurID authentication might stop working after upgrading to vCenter Server 7.0**

  If you have configured vCenter Server for either Smart Card or RSA SecurID authentication, see the VMware knowledge base article at https://kb.vmware.com/s/article/78057 before starting the vSphere 7.0 upgrade process. If you do not perform the workaround as described in the KB, you might see the following error messages and Smart Card or RSA SecurID authentication does not work.

  "Smart card authentication may stop working. Smart card settings may not be preserved, and smart card authentication may stop working."

  or

  "RSA SecurID authentication may stop working. RSA SecurID settings may not be preserved, and RSA SecurID authentication may stop working."

  Workaround: Before upgrading to vSphere 7.0, see the VMware knowledge base article at https://kb.vmware.com/s/article/78057.

- **You might not be able to apply or remove NSX while you add ESXi hosts by using a vSphere Lifecycle Manager image to a cluster with enabled VMware vSphere High Availability**

If you start an operation to apply or remove NSX while adding multiple ESXi hosts by using a vSphere Lifecycle Manager image to a vSphere HA-enabled cluster, the NSX-related operations might fail with an error in the vSphere Client such as:

`vSphere HA agent on some of the hosts on cluster <cluster_name> is neither vSphere HA master agent nor connected to vSphere HA master agent. Verify that the HA configuration is correct.`

The issue occurs because vSphere Lifecycle Manager configures vSphere HA for the ESXi hosts being added to the cluster one at a time. If you run an operation to apply or remove NSX while vSphere HA configure operations are still in progress, NSX operations might queue up between the vSphere HA configure operations for two different ESXi hosts. In such a case, the NSX operation fails with a cluster health check error, because the state of the cluster at that point does not match the expected state that all ESXi hosts have vSphere HA configured and running. The more ESXi hosts you add to a cluster at the same time, the more likely the issue is to occur.

Workaround: Deactivate and enable Sphere HA on the cluster. Proceed with the operations to apply or remove NSX.

- **After an upgrade of a vCenter Server 7.0 system, you cannot see the IP addresses of pods in the vSphere Pod Summary tab of the vSphere Client**

  If you upgrade your vCenter Server 7.0 system to a later version, you can no longer see the IP addresses of pods in the **vSphere Pod Summary** tab of the vSphere Client.

  Workaround: Use the Kubernetes CLI Tools for vSphere to review details of pods:

  a As a prerequisite, copy the pod and namespace names.
    - In the vSphere Client, navigate to **Workload Management** > **Clusters.**
    - Copy the IP displayed in the **Control Plane Node IP Address** tab.
    - You can navigate to `https://<control_plane_node_IP_address>` and download the Kubernetes CLI Tools, `kubectl` and `kubectl-vsphere`.

      Alternatively, follow the steps in Download and Install the Kubernetes CLI Tools for vSphere.

  b Use the CLI plug-in for vSphere to review the pod details.

    1 Log in to the Supervisor cluster by using the command

      ```
      kubectl vsphere login --server=https://<server_adress> --vsphere-username <your user account name> --insecure-skip-tls-verify
      ```

    2 By using the names copied in step 1, run the commands for retrieving the pod details:

      ```
      kubectl config use-context <namespace_name>
      ```

      and

      ```
      kubectl describe pod <pod_name> -n <namespace_name>
      ```

As a result, you can see the IP address in an output similar to:

```
$ kubectl describe pod helloworld -n my-podvm-ns ...

Status: Running

IP: 10.0.0.10

IPs:

IP: 10.0.0.10 ...
```

- **Upgrading a vCenter Server with an external Platform Services Controller from 6.7u3 to 7.0 fails with VMAFD error**

  When you upgrade a vCenter Server deployment using an external Platform Services Controller, you converge the Platform Services Controller into a vCenter Server appliance. If the upgrade fails with the error `install.vmafd.vmdir_vdcpromo_error_21`, the VMAFD firstboot process has failed. The VMAFD firstboot process copies the VMware Directory Service Database (data.mdb) from the source Platform Services Controller and replication partner vCenter Server appliance.

  Workaround: Deactivate TCP Segmentation Offload (TSO) and Generic Segmentation Offload (GSO) on the Ethernet adapter of the source Platform Services Controller or replication partner vCenter Server appliance before upgrading a vCenter Server with an external Platform Services Controller. See Knowledge Base article: https://kb.vmware.com/s/article/74678

- **vCenter Server system upgrades fail in the pre-check stage**

  Upgrades of your vCenter Server system might fail in the pre-check stage due to a limit in the authorization (Authz) connections. In the `/var/log/vmware/vpxd-svcs/vpxd-svcs*.log` file you see entries such as:

  `Session count for user [after add]: <DOMAIN-NAME>\machine-xxxx is 200Session limit reached for user: <DOMAIN-NAME>\machine-xxxx with 200 sessions.`

  You might also see delayed response from the vSphere Client to load the inventory.

  Workaround: Restart vmware-vpxd-svcs in your vCenter Server system by using the command `service-control --restart vmware-vpxd-svcs`. Use the command only when no other activity runs in the vCenter Server system to avoid any interruptions to the workflow. For more information, see VMware knowledge base article 81953.

- **Upgrading vCenter Server using the CLI incorrectly preserves the Transport Security Layer (TLS) configuration for the vSphere Authentication Proxy service**

If the vSphere Authentication Proxy service (`vmcam`) is configured to use a particular TLS protocol other than the default TLS 1.2 protocol, this configuration is preserved during the CLI upgrade process. By default, vSphere supports the TLS 1.2 encryption protocol. If you must use the TLS 1.0 and TLS 1.1 protocols to support products or services that do not support TLS 1.2, use the TLS Configurator Utility to enable or deactivate different TLS protocol versions.

Workaround: Use the TLS Configurator Utility to configure the `vmcam` port. To learn how to manage TLS protocol configuration and use the TLS Configurator Utility, see the *VMware Security* documentation.

- **You do not see a precheck error when patching to vCenter Server 7.0 Update 3c by using CLI**

  Due to the name change in the Intel i40en driver to i40enu and back to i40en, vCenter Server 7.0 Update 3c adds an upgrade precheck to make sure that ESXi hosts affected from the change are properly upgraded. In some cases, if such hosts exist in your system, patching from an vCenter Server version earlier than 7.0 Update 3 to a version later than 7.0 Update 3 by using CLI might fail with the error `Installation failed. Retry to resume from the current state. Or please collect the VC support bundle.`.

  However, instead of this error, you must see the precheck error message.

  Workaround: If you do not see the precheck error and patching your system to vCenter Server 7.0 Update 3c fails, make sure all ESXi hosts are upgraded to ESXi 7.0 Update 3c or higher, by using either a baseline created from an ISO or a single image, before upgrading vCenter Server. Do not use patch baselines based on the rollup bulletin. You can find additional debug log information at `/var/log/vmware/applmgmt`. For more details, see VMware knowledge base articles 87319 and 86447.

- **Smart card and RSA SecurID settings may not be preserved during vCenter Server upgrade**

  Authentication using RSA SecurID will not work after upgrading to vCenter Server 7.0. An error message will alert you to this issue when attempting to login using your RSA SecurID login.

  Workaround: Reconfigure the smart card or RSA SecureID.

- **Upgrade to vSphere 7.0 Update 3c might require additional steps to force a full host-sync**

  The supported upgrade sequence for vSphere systems is first to upgrade vCenter Server and then ESXi. However, in certain environments with ESXi hosts of version 7.0 Update 2c and later, you need to update ESXi first to 7.0 Update 3c and then vCenter Server. Such an upgrade sequence requires additional steps to force a full host-sync.

  Workaround: Log in to the appliance shell as a user with super administrative privileges (for example, `root`) and follow these steps:

  a Stop the vpxd service.

b Run the command `/opt/vmware/vpostgres/current/bin/psql -U postgres -d VCDB -c "update VPX_HOST_SYNC_GEN set master_gen=0 where host_id in (select id from VPX_HOST where product_version like '7.0.3%')"`.

c Start the vpxd service.

- **Migration of vCenter Server for Windows to vCenter Server appliance 7.0 fails with network error message**

  Migration of vCenter Server for Windows to vCenter Server appliance 7.0 fails with the error message `IP already exists in the network`. This prevents the migration process from configuring the network parameters on the new vCenter Server appliance. For more information, examine the log file: `/var/log/vmware/upgrade/UpgradeRunner.log`

  Workaround:

  a Verify that all Windows Updates have been completed on the source vCenter Server for Windows instance, or deactivate automatic Windows Updates until after the migration finishes.

  b Retry the migration of vCenter Server for Windows to vCenter Server appliance 7.0.

