VMware ESXi 8.0 Update 1c Release Notes

VMware vSphere 8.0

ESXi 8.0 Update 1

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

https://docs.vmware.com/

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

Contents

Introduction

1

ESXi 8.0 Update 1c | 27 JUL 2023 | Build 22088125

Check for additions and updates to these release notes.

What's New

2

ESXi 8.0 Update 1c supports vSphere Quick Boot on the following servers:

- Cisco Systems Inc
  - UCSC-C225-M6S
- Dell Inc.
  - R660 vSAN Ready Node
  - R760 vSAN Ready Node
  - PowerEdge R660xs
  - PowerEdge R760xd2
  - PowerEdge R760xa
  - PowerEdge R760xs
  - PowerEdge R860
  - PowerEdge R960
  - PowerEdge T560
- HPE
  - Alletra 4120
  - HPE Cray XD220v
  - ProLiant DL320 Gen11
  - ProLiant DL360 Gen11
  - ProLiant DL380 Gen11
  - ProLiant DL380a Gen11
  - ProLiant DL560 Gen11
  - ProLiant ML110 Gen11
  - ProLiant ML350 Gen11

- Lenovo
  - ThinkSystem SR630 V3
  - ThinkSystem SR650 V3

Earlier Releases of ESXi 8.0

3

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

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

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

Patches Contained in This Release

4

Read the following topics next:

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

Patch for VMware ESXi 8.0 Update 1c

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

Build Details

Download Filename:	VMware-ESXi-8.0U1c-22088125-depot.zip
Build:	22088125
Download Size:	949.1 MB
sha256checksum:	ab25ede2b6c40d6bb551f41ae43187fe124cc83853e9df25cd79ccf6836546a8
Host Reboot Required:	Yes
Virtual Machine Migration or Shutdown Required:	Yes

Components

Component	Bulletin	Category	Severity
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.25.22088125	Bugfix	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.25.22088125	Bugfix	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.25.22088125	Bugfix	Critical

Component	Bulletin	Category	Severity
Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi	Mellanox-nmlx5_4.23.0.36-15vmw.801.0.25.22088125	Bugfix	Critical
Broadcom NetXtreme I ESX VMKAPI ethernet driver	Broadcom-ntg3_4.1.10.0-5vmw.801.0.25.22088125	Bugfix	Critical
VMware NVMe over TCP Driver	VMware-NVMeoF-TCP_1.0.1.7-1vmw.801.0.25.22088125	Bugfix	Critical
ESXi Component - core ESXi VIBs	ESXi_8.0.1-0.20.22082334	Security	Critical
ESXi Install/Upgrade Component	esx-update_8.0.1-0.20.22082334	Security	Critical
ESXi Install/Upgrade Component	esxio-update_8.0.1-0.20.22082334	Security	Critical
ESXi Tools Component	VMware-VM-Tools_12.2.5.21855600-22082334	Security	Critical

Rollup Bulletin

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

Bulletin ID	Category	Severity	Detail
ESXi80U1c-22088125	Bugfix	Critical	Security fixes and Bug fixes
ESXi80U1sc-22082334	Security	Critical	Security only fixes

Image Profiles

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

Image Profile Name
ESXi-8.0U1c-22088125-standard
ESXi-8.0U1c-22088125-no-tools
ESXi-8.0U1sc-22082334-standard
ESXi-8.0U1sc-22082334-no-tools

ESXi Image

Name and Version	Release Date	Category	Detail
ESXi 8.0 U1c - 22088125	07/27/2023	Bugfix	Security and Bugfix image
ESXi 8.0 U1sc - 22082334	07/27/2023	Security	Security only image

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

Patch Download and Installation

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

Resolved Issues

5

Read the following topics next:

- ESXi_8.0.1-0.25.22088125
- esx-update_8.0.1-0.25.22088125
- esxio-update_8.0.1-0.25.22088125
- Mellanox-nmlx5_4.23.0.36-15vmw.801.0.25.22088125
- Broadcom-ntg3_4.1.10.0-5vmw.801.0.25.22088125
- VMware-NVMeoF-TCP_1.0.1.7-1vmw.801.0.25.22088125
- ESXi_8.0.1-0.20.22082334
- esx-update_8.0.1-0.20.22082334
- esxio-update_8.0.1-0.20.22082334
- VMware-VM-Tools_12.2.5.21855600-22082334
- ESXi-8.0U1c-22088125-standard
- ESXi-8.0U1c-22088125-no-tools
- ESXi-8.0U1sc-22082334-standard
- ESXi-8.0U1sc-22082334-no-tools
- ESXi 8.0 U1c - 22088125
- ESXi 8.0 U1sc - 22082334