- **Hot-patched ESXi hosts with both i40en and i40enu Intel network drivers installed might fail to configure vSphere High Availability after upgrade to vCenter Server 7.0 Update 3c**

  Due to the name change in the Intel i40en driver to i40enu and back to i40en, vCenter Server 7.0 Update 3c adds an upgrade precheck to make sure that ESXi hosts affected from the change are properly upgraded. However, if you apply an ESXi hot patch that is released after vCenter Server 7.0 Update 3c and then upgrade your system to vCenter Server 7.0 Update 3c, the hot patch might not be listed in the precheck. As a result, you might not follow the proper steps to the upgrade and vSphere HA might fail to configure on such hosts.

  Workaround: Upgrade the hot-patched ESXi hosts to version 7.0 Update 3c.

- **When you configure the number of virtual functions for an SR-IOV device by using the max_vfs module parameter, the changes might not take effect**

  In vSphere 7.0, you can configure the number of virtual functions for an SR-IOV device by using the Virtual Infrastructure Management (VIM) API, for example, through the vSphere Client. The task does not require reboot of the ESXi host. After you use the VIM API configuration, if you try to configure the number of SR-IOV virtual functions by using the `max_vfs` module parameter, the changes might not take effect because they are overridden by the VIM API configuration.

  Workaround: None. To configure the number of virtual functions for an SR-IOV device, use the same method every time. Use the VIM API or use the `max_vfs` module parameter and reboot the ESXi host.

- **Upgraded vCenter Server appliance instance does not retain all the secondary networks (NICs) from the source instance**

  During a major upgrade, if the source instance of the vCenter Server appliance is configured with multiple secondary networks other than the VCHA NIC, the target vCenter Server instance will not retain secondary networks other than the VCHA NIC. If the source instance is configured with multiple NICs that are part of VDS port groups, the NIC configuration will not be preserved during the upgrade. Configurations for vCenter Server appliance instances that are part of the standard port group will be preserved.

  Workaround: None. Manually configure the secondary network in the target vCenter Server appliance instance.

- **After upgrading or migrating a vCenter Server with an external Platform Services Controller, users authenticating using Active Directory lose access to the newly upgraded vCenter Server instance**

  After upgrading or migrating a vCenter Server with an external Platform Services Controller, if the newly upgraded vCenter Server is not joined to an Active Directory domain, users authenticating using Active Directory will lose access to the vCenter Server instance.

  Workaround: Verify that the new vCenter Server instance has been joined to an Active Directory domain. See Knowledge Base article: https://kb.vmware.com/s/article/2118543

- **Migrating a vCenter Server for Windows with an external Platform Services Controller using an Oracle database fails**

  If there are non-ASCII strings in the Oracle events and tasks table the migration can fail when exporting events and tasks data. The following error message is provided: UnicodeDecodeError

  Workaround: None.

- **After an ESXi host upgrade, a Host Profile compliance check shows non-compliant status while host remediation tasks fail**

  The non-compliant status indicates an inconsistency between the profile and the host.

  This inconsistency might occur because ESXi 7.0 does not allow duplicate claim rules, but the profile you use contains duplicate rules. For example, if you attempt to use the Host Profile that you extracted from the host before upgrading ESXi 6.5 or ESXi 6.7 to version 7.0 and the Host Profile contains any duplicate claim rules of system default rules, you might experience the problems.

  Workaround:

  a Remove any duplicate claim rules of the system default rules from the Host Profile document.

  b Check the compliance status.

  c Remediate the host.

d If the previous steps do not help, reboot the host.

- **Error message displays in the vCenter Server Management Interface**

  After installing or upgrading to vCenter Server 7.0, when you navigate to the Update panel within the vCenter Server Management Interface, the error message "Check the URL and try again" displays. The error message does not prevent you from using the functions within the Update panel, and you can view, stage, and install any available updates.

  Workaround: None.

- **Patching witness or passive nodes of environments with VMware vCenter Server High Availability enabled might fail**

  In environments with vCenter Server High Availability enabled, patching a witness or passive node might fail with a message similar to:

  `RuntimeError: unidentifiable C++ exception.`

  Workaround: Deactivate vCenter Server High Availability. Apply patches to your vCenter Server system. Re-enable vCenter Server High Availability.

- **After patching your vCenter Server system to vCenter Server 7.0.0a, the TLS version of the VC Storage Clients might revert to the default**

  If you have a TLS configuration for the VC Storage Clients service different from the default TLS 1.2 only, the TLS version might revert to the default after patching your vCenter Server system to vCenter Server 7.0.0a.

  Workaround: Use the TLS Configuration utility to enable or deactivate TLS versions on your vCenter Server system after the update.

- **After updating your system to vCenter Server 7.0.0b, you see systemd core dump in the /var/core folder**

  After updating your system to vCenter Server 7.0.0b from either vCenter Server 7.0.0a or vCenter Server 7.0, in the `/var/core` folder you see systemd core dump, such as `core.systemd-journal.393` and `core.systemd-udevd.405`. The core dump is harmless and can be removed.

  Workaround: None

- **After updating your vCenter Server system to 7.0.0b, the vCenter Server version is not updated in the Direct Console User Interface (DCUI)**

  After updating your system to vCenter Server 7.0.0b from vCenter Server 7.0.0a or vCenter Server 7.0, you still see the previous vCenter Server version in the DCUI.

  Workaround: After you complete the update, to refresh the vCenter Server version, in the appliance shell, run the command `/usr/lib/applmgmt/dcui/notify`.

- **Update Planner fails with error Configured repository is not accessible due to network connectivity or incorrect URL**

  If you use Update Planner, which is part of vSphere Lifecycle Manager used to facilitate vCenter Server updates, you might see the following error in the vSphere Client:

  `Configured repository is not accessible due to network connectivity or incorrect URL. Verify the repository settings.`

  The issue occurs when you use a custom local repository, such as `https:///uploads/dpe/` or a DBC path, to store the extracted . If the custom repository for URL-based patching has an authentication policy, Update Planner might not be able to fetch the list of available updates.

  Workaround: Configure the custom repository such that authentication is not needed to access the custom repository URL.

- **After upgrading to vCenter Server 7.0.0b, you see vSphere HA errors on vSphere Lifecycle Manager Image based clusters**

  After upgrading to vCenter Server 7.0.0b, on vSphere Lifecycle Manager Image based clusters that are configured with vSphere HA, you might see error messages about the vSphere HA configuration after logging in for the first time to the environment. In the vSphere Client, you see messages such as:

  `Cannot complete the configuration of the vSphere HA agent on the host.` or

  `Applying HA VIBs on the cluster encountered a failure.`

  The issue occurs because exports of the image depot might take long and cause a timeout of the task. In the`/storage/log/vmware/vmware-updatemgr/vum-server/vmware-vum-server.log` you see this message: `Export taking too long (Failure case)`

  Workaround: This is a transient issue that resolves in 10 minutes after the vCenter Server is up and running. The issue does not affect any functionality. vSphere HA on the affected clusters operates as expected. All operations related to virtual machines, such as power on and migration, work across the vSphere HA-enabled clusters while this error recovery is still in progress.

- **You do not see a precheck error when patching to vCenter Server 7.0 Update 3d by using CLI**

  Due to the name change in the Intel i40en driver to i40enu and back to i40en, vCenter Server 7.0 Update 3d and later add an upgrade precheck to make sure that ESXi hosts affected from the change are properly upgraded. In some cases, if such hosts exist in your system, patching from a vCenter Server version earlier than 7.0 Update 3 to a version later than 7.0 Update 3 by using CLI might fail with the error `Installation failed. Retry to resume from the current state. Or please collect the VC support bundle.`. However, instead of this error, you must see the precheck error message.

Workaround: If you do not see the precheck error and patching your system to vCenter Server 7.0 Update 3d fails, make sure all ESXi hosts are upgraded to ESXi 7.0 Update 3d, by using either a baseline created from an ISO or a single image, before upgrading vCenter Server. Do not use patch baselines based on the rollup bulletin. You can find additional debug log information at `/var/log/vmware/applmgmt`. For more details, see VMware knowledge base articles 87319 and 86447.

Security Features Issues

- **Encrypted virtual machine fails to power on when HA-enabled Trusted Cluster contains an unattested host**

  In VMware® vSphere Trust Authority™, if you have enabled HA on the Trusted Cluster and one or more hosts in the cluster fails attestation, an encrypted virtual machine cannot power on.

  Workaround: Either remove or remediate all hosts that failed attestation from the Trusted Cluster.

- **Encrypted virtual machine fails to power on when DRS-enabled Trusted Cluster contains an unattested host**

  In VMware® vSphere Trust Authority™, if you have enabled DRS on the Trusted Cluster and one or more hosts in the cluster fails attestation, DRS might try to power on an encrypted virtual machine on an unattested host in the cluster. This operation puts the virtual machine in a locked state.

  Workaround: Either remove or remediate all hosts that failed attestation from the Trusted Cluster.

- **Migrating or cloning encrypted virtual machines across vCenter Server instances fails when attempting to do so using the vSphere Client**

  If you try to migrate or clone an encrypted virtual machine across vCenter Server instances using the vSphere Client, the operation fails with the following error message: "The operation is not allowed in the current state."