ESXi_8.0.1-0.25.22088125

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

Affected Software	N/A
Affected VIBs	■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-base_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 ■ VMware_bootbank_vdfs_8.0.1-0.25.22088125 ■ VMware_bootbank_bmcal_8.0.1-0.25.22088125 ■ VMware_bootbank_clusterstore_8.0.1-0.25.22088125 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 ■ VMware_bootbank_gc_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 ■ VMware_bootbank_vsan_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-base_8.0.1-0.25.22088125 ■ VMware_bootbank_trx_8.0.1-0.25.22088125 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 ■ VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 ■ VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 ■ VMware_bootbank_crx_8.0.1-0.25.22088125
PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111
CVE numbers	N/A

Updates the `esxio-combiner-esxio, esxio-base, esxio-dvfilter-generic-fastpath, vdfs, bmcal, clusterstore, native-misc-drivers-esxio, gc-esxio, esxio, bmcal-esxio, esxio-combiner, gc, esx-xserver, vsan, esx-base, trx, native-misc-drivers, cpu-microcode, vsanhealth`, and `crx` VIBs to resolve the following issues:

- **Transient vSAN health check warning: Network configuration is out of sync**

  vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

  This issue is resolved in this release.

- **A race condition between two operations could potentially lead to inconsistent metadata**

  In extremely rare cases, a race condition between an unmap operation on a VMDK and a concurrent snapshot creation operation on the same object could cause inconsistent vSAN metadata for that object in a vSAN host. This inconsistency could potentially lead to VM failure or host failure.

  This issue is resolved in this release.

- **ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU**

  The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

  ```
  ^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't have a heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked up.^[[0m

  ^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI: RIPOFF(base):RBP:CS
  ```

  This issue is resolved in this release.

- **ESXi ConfigStore database fills up and writes fail**

  Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

  ```
  2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend visorfs file /etc/vmware/configstore/current-store-1-journal because its ramdisk (configstore) is full.
  ```

  This issue is resolved in this release.

- **Delay before evacuating vSAN storage device in UNHEALTHY state**

After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

This issue is resolved in this release.

- **vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility**

  This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

  This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

- **The vSAN iSCSI Target service might fail due to a rare race condition**

  When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

  This issue is resolved in this release.

- **If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later**

  If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

  This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

- **When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen**

  Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

  ```
  0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack:
  0x452900262dc8, 0x4302f6c06508, 0x4200138fee8b,
  0x420013df1300, 0x452900262cf0  0x452900262dc0:
  [0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
  ```

```
0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514  0x452900262e30:
[0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
0x0, 0x0  0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0  0x452900262f20:
[0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068, 0xf50,
0xf50, 0x0  0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover+0x68
stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088  0x45397b61bd30:
[0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack: 0x1,
0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757  0x45397b61bd50:
[0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack:
0x100005cf757, 0x0, 0x0, 0x4200139b36d9, 0x0  0x45397b61bd80:
[0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover+0x59 stack: 0x39,
0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0  0x45397b61bdc0:
[0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0,
0xa, 0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0
```

This issue is resolved in this release.

- **vSAN cluster shutdown fails on a cluster with IPV6 disabled**

  This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: `'NoneType' object is not iterable`.

  This issue is resolved in this release.

- **You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled**

  After the hostd service starts, for example after an ESXi host reboot, it might create a SNMP directory under `/var/spool` and you see many `.trp` files to pile up in this directory.

  This issue is resolved in this release. The fix makes sure that the directory `/var/spool/snmp` exists only when SNMP is enabled.

- **The hostd service repeatedly fails and the ESXi host disconnects from the vCenter system**

  If for any reason an ESXi host is temporarily in a state of insufficient memory, the hostd service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

  This issue is resolved in this release.