  Workaround: You must use the vSphere APIs to migrate or clone encrypted virtual machines across vCenter Server instances.

Networking Issues

- **PR3304411: After batch cloning and customizing Windows virtual machines, the VMs network might not be configured**

When cloning Windows VMs in bulk, a MAC address conflict might occur, causing vCenter to reassign new MAC addresses to the VMs. As a result, the guest customization process fails to configure properly the VMs network because it cannot match the original MAC address to the corresponding network card.

This issue is resolved in this release.

- **When you use the VMware Remote Console, the envoy service might intermittently fail**

  An issue with the envoy service specific to the VMware Remote Console might lead to intermittent failures of the service. As a result, the vCenter Server Management Interface or vCenter Server APIs might also become unavailable.

  Workaround: Use the vSphere Client as an alternative to the VMware Remote Console.

- **Reduced throughput in networking performance on Intel 82599/X540/X550 NICs**

  The new queue-pair feature added to ixgben driver to improve networking performance on Intel 82599EB/X540/X550 series NICs might reduce throughput under some workloads in vSphere 7.0 as compared to vSphere 6.7.

  Workaround: To achieve the same networking performance as vSphere 6.7, you can deactivate the queue-pair with a module parameter. To deactivate the queue-pair, run the command:

  ```
  # esxcli system module parameters set -p "QPair=0,0,0,0..." -m ixgben
  ```

  After running the command, reboot.

- **If you try to deactivate vSphere with Tanzu on a vSphere cluster, the operation stops with an error**

  If some virtual machines outside of a Supervisor Cluster reside on any of the NSX segment port groups on the cluster, the cleanup script cannot delete such ports and deactivate vSphere with Tanzu on the cluster. In the vSphere Client, you see the error `Cleanup requests to NSX Manager failed` and the operation stops at `Removing` status. In the `/var/log/vmware/wcp/wcpsvc.log` file, you see an error message such as

  ```
  Segment path=[...] has x VMs or VIFs attached. Disconnect all VMs and VIFs before deleting a segment.
  ```

  Workaround: Delete the virtual machines indicated in the `/var/log/vmware/wcp/wcpsvc.log` file from the segment. Wait for the operation to restore.

- **After upgrading to NSX 6.4.7, when a static IPv6 address is assigned to workload VMs on an IPv6 network, the VMs are unable to ping the IPv6 gateway interface of the edge**

  This issue occurs after upgrading the vSphere Distributed Switches from 6.x to 7.0.

  Workaround 1:

  Select the VDS where all the hosts are connected, go to the **Edit** setting, and under **Multicast** option switch to basic.

Workaround 2:

Add the following rules on the edge firewall:

Ping allow rule.

Multicast Listener Discover (MLD) allow rule, which are icmp6, type 130 (v1) and type 143 (v2).

- **High throughput virtual machines may experience degradation in network performance when Network I/O Control (NetIOC) is enabled**

  Virtual machines requiring high network throughput can experience throughput degradation when upgrading from vSphere 6.7 to vSphere 7.0 with NetIOC enabled.

  Workaround: Adjust the `ethernetx.ctxPerDev` setting to enable multiple worlds.

- **IPv6 traffic fails to pass through VMkernel ports using IPsec**

  When you migrate VMkernel ports from one port group to another, IPv6 traffic does not pass through VMkernel ports using IPsec.

  Workaround: Remove the IPsec security association (SA) from the affected server, and then reapply the SA. To learn how to set and remove an IPsec SA, see the *vSphere Security* documentation.

- **Higher ESX network performance with a portion of CPU usage increase**

  ESX network performance may increase with a portion of CPU usage.

  Workaround: Remove and add the network interface with only 1 rx dispatch queue. For example:

  ```
  esxcli network ip interface remove --interface-name=vmk1
  ```

  ```
  esxcli network ip interface add --interface-name=vmk1 --num-rxqueue=1
  ```

- **VM might lose Ethernet traffic after hot-add, hot-remove or storage vMotion**

  A VM might stop receiving Ethernet traffic after a hot-add, hot-remove or storage vMotion. This issue affects VMs where the uplink of the VNIC has SR-IOV enabled. PVRDMA virtual NIC exhibits this issue when the uplink of the virtual network is a Mellanox RDMA capable NIC and RDMA namespaces are configured.

  Workaround: You can hot-remove and hot-add the affected Ethernet NICs of the VM to restore traffic. On Linux guest operating systems, restarting the network might also resolve the issue. If these workarounds have no effect, you can reboot the VM to restore network connectivity.

- **Change of IP address for a VCSA deployed with static IP address requires that you create the DNS records in advance**

With the introduction of the DDNS, the DNS record update only works for VCSA deployed with DHCP configured networking. While changing the IP address of the vCenter server via VAMI, the following error is displayed:

The specified IP address does not resolve to the specified hostname.

Workaround: There are two possible workarounds.

a Create an additional DNS entry with the same FQDN and desired IP address. Log in to the VAMI and follow the steps to change the IP address.

b Log in to the VCSA using ssh. Execute the following script:

```
./opt/vmware/share/vami/vami_config_net
```

Use option 6 to change the IP adddress of eth0. Once changed, execute the following script:

```
./opt/likewise/bin/lw-update-dns
```

Restart all the services on the VCSA to update the IP information on the DNS server.

- **It may take several seconds for the NSX Distributed Virtual Port Group (NSX DVPG) to be removed after deleting the corresponding logical switch in NSX Manager.**

  As the number of logical switches increases, it may take more time for the NSX DVPG in vCenter Server to be removed after deleting the corresponding logical switch in NSX Manager. In an environment with 12000 logical switches, it takes approximately 10 seconds for an NSX DVPG to be deleted from vCenter Server.

  Workaround: None.

- **Hostd runs out of memory and fails if a large number of NSX Distributed Virtual port groups are created.**

  In vSphere 7.0, NSX Distributed Virtual port groups consume significantly larger amounts of memory than opaque networks. For this reason, NSX Distributed Virtual port groups can not support the same scale as an opaque network given the same amount of memory.

  Workaround:To support the use of NSX Distributed Virtual port groups, increase the amount of memory in your ESXi hosts. If you verify that your system has adequate memory to support your VMs, you can directly increase the memory of `hostd` using the following command.

  ```
  localcli --plugin-dir /usr/lib/vmware/esxcli/int/ sched group setmemconfig --group-path host/vim/vmvisor/hostd --units mb --min 2048 --max 2048
  ```

  Note that this will cause `hostd` to use memory normally reserved for your environment's VMs. This may have the affect of reducing the number of VMs your ESXi host can support.

- **DRS may incorrectly launch vMotion if the network reservation is configured on a VM**

If the network reservation is configured on a VM, it is expected that DRS only migrates the VM to a host that meets the specified requirements. In a cluster with NSX transport nodes, if some of the transport nodes join the transport zone by NSX-T Virtual Distributed Switch (N-VDS), and others by vSphere Distributed Switch (VDS) 7.0, DRS may incorrectly launch vMotion. You might encounter this issue when:

- The VM connects to an NSX logical switch configured with a network reservation.
- Some transport nodes join transport zone using N-VDS, and others by VDS 7.0, or, transport nodes join the transport zone through different VDS 7.0 instances.

Workaround: Make all transport nodes join the transport zone by N-VDS or the same VDS 7.0 instance.

- **When adding a VMkernel NIC (vmknic) to an NSX portgroup, vCenter Server reports the error "Connecting VMKernel adapter to a NSX Portgroup on a Stateless host is not a supported operation. Please use Distributed Port Group instead."**
  - For stateless ESXi on Distributed Virtual Switch (VDS), the vmknic on a NSX port group is blocked. You must instead use a Distributed Port Group.
  - For stateful ESXi on VDS, vmknic on NSX port group is supported, but vSAN may have an issue if it is using vmknic on a NSX port group.

  Workaround: Use a Distributed Port Group on the same VDS.

- **Enabling SRIOV from vCenter for QLogic 4x10GE QL41164HFCU CNA might fail**

  If you navigate to the **Edit Settings** dialog for physical network adapters and attempt to enable SR-IOV, the operation might fail when using QLogic 4x10GE QL41164HFCU CNA. Attempting to enable SR-IOV might lead to a network outage of the ESXi host.

  Workaround: Use the following command on the ESXi host to enable SRIOV:

  ```
  esxcfg-module
  ```

Storage Issues

- **VMFS datastores are not mounted automatically after disk hot remove and hot insert on HPE Gen10 servers with SmartPQI controllers**

  When SATA disks on HPE Gen10 servers with SmartPQI controllers without expanders are hot removed and hot inserted back to a different disk bay of the same machine, or when multiple disks are hot removed and hot inserted back in a different order, sometimes a new local name is assigned to the disk. The VMFS datastore on that disk appears as a snapshot and will not be mounted back automatically because the device name has changed.

  Workaround: None. SmartPQI controller does not support unordered hot remove and hot insert operations.

- **ESXi might terminate I/O to NVMeOF devices due to errors on all active paths**

  Occasionally, all active paths to NVMeOF device register I/O errors due to link issues or controller state. If the status of one of the paths changes to Dead, the High Performance Plug-in (HPP) might not select another path if it shows high volume of errors. As a result, the I/O fails.

  Workaround: Deactivate the configuration option **/Misc/HppManageDegradedPaths** to unblock the I/O.

- **VOMA check on NVMe based VMFS datastores fails with error**

  VOMA check is not supported for NVMe based VMFS datastores and will fail with the error:

  ```
  ERROR: Failed to reserve device. Function not implemented
  ```

  Example:

  ```
  # voma -m vmfs -f check -d /vmfs/devices/disks/: <partition#>
  Running VMFS Checker version 2.1 in check mode
  Initializing LVM metadata, Basic Checks will be done

  Checking for filesystem activity
  Performing filesystem liveness check..|Scanning for VMFS-6 host activity (4096 bytes/HB,
  1024 HBs).
  ERROR: Failed to reserve device. Function not implemented
  Aborting VMFS volume check
  VOMA failed to check device : General Error
  ```

  Workaround: None. If you need to analyse VMFS metadata, collect it using the `-l` option, and pass to VMware customer support. The command for collecting the dump is:

  ```
  voma -l -f dump -d /vmfs/devices/disks/:<partition#>
  ```

- **Using the VM reconfigure API to attach an encrypted First Class Disk to an encrypted virtual machine might fail with error**

  If an FCD and a VM are encrypted with different crypto keys, your attempts to attach the encrypted FCD to the encrypted VM using the `VM reconfigure API` might fail with the error message:

  ```
  Cannot decrypt disk because key or password is incorrect.
  ```

  Workaround: Use the `attachDisk API` rather than the `VM reconfigure API` to attach an encrypted FCD to an encrypted VM.

- **ESXi host might get in non responding state if a non-head extent of its spanned VMFS datastore enters the Permanent Device Loss (PDL) state**

This problem does not occur when a non-head extent of the spanned VMFS datastore fails along with the head extent. In this case, the entire datastore becomes inaccessible and no longer allows I/Os.

In contrast, when only a non-head extent fails, but the head extent remains accessible, the datastore heartbeat appears to be normal. And the I/Os between the host and the datastore continue. However, any I/Os that depend on the failed non-head extent start failing as well. Other I/O transactions might accumulate while waiting for the failing I/Os to resolve, and cause the host to enter the non responding state.

Workaround: Fix the PDL condition of the non-head extent to resolve this issue.

- **After recovering from APD or PDL conditions, VMFS datastore with enabled support for clustered virtual disks might remain inaccessible**

  You can encounter this problem only on datastores where the clustered virtual disk support is enabled. When the datastore recovers from an All Paths Down (APD) or Permanent Device Loss (PDL) condition, it remains inaccessible. The VMkernel log might show multiple `SCSI3 reservation conflict` messages similar to the following:

  ```
  2020-02-18T07:41:10.273Z cpu22:1001391219)ScsiDeviceIO: vm 1001391219:
  SCSIDeviceCmdCompleteCB:2972: Reservation conflict retries 544 for command
  0x45ba814b8340 (op: 0x89) to device "naa.624a9370b97601e346f64ba900024d53"
  ```

  The problem can occur because the ESXi host participating in the cluster loses SCSI reservations for the datastore and cannot always reacquire them automatically after the datastore recovers.

  Workaround: Manually register the reservation using the following command:

  ```
  vmkfstools -L registerkey /vmfs/devices/disks/<device name>
  ```

  where the `<device name>` is the name of the device on which the datastore is created.

- **Virtual NVMe Controller is the default disk controller for Windows 10 guest operating systems**

  The Virtual NVMe Controller is the default disk controller for the following guest operating systems when using Hardware Version 15 or later:

  Windows 10

  Windows Server 2016

  Windows Server 2019

  Some features might not be available when using a Virtual NVMe Controller. For more information, see https://kb.vmware.com/s/article/2147714

  Note: Some clients use the previous default of LSI Logic SAS. This includes ESXi host client and PowerCLI.

Workaround: If you need features not available on Virtual NVMe, switch to VMware Paravirtual SCSI (PVSCSI) or LSI Logic SAS. For information on using VMware Paravirtual SCSI (PVSCSI), see https://kb.vmware.com/s/article/1010398

- **After an ESXi host upgrade to vSphere 7.0, presence of duplicate core claim rules might cause unexpected behavior**

  Claim rules determine which multipathing plugin, such as NMP, HPP, and so on, owns paths to a particular storage device. ESXi 7.0 does not support duplicate claim rules. However, the ESXi 7.0 host does not alert you if you add duplicate rules to the existing claim rules inherited through an upgrade from a legacy release. As a result of using duplicate rules, storage devices might be claimed by unintended plugins, which can cause unexpected outcome.

  Workaround: Do not use duplicate core claim rules. Before adding a new claim rule, delete any existing matching claim rule.

- **A CNS query with the compliance status filter set might take unusually long time to complete**

  The CNS QueryVolume API enables you to obtain information about the CNS volumes, such as volume health and compliance status. When you check the compliance status of individual volumes, the results are obtained quickly. However, when you invoke the CNS QueryVolume API to check the compliance status of multiple volumes, several tens or hundreds, the query might perform slowly.

  Workaround: Avoid using bulk queries. When you need to get compliance status, query one volume at a time or limit the number of volumes in the query API to 20 or fewer. While using the query, avoid running other CNS operations to get the best performance.

- **New Deleted CNS volumes might temporarily appear as existing in the CNS UI**

  After you delete an FCD disk that backs a CNS volume, the volume might still show up as existing in the CNS UI. However, your attempts to delete the volume fail. You might see an error message similar to the following:

  ```
  The object or item referred to could not be found.
  ```

  Workaround: The next full synchronization will resolve the inconsistency and correctly update the CNS UI.

- **New Under certain circumstances, while a CNS operation fails, the task status appears as successful in the vSphere Client**

  This might occur when, for example, you use an incompliant storage policy to create a CNS volume. The operation fails, while the vSphere Client shows the task status as successful.

  Workaround: The successful task status in the vSphere Client does not guarantee that the CNS operation succeeded. To make sure the operation succeeded, verify its results.

- **New Unsuccessful delete operation for a CNS persistent volume might leave the volume undeleted on the vSphere datastore**

  This issue might occur when the CNS Delete API attempts to delete a persistent volume that is still attached to a pod. For example, when you delete the Kubernetes namespace where the pod runs. As a result, the volume gets cleared from CNS and the CNS query operation does not return the volume. However, the volume continues to reside on the datastore and cannot be deleted through the repeated CNS Delete API operations.

  Workaround: None.

vCenter Server and vSphere Client Issues

- **PR 3262556: vCenter upgrade fails with "VP last sync time not updated after waiting for 5 minutes after vCenter update"**

  In rare cases, during a vCenter upgrade, a race condition between vSphere API for Storage Awareness (VASA) providers might cause the upgrade to fail. In the `pod-service.log` file, you see an error such as `vSAN VP last sync time not updated after waiting for 5 minutes after vCenter update`.

  This issue is resolved in this release. The fix prevents the race condition.

- **You do not see the VLAN ID in the port properties page of a VMkernel network adapter in the vSphere Client**

  In the vSphere Client, when you follow the path **Host > Add Networking**, after you select a VMkernel network adapter and an existing standard switch, in the **Port properties** page you might not see the VLAD ID. The issue occurs because the VLAN ID component is part of the `network-ui` module, which is lazy loaded and needs a trigger to refresh.

  Workaround: Go to **Host > Configure > Networking > Virtual Switches** to trigger the loading of the `network ui` module and the VLAD ID appears.

- **Vendor providers go offline after a PNID change**

  When you change the vCenter IP address (PNID change), the registered vendor providers go offline.

  Workaround: Re-register the vendor providers.

- **Cross vCenter migration of a virtual machine fails with an error**

  When you use cross vCenter vMotion to move a VM's storage and host to a different vCenter server instance, you might receive the error `The operation is not allowed in the current state.`

  This error appears in the UI wizard after the Host Selection step and before the Datastore Selection step, in cases where the VM has an assigned storage policy containing host-based rules such as encryption or any other IO filter rule.

Workaround: Assign the VM and its disks to a storage policy without host-based rules. You might need to decrypt the VM if the source VM is encrypted. Then retry the cross vCenter vMotion action.

- **Storage Sensors information in Hardware Health tab shows incorrect values on vCenter UI, host UI, and MOB**

  When you navigate to **Host > Monitor > Hardware Health > Storage Sensors** on vCenter UI, the storage information displays either incorrect or unknown values. The same issue is observed on the host UI and the MOB path `"runtime.hardwareStatusInfo.storageStatusInfo"` as well.

  Workaround: None.