- **During service insertion for NSX-managed workload VMs, some VMS might become intermittently unresponsive and the virtual device might reset**

During service insertion for NSX-managed workload VMs, a packet list might be reinjected from the input chain of one switchport to another switchport. In such cases, the source switchport does not correspond to the actual portID of the input chain and the virtual device does not get completion status for the transmitted frame. As a result, when you run a service insertion task, some VMs might become intermittently unresponsive due to network connectivity issues.

This issue is resolved in this release.

- **Performance of certain nested virtual machines on AMD CPUs might degrade**

  Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

  This issue is resolved in this release.

- **Changing the mode of the virtual disk** on **a running virtual machine might cause the VM to fail**

  If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the vmware.log, you see errors such as:

  ```
  msg.disk.notConfigured2] Failed to configure disk 'scsi0:4'. The virtual machine cannot be powered on with an unconfigured disk.
  ```

  ```
  [msg.checkpoint.continuesync.error] An operation required the virtual machine to quiesce and the virtual machine was unable to continue running.
  ```

  This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

- **ESXi NVMe/TCP initiator fails to recover paths after target failure recovery**

  When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

  This issue is resolved in this release.

- **ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host**

  Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

  This issue is resolved in this release.

- **A Logical Volume Manager (LVM) disk goes offline during datastore expansion**

If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

This issue is resolved in this release.

- **The durable name of a SCSI LUN might not be set**

  The durable name property for a SCSI-3 compliant device comes from pages `80h` and `83h` of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

  This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages `80h` and `83h`, if that data is not required for the generation of durable name.

- **Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel**

  In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

  This issue is resolved in this release.

- **NVMe over Fabrics controller might unexpectedly disconnect during discovery**

  If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

  This issue is resolved in this release.

- **If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures**

  When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

  This issue is resolved in this release. For more details, see VMware knowledge base article 92276.

- **If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline**

  While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

  This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

- **Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues**

  Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

  This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

  a On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings** > **System Image Cache Configuration**

  b In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.

  c Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

- **VMware VIB installation might fail during concurrent vendor package installations**

  When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration`.

  This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

- **Certain applications might take too many ESXi file handles and cause performance aggravation**

  In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

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

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

  This issue is resolved in this release.

esx-update_8.0.1-0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	■ VMware_bootbank_loadesx_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-update_8.0.1-0.25.22088125
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesx` and `esx-update` VIBs.

esxio-update_8.0.1-0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	■ VMware_bootbank_loadesxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-update_8.0.1-0.25.22088125

PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

Mellanox-nmlx5_4.23.0.36-15vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 ■ VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 ■ VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 ■ VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125
PRs Fixed	3152476
CVE numbers	N/A

Updates the `nmlx5-rdma`, `nmlx5-core`, `nmlx5-rdma-esxio`, and `nmlx5-core-esxio` VIBs to resolve the following issue:

- **In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs**

  In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (nmlx5) driver might be low and you might see uneven traffic flow across RSS queues.

  This issue is resolved in this release.

Broadcom-ntg3_4.1.10.0-5vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes

Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125
PRs Fixed	3187446
CVE numbers	N/A

Updates the `ntg3` VIB to resolve the following issue:

- **After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network**

  Changes in the ntg3 driver version `4.1.9.0-4vmw` might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

  This issue is resolved in this release.

VMware-NVMeoF-TCP_1.0.1.7-1vmw.801.0.25.22088125

Patch Category	Bugfix
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125
PRs Fixed	3219191
CVE numbers	N/A

Updates the `nvmetcp` VIB to resolve the following issue:

- **Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations** vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

This issue is resolved in this release.