- **vSphere UI host advanced settings shows the current product locker location as empty with an empty default**

  vSphere UI host advanced settings shows the current product locker location as empty with an empty default. This is inconsistent as the actual product location `symlink` is created and valid. This causes confusion to the user. The default cannot be corrected from UI.

  Workaround: User can use the esxcli command on the host to correct the current product locker location default as below.

  1. Remove the existing Product Locker Location setting with: `"esxcli system settings advanced remove -o ProductLockerLocation"`

  2. Re-add the Product Locker Location setting with the appropriate default:

  2.a. If the ESXi is a full installation, the default value is `"/locker/packages/vmtoolsRepo" export PRODUCT_LOCKER_DEFAULT="/locker/packages/vmtoolsRepo"`

  2.b. If the ESXi is a PXEboot configuration such as autodeploy, the default value is: `"/vmtoolsRepo" export PRODUCT_LOCKER_DEFAULT="/vmtoolsRepo"`

  Run the following command to automatically figure out the location: `export PRODUCT_LOCKER_DEFAULT=`readlink /productLocker``

  Add the setting: `esxcli system settings advanced add -d "Path to VMware Tools repository" -o ProductLockerLocation -t string -s $PRODUCT_LOCKER_DEFAULT`

  You can combine all the above steps in step 2 by issuing the single command:

  ```
  esxcli system settings advanced add -d "Path to VMware Tools repository" -o ProductLockerLocation -t string -s `readlink /productLocker`
  ```

- **Skyline Health page displays garbage characters**

  In the vSphere Client, when you navigate to vCenter Server or select an ESXi host in the vSphere Client navigator and click Monitor > Skyline Health, the page displays garbage characters in the following locales: Korean, Japanese, German and French.

Workaround: Switch to English locale.

Virtual Machine Management Issues

- **You cannot add or modify an existing network adapter on a virtual machine**

  If you try to add or modify an existing network adapter on a virtual machine, the Reconfigure Virtual Machine task might fail with an error such as `Cannot complete operation due to concurrent modification by another operation` in the vSphere Client. In the `/var/log/hostd.log` file of the ESXi host where the virtual machine runs, you see logs such as:

  ```
  2020-07-28T07:47:31.621Z verbose hostd[2102259] [Originator@6876 sub=Vigor.Vmsvc.vm:/vmfs/volumes/vsan:526bc94351cf8f42-41153841cab2f9d9/bad71f5f-d85e-a276-4cf6-246e965d7154/interop_l2vpn_vmotion_VM_1.vmx] NIC: connection control message: Failed to connect virtual device 'ethernet0'.
  ```

  In the `vpxa.log` file, you see entries similar to: `2020-07-28T07:47:31.941Z info vpxa[2101759] [Originator@6876 sub=Default opID=opId-59f15-19829-91-01-ed] [VpxLRO] -- ERROR task-138 -- vm-13 -- vim.VirtualMachine.reconfigure: vim.fault.GenericVmConfigFault:`

  Workaround: For each ESXi host in your cluster do the following:

  a Connect to the ESXi host by using SSH and run the command

    ```
    esxcli system module parameters set -a -p dvfiltersMaxFilters=8192 -m dvfilter
    ```

  b Put the ESXi host in Maintenance Mode.

  c Reboot the ESXi host.

  For more information, see VMware knowledge base article 80399.

- **ESXi 6.5 hosts with AMD Opteron Generation 3 (Greyhound) processors cannot join Enhanced vMotion Compatibility (EVC) AMD REV E or AMD REV F clusters on a vCenter Server 7.0 Update 1 system**

  In vCenter Server 7.0 Update 1, vSphere cluster services, such as vSphere DRS and vSphere HA, run on ESX agent virtual machines to make the services functionally independent of vCenter Server. However, the CPU baseline for AMD processors of the ESX agent virtual machines have POPCNT SSE4A instructions, which prevents ESXi 6.5 hosts with AMD Opteron Generation 3 (Greyhound) processors to enable EVC mode AMD REV E and AMD REV F on a vCenter Server 7.0 Update 1 system.

  Workaround: None

- **The postcustomization section of the customization script runs before the guest customization**

When you run the guest customization script for a Linux guest operating system, the `precustomization` section of the customization script that is defined in the customization specification runs before the guest customization and the `postcustomization` section runs after that. If you enable Cloud-Init in the guest operating system of a virtual machine, the `postcustomization` section runs before the customization due to a known issue in Cloud-Init.

Workaround: Deactivate Cloud-Init and use the standard guest customization.

- **Group migration operations in vSphere vMotion, Storage vMotion, and vMotion without shared storage fail with error**

  When you perform group migration operations on VMs with multiple disks and multi-level snapshots, the operations might fail with the error `com.vmware.vc.GenericVmConfigFault Failed waiting for data. Error 195887167. Connection closed by remote host, possibly due to timeout.`

  Workaround: Retry the migration operation on the failed VMs one at a time.

- **Deploying an OVF or OVA template from a URL fails with a 403 Forbidden error**

  The URLs that contain an HTTP query parameter are not supported. For example, `http://webaddress.com?file=abc.ovf` or the Amazon pre-signed S3 URLs.

  Workaround: Download the files and deploy them from your local file system.

- **Importing or deploying local OVF files containing non-ASCII characters in their name might fail with an error**

  When you import local `.ovf` files containing non-ASCII characters in their name, you might receive `400 Bad Request Error`. When you use such `.ovf` files to deploy a virtual machine in the vSphere Client, the deployment process stops at 0%. As a result, you might receive `400 Bad Request Error` or `500 Internal Server Error`.

  Workaround:

  a Remove the non-ASCII characters from the `.ovf` and `.vmdk` file names.
    - To edit the `.ovf` file, open it with a text editor.
    - Search the non-ASCII `.vmdk` file name and change it to ASCII.

  b Import or deploy the saved files again.

- **New The third level of nested objects in a virtual machine folder is not visible**

  Perform the following steps:

  a Navigate to a data center and create a virtual machine folder.

  b In the virtual machine folder, create a nested virtual machine folder.

c In the second folder, create another nested virtual machine, virtual machine folder, vApp, or VM Template.

As a result, from the VMs and Templates inventory tree you cannot see the objects in the third nested folder.

Workaround: To see the objects in the third nested folder, navigate to the second nested folder and select the VMs tab.

- **vSAN file services operations fail on vSphere Lifecycle Manager-enabled clusters**

  During a change in the state of an ESXi host, vSAN file services operations might fail on vSphere Lifecycle Manager-enabled clusters due to a race condition with the vSphere ESX Agent Manager (EAM). The problem happens during upgrades and operations, such as power on or power off, booting, or when the host exits maintenance or standby mode. The race condition occurs when an endpoint has been unavailable before the change of state of the ESXi host. In such cases, the EAM starts a remediation process that cannot be resolved and fails operations from other services, such as the vSAN file services.

  Workaround: Restart the vSphere ESX Agent Manager.

vSphere HA and Fault Tolerance Issues

- **VMs in a cluster might be orphaned after recovering from storage inaccessibility such as a cluster wide APD**

  Some VMs might be in orphaned state after cluster wide APD recovers, even if HA and VMCP are enabled on the cluster.

  This issue might be encountered when the following conditions occur simultaneously:

  - All hosts in the cluster experience APD and do not recover until VMCP timeout is reached.
  - HA primary initiates failover due to APD on a host.
  - Power on API during HA failover fails due to one of the following:
    - APD across the same host
    - Cascading APD across the entire cluster
    - Storage issues
    - Resource unavailability
  - FDM unregistration and VCs steal VM logic might initiate during a window where FDM has not unregistered the failed VM and VC's host synchronization responds that multiple hosts are reporting the same VM. Both FDM and VC unregister the different registered copies of the same VM from different hosts, causing the VM to be orphaned.

  Workaround: You must unregister and reregister the orphaned VMs manually within the cluster after the APD recovers.

If you do not manually reregister the orphaned VMs, HA attempts failover of the orphaned VMs, but it might take between 5 to 10 hours depending on when APD recovers.

The overall functionality of the cluster is not affected in these cases and HA will continue to protect the VMs. This is an anomaly in what gets displayed on VC for the duration of the problem.

vSphere Lifecycle Manager Issues

- **vSphere Lifecycle Manager and vSAN File Services cannot be simultaneously enabled on a vSAN cluster in vSphere 7.0 release**

  If vSphere Lifecycle Manager is enabled on a cluster, vSAN File Services cannot be enabled on the same cluster and vice versa. In order to enable vSphere Lifecycle Manager on a cluster, which has VSAN File Services enabled already, first deactivate vSAN File Services and retry the operation. Please note that if you transition to a cluster that is managed by a single image, vSphere Lifecycle Manager cannot be deactivated on that cluster.

  Workaround: None.

- **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.

- **When a hardware support manager is unavailable, vSphere High Availability (HA) functionality is impacted**

  If hardware support manager is unavailable for a cluster that you manage with a single image, where a firmware and drivers addon is selected and vSphere HA is enabled, the vSphere HA functionality is impacted. You may experience the following errors.