ESXi_8.0.1-0.20.22082334

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

VIBs Included	- VMware_bootbank_bmcal-esxio_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 - VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsan_8.0.1-0.20.22082334 - VMware_bootbank_bmcal_8.0.1-0.20.22082334 - VMware_bootbank_esx-base_8.0.1-0.20.22082334 - VMware_bootbank_trx_8.0.1-0.20.22082334 - VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 - VMware_bootbank_gc_8.0.1-0.20.22082334 - VMware_bootbank_esxio_8.0.1-0.20.22082334 - VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 - VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334 - VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 - VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 - VMware_bootbank_clusterstore_8.0.1-0.20.22082334 - VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 - VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 - VMware_bootbank_esxio-base_8.0.1-0.20.22082334 - VMware_bootbank_vdfs_8.0.1-0.20.22082334 - VMware_bootbank_crx_8.0.1-0.20.22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149
CVE numbers	N/A

Updates the `bmcal-esxio, native-misc-drivers, esx-dvfilter-generic-fastpath, vsan, bmcal, esx-base, trx, gc-esxio, gc, esxio, esx-xserver, cpu-microcode, esxio-combiner-esxio, esxio-dvfilter-generic-fastpath, vsanhealth, vds-vsip, clusterstore, native-misc-drivers-esxio, esxio-combiner, esxio-base, vdfs`, and *crx* VIBs to resolve the following issues:

- **ESXi 8.0 Update 1c provides the following security updates:**
  - The Envoy proxy is updated to version v1.23.9.
  - The ESXi userworld libxml2 library is updated to version 2.10.4.
  - The cURL library is updated to version 8.0.1.
  - The Go library is updated to version 1.19.9.

- The etcd package is updated to 3.4.25.

esx-update_8.0.1-0.20.22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMware_bootbank_esx-update_8.0.1-0.20.22082334 ■ VMware_bootbank_loadesx_8.0.1-0.20.22082334
PRs Fixed	N/A
CVE numbers	N/A

Updates the `esx-update` and `loadesx` VIBs.

esxio-update_8.0.1-0.20.22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	Yes
Virtual Machine Migration or Shutdown Required	Yes
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMware_bootbank_loadesxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-update_8.0.1-0.20.22082334
PRs Fixed	N/A
CVE numbers	N/A

Updates the `loadesxio` and `esxio-update` VIBs.

VMware-VM-Tools_12.2.5.21855600-22082334

Patch Category	Security
Patch Severity	Critical
Host Reboot Required	No
Virtual Machine Migration or Shutdown Required	No
Affected Hardware	N/A
Affected Software	N/A
VIBs Included	■ VMware_locker_tools-light_12.2.5.21855600-22082334
PRs Fixed	3186166
CVE numbers	N/A

- Updates the `tools-light` VIB.
  - The following VMware Tools ISO images are bundled with ESXi 8.0 Update 1c:
    - **windows.iso**: VMware Tools 12.2.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
    - **linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

    The following VMware Tools ISO images are available for download:

    - VMware Tools 11.0.6:
      - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
    - VMware Tools 10.0.12:
      - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
      - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
    - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
    - **darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article 88698 for details.

    Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

    - Updating VMware Tools
    - VMware Tools for hosts provisioned with Auto Deploy

- What Every vSphere Admin Must Know About VMware Tools
- Earlier versions of VMware Tools
- VMware Tools 12.2.5 Release Notes

ESXi-8.0U1c-22088125-standard

Profile Name	ESXi-8.0U1c-22088125-standard
Build	For build information, see Patches Contained in This Release.
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_vdfs_8.0.1-0.25.22088125 ■ VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125 ■ VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125 ■ VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_gc_8.0.1-0.25.22088125 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 ■ VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.25.22088125 ■ VMware_bootbank_vds-vsip_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 ■ VMware_bootbank_loadesxio_8.0.1-0.25.22088125 ■ VMware_bootbank_loadesx_8.0.1-0.25.22088125 ■ VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_esxio-update_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-update_8.0.1-0.25.22088125 ■ VMware_bootbank_vsan_8.0.1-0.25.22088125 ■ VMware_bootbank_crx_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_bmcal_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 ■ VMware_bootbank_trx_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-base_8.0.1-0.25.22088125 ■ VMware_bootbank_clusterstore_8.0.1-0.25.22088125 ■ Mware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-base_8.0.1-0.25.22088125 ■ VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_locker_tools-light_12.2.5.21855600-22082334

PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

This patch updates the following issues:

- **Transient vSAN health check warning: Network configuration is out of sync**

  vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

  This issue is resolved in this release.

- **vSAN cache overwrite issue might cause inconsistent metadata**

  In rare cases, racing between unmap and snapshot operations might cause inconsistent vSAN metadata in a vSAN host. Such inconsistency can lead to unpredictable consequences, including VM failures or host failures.