  - Configuring vSphere HA on a cluster fails.
  - Cannot complete the configuration of the vSphere HA agent on a host: `Applying HA VIBs on the cluster encountered a failure.`
  - Remediating vSphere HA fails: `A general system error occurred: Failed to get Effective Component map.`
  - Disabling vSphere HA fails: Delete Solution task failed. `A general system error occurred: Cannot find hardware support package from depot or hardware support manager.`

Workaround:

- If the hardware support manager is temporarily unavailable, perform the following steps.

a Reconnect the hardware support manager to vCenter Server.

b Select a cluster from the Hosts and Cluster menu.

c Select the Configure tab.

d Under Services, click vSphere Availability.

e Re-enable vSphere HA.

- If the hardware support manager is permanently unavailable, perform the following steps.

a Remove the hardware support manager and the hardware support package from the image specification

b Re-enable vSphere HA.

c Select a cluster from the Hosts and Cluster menu.

d Select the Updates tab.

e Click Edit .

f Remove the firmware and drivers addon and click Save.

g Select the Configure tab.

h Under Services, click vSphere Availability.

i Re-enable vSphere HA.

- **I/OFilter is not removed from a cluster after a remediation process in vSphere Lifecycle Manager**

  Removing I/OFilter from a cluster by remediating the cluster in vSphere Lifecycle Manager, fails with the following error message: `iofilter XXX already exists`. The iofilter remains listed as installed.

  Workaround:

  a Call IOFilter API `UninstallIoFilter_Task` from the vCenter Server managed object (IoFilterManager).

  b Remediate the cluster in vSphere Lifecycle Manager.

  c Call IOFilter API `ResolveInstallationErrorsOnCluster_Task` from the vCenter Server managed object (IoFilterManager) to update the database.

- **While remediating a vSphere HA enabled cluster in vSphere Lifecycle Manager, disabling and re-enabling vSphere HA causes a vSphere HA error state**

Disabling and re-enabling vSphere HA during remediation process of a cluster, may fail the remediation process due to vSphere HA health checks reporting that hosts don't have vSphere HA VIBs installed. You may see the following error message: `Setting desired image spec for cluster failed`.

Workaround: After the cluster remediation operation has finished, deactivate and re-enable vSphere HA for the cluster.

- **Checking for recommended images in vSphere Lifecycle Manager has slow performance in large clusters**

  In large clusters with more than 16 hosts, the recommendation generation task could take more than an hour to finish or may appear to hang. The completion time for the recommendation task depends on the number of devices configured on each host and the number of image candidates from the depot that vSphere Lifecycle Manager needs to process before obtaining a valid image to recommend.

  Workaround: None.

- **Checking for hardware compatibility in vSphere Lifecycle Manager has slow performance in large clusters**

  In large clusters with more than 16 hosts, the validation report generation task could take up to 30 minutes to finish or may appear to hang. The completion time depends on the number of devices configured on each host and the number of hosts configured in the cluster.

  Workaround: None

- **Incomplete error messages in non-English languages are displayed, when remediating a cluster in vSphere Lifecycle Manager**

  You can encounter incomplete error messages for localized languages in the vCenter Server user interface. The messages are displayed, after a cluster remediation process in vSphere Lifecycle Manager fails. For example, your can observe the following error message.

  The error message in English language: `Virtual machine 'VMC on DELL EMC -FileServer' that runs on cluster 'Cluster-1' reported an issue which prevents entering maintenance mode: Unable to access the virtual machine configuration: Unable to access file[local-0] VMC on Dell EMC - FileServer/VMC on Dell EMC - FileServer.vmx`

  The error message in French language: `La VM « VMC on DELL EMC -FileServer », située sur le cluster « {Cluster-1} », a signalé un problème empêchant le passage en mode de maintenance : Unable to access the virtual machine configuration: Unable to access file[local-0] VMC on Dell EMC - FileServer/VMC on Dell EMC - FileServer.vmx`

  Workaround: None.

- **When you convert a cluster that uses baselines to a cluster that uses a single image, a warning is displayed that vSphere HA VIBs will be removed**

Converting a vSphere HA enabled cluster that uses baselines to a cluster that uses a single image, may result a warning message displaying that `vmware-fdm` component will be removed.

Workaround: This message can be ignored. The conversion process installs the `vmware-fdm` component.

- **If vSphere Update Manager is configured to download patch updates from the Internet through a proxy server, after upgrade to vSphere 7.0 that converts Update Manager to vSphere Lifecycle Manager, downloading patches from VMware patch repository might fail**

  In earlier releases of vCenter Server you could configure independent proxy settings for vCenter Server and vSphere Update Manager. After an upgrade to vSphere 7.0, vSphere Update Manager service becomes part of the vSphere Lifecycle Manager service. For the vSphere Lifecycle Manager service, the proxy settings are configured from the vCenter Server appliance settings. If you had configured Update Manager to download patch updates from the Internet through a proxy server but the vCenter Server appliance had no proxy setting configuration, after a vCenter Server upgrade to version 7.0, the vSphere Lifecycle Manager fails to connect to the VMware depot and is unable to download patches or updates.

  Workaround: Log in to the vCenter Server Appliance Management Interface, https://*vcenter-server-appliance-FQDN-or-IP-address*:5480, to configure proxy settings for the vCenter Server appliance and enable vSphere Lifecycle Manager to use proxy.

- **If you use a Java client to review remediation tasks, you cannot extract the results from the remediation operations**

  If you use a Java client to review remediation tasks, extracting the results might fail with a `ConstraintValidationException` error. The issue occurs when an ESXi host fails to enter maintenance mode during the remediation and gets a status SKIPPED, but at the same time wrongly gets an In Progress flag for the consecutive remediation operations. This causes the `ConstraintValidationException` error on the Java Clients and you cannot extract the result of the remediation operation.

  Workaround: Fix the underlying issues that prevent ESXi hosts to enter Maintenance Mode and retry the remediation operation.

- **The general vSphere Lifecycle Manager depot and local depots in Remote Office and Branch Office (ROBO) deployments might not be in sync**

  ROBO clusters that have limited or no access to the Internet or limited connectivity to vCenter Server can download an image from a depot that is local for them instead of accessing the vSphere Lifecycle Manager depot in vCenter Server. However, vSphere Lifecycle Manager generates software recommendations in the form of pre-validated images only on a central level and a recommended image content might not be available at a depot override.

  Workaround: If you decide to use a recommended image, make sure the content between depot overrides and the central depot are in sync.

- **Cluster remediation by using the vSphere Lifecycle Manager might fail on ESXi hosts with enabled lockdown mode**

  If a cluster has ESXi hosts with enabled lockdown mode, remediation operations by using the vSphere Lifecycle Manager might skip such hosts. In the log files, you see messages such as `Host scan task failed` and `com.vmware.vcIntegrity.lifecycle.EsxImage.UnknownError An unknown error occurred while performing the operation..`

  Workaround: Add the root user to the exception list for lockdown mode and retry the cluster remediation.

- **You cannot use an ESXi image with OEM content to create clusters by using vSphere Lifecycle Manager workflows after an update**

  When you update an ESXi image with OEM content, for example Dell ESXi 7.0 Update 2a to Dell ESXi 7.0 Update 3d, some reserved VIBs, such that are present in the VMware base image but overridden by async VIBs packaged by the OEM, might be deleted. The same issue might occur after updating your system by using a non-critical baseline and then update to a higher version by using an ESXi image with OEM content. For example, if you update to 7.0 Update 2 by using a non-critical baseline and then update to 7.0 Update 3 by using a Dell ESXi 7.0 Update 3 image, reserved VIBs might also be deleted. As a result, you cannot use ESXi hosts with the updated version to create a vSphere Lifecycle Manager cluster managed by a single image. The operation fails and in the backtrace, you see errors such as:

  ```
  2021-11-24T09:42:49Z lifecycle: 2101166: HostSeeding:956 ERROR Extract depot failed: ('VMW_bootbank_bnxtroce_216.0.58.0-23vmw.703.0.0.18644231', 'Failed to add reserved VIB VMW_bootbank_bnxtroce_216.0.58.0-23vmw.703.0.0.18644231: not found in the reserved VIB cache storage')2021-11-24T09:42:50Z lifecycle: 2101166: imagemanagerctl:373 ERROR Extract depot failed.021-11-24T09:42:50Z lifecycle: 2101166: imagemanagerctl:152 ERROR [ReservedVibExtractError] ('VMW_bootbank_bnxtroce_216.0.58.0-23vmw.703.0.0.18644231', 'Failed to add reserved VIB VMW_bootbank_bnxtroce_216.0.58.0-23vmw.703.0.0.18644231: not found in the reserved VIB cache storage')
  ```

  Workaround: Use interactive or scripted upgrade instead of vSphere Lifecycle Manager workflows.

- **After you upgrade to vCenter Server 7.0.0b, in the vSphere Lifecycle Manager home view in the vSphere Client, you do not see the Show only rollup updates toggle button**

  In vCenter Server 7.0.0b, you can use the **Show only rollup updates** toggle button to filter and select patches that you want to include in a baseline when you use the vSphere Lifecycle Manager.

  The button is available in the **Updates** tab on the **Lifecycle Manager** pane, **Menu** > **Lifecycle Manager**, which is the vSphere Lifecycle Manager home view in the vSphere Client. The button is also available in the **Select Patches Manually** page on the **Baselines** tab in the **Create Baseline** wizard, which opens when you select **New** > **Baseline**.

However, the **Show only rollup updates** toggle button might not be visible after you upgrade to vCenter Server 7.0.0b.

Workaround: After an upgrade to vCenter Server 7.0.0b, restart the vSphere Client. For more information, see Start, Stop, and Restart Services.

- **The Show only rollup updates toggle button is always turned on when you open a tab in the vSphere Lifecycle Manager home view in the vSphere Client**

  In vCenter Server 7.0.0b, you can use the **Show only rollup updates** toggle button to filter and select patches that you want to include in a baseline when you use the vSphere Lifecycle Manager.

  The button is available in the **Updates** tab on the **Lifecycle Manager** pane, **Menu** > **Lifecycle Manager**, which is the vSphere Lifecycle Manager home view in the vSphere Client. The button is also available in the **Select Patches Manually** page on the **Baselines** tab in the **Create Baseline** wizard, which opens when you select **New** > **Baseline**.

  However, the toggle button appears always turned on when you navigate to either of the **Updates** tab or the **Select Patches Manually** page. Even if you turn off the button when navigating away from the tab or page, it appears still turned on the next time you open them.

  Workaround: None

- **When you use the Update Planner, in the vSphere Client you might see Unexpected error occurred while fetching the updates**

  When you use Update Planner, which is part of vSphere Lifecycle Manager, used to facilitate vCenter Server updates, you might see the following error in the vSphere Client:

  ```
  Unexpected error occurred while fetching the updates
  ```

  The issue occurs when you use a custom HTTPS port that prevents you from running interoperability reports by using the vSphere Client.

  Workaround: Manually invoke the API. For more details, see the vSphere Automation API.

Miscellaneous Issues

- **When applying a host profile with version 6.5 to a ESXi host with version 7.0, the compliance check fails**

  Applying a host profile with version 6.5 to a ESXi host with version 7.0, results in Coredump file profile reported as not compliant with the host.

  Workaround: There are two possible workarounds.

  a When you create a host profile with version 6.5, set an advanced configuration option VMkernel.Boot.autoCreateDumpFile to false on the ESXi host.

b When you apply an existing host profile with version 6.5, add an advanced configuration option VMkernel.Boot.autoCreateDumpFile in the host profile, configure the option to a fixed policy, and set value to false.

- **Cannot migrate container volume due to insufficient space**

  The Migrate volume dialog in vSphere Client considers the current size of container volumes and the free space available on the target datastores. If the free space of a datastore is less than the volume size, then you cannot select the datastore for migration. vSphere preserves the storage policy as it migrates the volume, but the file size of the volume might change depending on the destination datastore type and its supported features. The vSphere Client might show the datastore has insufficient space to perform the migration, even though enough space is available.

  Workaround: None

- **The Actions drop-down menu does not contain any items when your browser is set to language different from English**

  When your browser is set to language different from English and you click the **Switch to New View** button from the virtual machine **Summary** tab of the vSphere Client inventory, the **Actions** drop-down menu in the **Guest OS** panel does not contain any items.

  Workaround: Select the **Actions** drop-down menu on the top of the virtual machine page.

- **Mellanox ConnectX-4 or ConnectX-5 native ESXi drivers might exhibit minor throughput degradation when Dynamic Receive Side Scaling (DYN_RSS) or Generic RSS (GEN_RSS) feature is turned on**

  Mellanox ConnectX-4 or ConnectX-5 native ESXi drivers might exhibit less than 5 percent throughput degradation when DYN_RSS and GEN_RSS feature is turned on, which is unlikely to impact normal workloads.

  Workaround: You can deactivate DYN_RSS and GEN_RSS feature with the following commands:

  ```
  # esxcli system module parameters set -m nmlx5_core -p "DYN_RSS=0 GEN_RSS=0"

  # reboot
  ```

- **RDMA traffic between two VMs on the same host might fail in PVRDMA environment**

  In a vSphere 7.0 implementation of a PVRDMA environment, VMs pass traffic through the HCA for local communication if an HCA is present. However, loopback of RDMA traffic does not work on qedrntv driver. For instance, RDMA Queue Pairs running on VMs that are configured under same uplink port cannot communicate with each other.

  In vSphere 6.7 and earlier, HCA was used for local RDMA traffic if SRQ was enabled. vSphere 7.0 uses HCA loopback with VMs using versions of PVRDMA that have SRQ enabled with a minimum of HW v14 using RoCE v2.

The current version of Marvell FastLinQ adapter firmware does not support loopback traffic between QPs of the same PF or port.

Workaround: Required support is being added in the out-of-box driver certified for vSphere 7.0. If you are using the inbox qedrntv driver, you must use a 3-host configuration and migrate VMs to the third host.

- **Unreliable Datagram traffic QP limitations in qedrntv driver**

  There are limitations with the Marvell FastLinQ qedrntv RoCE driver and Unreliable Datagram (UD) traffic. UD applications involving bulk traffic might fail with qedrntv driver. Additionally, UD QPs can only work with DMA Memory Regions (MR). Physical MRs or FRMR are not supported. Applications attempting to use physical MR or FRMR along with UD QP fail to pass traffic when used with qedrntv driver. Known examples of such test applications are `ibv_ud_pingpong` and `ib_send_bw`.

  Standard RoCE and RoCEv2 use cases in a VMware ESXi environment such as iSER, NVMe-oF (RoCE) and PVRDMA are not impacted by this issue. Use cases for UD traffic are limited and this issue impacts a small set of applications requiring bulk UD traffic.

  Marvell FastLinQ hardware does not support RDMA UD traffic offload. In order to meet the VMware PVRDMA requirement to support GSI QP, a restricted software only implementation of UD QP support was added to the qedrntv driver. The goal of the implementation is to provide support for control path GSI communication and is not a complete implementation of UD QP supporting bulk traffic and advanced features.

  Since UD support is implemented in software, the implementation might not keep up with heavy traffic and packets might be dropped. This can result in failures with bulk UD traffic.

  Workaround: Bulk UD QP traffic is not supported with qedrntv driver and there is no workaround at this time. VMware ESXi RDMA (RoCE) use cases like iSER, NVMe, RDMA and PVRDMA are unaffected by this issue.

- **Servers equipped with QLogic 578xx NIC might fail when frequently connecting or disconnecting iSCSI LUNs**

  If you trigger QLogic 578xx NIC iSCSI connection or disconnection frequently in a short time, the server might fail due to an issue with the qfle3 driver. This is caused by a known defect in the device's firmware.

  Workaround: None.

- **ESXi might fail during driver unload or controller disconnect operation in Broadcom NVMe over FC environment**

  In Broadcom NVMe over FC environment, ESXi might fail during driver unload or controller disconnect operation and display an error message such as: `@BlueScreen: #PF Exception 14 in world 2098707:vmknvmeGener IP 0x4200225021cc addr 0x19`

  Workaround: None.

- **ESXi does not display OEM firmware version number of i350/X550 NICs on some Dell servers**

  The inbox ixgben driver only recognizes firmware data version or signature for i350/X550 NICs. On some Dell servers the OEM firmware version number is programmed into the OEM package version region, and the inbox ixgben driver does not read this information. Only the 8-digit firmware signature is displayed.

  Workaround: To display the OEM firmware version number, install async ixgben driver version 1.7.15 or later.

- **X710 or XL710 NICs might fail in ESXi**

  When you initiate certain destructive operations to X710 or XL710 NICs, such as resetting the NIC or manipulating VMKernel's internal device tree, the NIC hardware might read data from non-packet memory.

  Workaround: Do not reset the NIC or manipulate vmkernel internal device state.

- **NVMe-oF does not guarantee persistent VMHBA name after system reboot**

  NVMe-oF is a new feature in vSphere 7.0. If your server has a USB storage installation that uses vmhba30+ and also has NVMe over RDMA configuration, the VMHBA name might change after a system reboot. This is because the VMHBA name assignment for NVMe over RDMA is different from PCIe devices. ESXi does not guarantee persistence.

  Workaround: None.

- **Backup fails for vCenter database size of 300 GB or greater**

  If the vCenter database size is 300 GB or greater, the file-based backup will fail with a timeout. The following error message is displayed: `Timeout! Failed to complete in 72000 seconds`

  Workaround: None.