  This issue is resolved in this release.

- **ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU**

  The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

  ```
  ^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't have a heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked up.^[[0m
  ```

  ```
  ^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI: RIPOFF(base):RBP:CS
  ```

This issue is resolved in this release.

- **ESXi ConfigStore database fills up and writes fail**

  Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

  ```
  2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend
  visorfs file /etc/vmware/configstore/current-store-1-journal because its ramdisk
  (configstore) is full.
  ```

  This issue is resolved in this release.

- **Delay before evacuating vSAN storage device in UNHEALTHY state**

  After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

  This issue is resolved in this release.

- **vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility**

  This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

  This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

- **The vSAN iSCSI Target service might fail due to a rare race condition**

  When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

  This issue is resolved in this release.

- **If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later**

  If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

- **When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen**

  Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

  ```
  0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack:
  0x452900262dc8, 0x4302f6c06508, 0x4200138fee8b,
  0x420013df1300, 0x452900262cf0  0x452900262dc0:
  [0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
  0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514  0x452900262e30:
  [0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
  0x0, 0x0  0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
  0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0  0x452900262f20:
  [0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068, 0xf50,
  0xf50, 0x0  0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover+0x68
  stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088  0x45397b61bd30:
  [0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack: 0x1,
  0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757  0x45397b61bd50:
  [0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack:
  0x100005cf757, 0x0, 0x0, 0x4200139b36d9, 0x0  0x45397b61bd80:
  [0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover+0x59 stack: 0x39,
  0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0  0x45397b61bdc0:
  [0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0,
  0xa, 0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0
  ```

  This issue is resolved in this release.

- **After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network**

  Changes in the ntg3 driver version `4.1.9.0-4vmw` might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

  This issue is resolved in this release.

- **vSAN cluster shutdown fails on a cluster with IPV6 disabled**

  This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: `'NoneType' object is not iterable`.

  This issue is resolved in this release.

- **You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled**

After the hostd service starts, for example after an ESXi host reboot, it might create a SNMP directory under `/var/spool` and you see many `.trp` files to pile up in this directory.

This issue is resolved in this release. The fix makes sure that the directory `/var/spool/snmp` exists only when SNMP is enabled.

- **The hostd service repeatedly fails and the ESXi host disconnects from the vCenter system**

  If for any reason an ESXi host is temporarily in a state of insufficient memory, the hostd service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

  This issue is resolved in this release.

- **Performance of certain nested virtual machines on AMD CPUs might degrade**

  Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

  This issue is resolved in this release.

- **Changing the mode of the virtual disk** on **a running virtual machine might cause the VM to fail**

  If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the vmware.log, you see errors such as:

  ```
  msg.disk.notConfigured2] Failed to configure disk 'scsi0:4'. The virtual machine cannot be powered on with an unconfigured disk.
  ```

  ```
  [msg.checkpoint.continuesync.error] An operation required the virtual machine to quiesce and the virtual machine was unable to continue running.
  ```

  This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

- **ESXi NVMe/TCP initiator fails to recover paths after target failure recovery**

  When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

  This issue is resolved in this release.

- **ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host**

  Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

This issue is resolved in this release.

- **In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs**

  In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (nmlx5) driver might be low and you might see uneven traffic flow across RSS queues.

  This issue is resolved in this release.

- **A Logical Volume Manager (LVM) disk goes offline during datastore expansion**

  If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

  This issue is resolved in this release.

- **The durable name of a SCSI LUN might not be set**

  The durable name property for a SCSI-3 compliant device comes from pages `80h` and `83h` of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

  This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages `80h` and `83h`, if that data is not required for the generation of durable name.

- **Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations**

  vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

  This issue is resolved in this release.

- **Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel**

  In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

This issue is resolved in this release.

- **NVMe over Fabrics controller might unexpectedly disconnect during discovery**

  If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

  This issue is resolved in this release.

- **If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures**

  When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

  This issue is resolved in this release. For more details, see VMware knowledge base article 92276.

- **If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline**

  While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

  This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

- **Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues**

  Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

  This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

  a On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings** > **System Image Cache Configuration** b In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.

c Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

- **VMware VIB installation might fail during concurrent vendor package installations**

  When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration`.