- **A restore of vCenter Server 7.0 which is upgraded from vCenter Server 6.x with External Platform Services Controller to vCenter Server 7.0 might fail**

  When you restore a vCenter Server 7.0 which is upgraded from 6.x with External Platform Services Controller to vCenter Server 7.0, the restore might fail and display the following error: `Failed to retrieve appliance storage list`

  Workaround: During the first stage of the restore process, increase the storage level of the vCenter Server 7.0. For example if the vCenter Server 6.7 External Platform Services Controller setup storage type is small, select storage type large for the restore process.

- **Enabled SSL protocols configuration parameter is not configured during a host profile remediation process**

`Enabled SSL protocols` configuration parameter is not configured during a host profile remediation and only the system default protocol `tlsv1.2` is enabled. This behavior is observed for a host profile with version 7.0 and earlier in a vCenter Server 7.0 environment.

Workaround: To enable TLSV 1.0 or TLSV 1.1 SSL protocols for SFCB, log in to an ESXi host by using SSH, and run the following ESXCLI command: `esxcli system wbem -P <protocol_name>`

- **Unable to configure Lockdown Mode settings by using Host Profiles**

  Lockdown Mode cannot be configured by using a security host profile and cannot be applied to multiple ESXi hosts at once. You must manually configure each host.

  Workaround: In vCenter Server 7.0, you can configure Lockdown Mode and manage Lockdown Mode exception user list by using a security host profile.

- **When a host profile is applied to a cluster, Enhanced vMotion Compatibility (EVC) settings are missing from the ESXi hosts**

  Some settings in the VMware config file `/etc/vmware/config` are not managed by Host Profiles and are blocked, when the config file is modified. As a result, when the host profile is applied to a cluster, the EVC settings are lost, which causes loss of EVC functionalities. For example, unmasked CPUs can be exposed to workloads.

  Workaround: Reconfigure the relevant EVC baseline on cluster to recover the EVC settings.

- **Using a host profile that defines a core dump partition in vCenter Server 7.0 results in an error**

  In vCenter Server 7.0, configuring and managing a core dump partition in a host profile is not available. Attempting to apply a host profile that defines a core dump partition, results in the following error: `No valid coredump partition found.`

  Workaround: None. In vCenter Server 7.0., Host Profiles supports only file-based core dumps.

- **You cannot paste by right clicking in the vi editor**

  If you use the vi text editor on a vCenter Server Appliance instance, you cannot paste previously copied content by right clicking. Instead, the vi editor changes mode to Visual.

  Workaround: From the vCenter Server Appliance shell, run the command`/usr/bin/sed -i.bak -E '2alet skip_defaults_vim=1' /etc/vimrc`.

- **HTTP requests from certain libraries to vSphere might be rejected**

  The HTTP reverse proxy in vSphere 7.0 enforces stricter standard compliance than in previous releases. This might expose pre-existing problems in some third-party libraries used by applications for SOAP calls to vSphere.

If you develop vSphere applications that use such libraries or include applications that rely on such libraries in your vSphere stack, you might experience connection issues when these libraries send HTTP requests to VMOMI. For example, HTTP requests issued from vijava libraries can take the following form:

```
POST /sdk HTTP/1.1
SOAPAction
Content-Type: text/xml; charset=utf-8
User-Agent: Java/1.8.0_221
```

The syntax in this example violates an HTTP protocol header field requirement that mandates a colon after SOAPAction. Hence, the request is rejected in flight.

Workaround: Developers leveraging noncompliant libraries in their applications can consider using a library that follows HTTP standards instead. For example, developers who use the vijava library can consider using the latest version of the yavijava library instead.

- **SNMP dynamic firewall ruleset is modified by Host Profiles during a remediation process**

  The SNMP firewall ruleset is a dynamic state, which is handled during runtime. When a host profile is applied, the configuration of the ruleset is managed simultaneously by Host Profiles and SNMP, which can modify the firewall settings unexpectedly.

  Workaround: There are two possible workarounds.

  - To allow the ruleset to manage itself dynamically, exclude the SNMP firewall ruleset option in the configuration of the host profile.
  - To proceed with the double management of the ruleset, when needed, correct the firewall ruleset state.

- **You might see a dump file when using Broadcom driver lsi_msgpt3, lsi_msgpt35 and lsi_mr3**

  When using the lsi_msgpt3, lsi_msgpt35 and lsi_mr3 controllers, there is a potential risk to see dump file lsuv2-lsi-drivers-plugin-util-zdump. There is an issue when exiting the storelib used in this plugin utility. There is no impact on ESXi operations, you can ignore the dump file.

  Workaround: You can safely ignore this message. You can remove the lsuv2-lsi-drivers-plugin with the following command:

  ```
  esxcli software vib remove -n lsuv2-lsiv2-drivers-plugin
  ```

- **You might see reboot is not required after configuring SR-IOV of a PCI device in vCenter, but device configurations made by third party extensions might be lost and require reboot to be re-applied.**

  In ESXi 7.0, SR-IOV configuration is applied without a reboot and the device driver is reloaded. ESXi hosts might have third party extensions perform device configurations that need to run after the device driver is loaded during boot. A reboot is required for those third party extensions to re-apply the device configuration.

Workaround: You must reboot after configuring SR-IOV to apply third party device configurations.

- **You see black or grey zones in the background of a parent window in the Direct Console User Interface (DCUI) after a child window closes**

  In the DCUI, when you close a child window by pressing the ESC or Enter keys, or the **Cancel** or **OK** buttons, the parent window appearance might change. The background color changes to grey or black for some part of the parent window. However, all required information from the DCUI is properly displayed and all operations performed in the DCUI complete successfully.

  Workaround: Wait for 1 minute without refreshing the current window in the DCUI or pressing any key.

Backup Issues

- **If you use the NFS and SMB protocols for file-based backup of vCenter Server, the backup fails after an update from vCenter Server 7.x to vCenter Server 7.0 Update 1**

  If you use the Network File System (NFS) and Server Message Block (SMB) protocols for file-based backup of vCenter Server, the backup fails after an update from an earlier version of vCenter Server 7.x to vCenter Server 7.0 Update 1. In the `applmgmt.log`, you see an error message such as `Failed to mount the remote storage`. The issue occurs because of Linux kernel updates that run during the patch process. The issue does not occur on fresh installations of vCenter Server 7.0 Update 1.

  Workaround: Reboot the vCenter Server appliance after the update is complete.

- **Microsoft Active Directory Federation Services (ADFS) logins might fail after restoring a vCenter Server Appliance**

  If you manually add a certificate to the vCenter Server JRE truststore or modify the `/etc/hosts` file when setting up ADFS, the changes are not preserved after restoring and might cause ADFS logins to fail.

  Workaround: Add the ADFS certificate to the vCenter Server JRE truststore after restoring your vCenter Server Appliance. For more information, see Import the Trusted Certificate of an External Identity Provider. Add the necessary host name mappings back to the `/etc/hosts` file after restoring your vCenter Server Appliance.

vSAN Issues

- **Higher than actual vSAN storage pool free space estimate might lead to unexpected storage shortages**

The `freeSpace` parameter of a vSAN storage pool might not exclude disks with errors and display the unused space estimate as higher than the actual. As a result, you might see unexpected storage shortages. This issue only affects vSAN storage pools.

Workaround: Manually compute the storage capacity by disregarding the `freeSpace` value of disks with errors from the overall storage capacity provided by the `freeSpace` parameter of the storage pool.

Server Configuration Issues

- **You see a message Error retrieving when trying to view the status of Key Management Server (KMS) instances in the vSphere Client**

  In the vSphere Client, you might see a message `Error retrieving` for several minutes when trying to view Key Management Server (KMS) instances. The issue occurs when a KMS instance in a standard key provider loses connectivity. Until all network requests to the affected KMS time out, which takes around 4 minutes, you cannot see the status of any KMS instance in your system, only the error message for the key provider. After the timeout, you can see the status of all KMS instances.

  Workaround: If you see the Error retrieving message, wait for 4 minutes.

- **Due to a default timeout setting of 2 minutes, log in to vCenter instances in Enhanced Linked Mode takes long when any of the vCenters is down**

  When you have a number of vCenter instances in Enhanced Linked Mode, if one of the instances is down for some reason, such as maintenance, log in to the other instances might take up to 2 minutes, which is the default timeout setting. Changing the timeout property that specifies the wait time for log in to linked vCenter instances is a complex task that requires manually editing the `LinkedVcGroup.login.timeout` property located in `vim-commons-vsphere.properties`. To simlify the task, starting with vCenter Server 7.0 Update 3l, the setting `LinkedVcGroup.login.timeout=120000` moves from `vim-commons-vsphere.properties` to the `webclient.properties` file. Editing this option allows you to reduce the wait time so that if one of the vCenter logins is taking more time, it does not affect the log in time for other instances.

  Workaround: Edit `/etc/vmware/vsphere-ui/webclient.properties` and change the value of `LinkedVcGroup.login.timeout` from `120000` to a smaller value in milliseconds, but consider that a value of `<= 0` is an infinite timeout.