  This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

- **Certain applications might take too many ESXi file handles and cause performance aggravation**

  In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

  This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

ESXi-8.0U1c-22088125-no-tools

Profile Name	ESXi-8.0U1c-22088125-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_vdfs_8.0.1-0.25.22088125 ■ VMW_bootbank_nvmetcp_1.0.1.7-1vmw.801.0.25.22088125 ■ VMW_bootbank_ntg3_4.1.10.0-5vmw.801.0.25.22088125 ■ VMW_bootbank_nmlx5-rdma-esxio_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_gc_8.0.1-0.25.22088125 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.25.22088125 ■ VMware_bootbank_cpu-microcode_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.25.22088125 ■ VMware_bootbank_vds-vsip_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.25.22088125 ■ VMware_bootbank_loadesxio_8.0.1-0.25.22088125 ■ VMware_bootbank_loadesx_8.0.1-0.25.22088125 ■ VMware_bootbank_gc-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-combiner_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-core_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_esxio-update_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-update_8.0.1-0.25.22088125 ■ VMware_bootbank_vsan_8.0.1-0.25.22088125 ■ VMware_bootbank_crx_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-rdma_4.23.0.36-15vmw.801.0.25.22088125 ■ VMware_bootbank_bmcal_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-xserver_8.0.1-0.25.22088125 ■ VMware_bootbank_trx_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio-base_8.0.1-0.25.22088125 ■ VMware_bootbank_clusterstore_8.0.1-0.25.22088125 ■ Mware_bootbank_native-misc-drivers-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_bmcal-esxio_8.0.1-0.25.22088125 ■ VMware_bootbank_esx-base_8.0.1-0.25.22088125 ■ VMware_bootbank_vsanhealth_8.0.1-0.25.22088125 ■ VMW_bootbank_nmlx5-core-esxio_4.23.0.36-15vmw.801.0.25.22088125

PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

This patch updates the following issues:

- **Transient vSAN health check warning: Network configuration is out of sync**

  vSAN Skyline health might randomly report that network configuration is out of sync. This transient issue occurs when the vSAN health service uses an outdated vCenter configuration to perform unicast check.

  This issue is resolved in this release.

- **vSAN cache overwrite issue might cause inconsistent metadata**

  In rare cases, racing between unmap and snapshot operations might cause inconsistent vSAN metadata in a vSAN host. Such inconsistency can lead to unpredictable consequences, including VM failures or host failures.

  This issue is resolved in this release.

- **ESX hosts might fail with a purple diagnostic screen and an error NMI IPI: Panic requested by another PCPU**

  The resource pool cache is a VMFS specific volume level cache that stores the resource clusters corresponding to the VMFS volume. While searching for priority clusters, the cache flusher workflow iterates through a large list of cached resource clusters, which can cause lockup of the physical CPUs. As a result, ESX hosts might fail with a purple diagnostic screen. In the `logDump` file, you see an error such as:

  ```
  ^[[7m2022-10-22T07:56:47.322Z cpu13:2101160)WARNING: Heartbeat: 827: PCPU 0 didn't
  have a heartbeat for 7 seconds, timeout is 14, 1 IPIs sent; *may* be locked
  up.^[[0m
  ```

  ```
  ^[[31;1m2022-10-22T07:56:47.322Z cpu0:2110633)ALERT: NMI: 710: NMI IPI:
  RIPOFF(base):RBP:CS
  ```

This issue is resolved in this release.

- **ESXi ConfigStore database fills up and writes fail**

  Stale data related to block devices might not be deleted in time from the ESXi ConfigStore database and cause an out of space condition. As a result, write operations to ConfigStore start to fail. In the backtrace, you see logs such as:

  ```
  2022-12-19T03:51:42.733Z cpu53:26745174)WARNING: VisorFSRam: 203: Cannot extend
  visorfs file /etc/vmware/configstore/current-store-1-journal because its ramdisk
  (configstore) is full.
  ```

  This issue is resolved in this release.

- **Delay before evacuating vSAN storage device in UNHEALTHY state**

  After detecting devices as UNHEALTHY, the Local Log-Structured Object Manager (LSOM) might wait for 10 minutes (LSOM_DEVICE_MONITORING_INTERVAL) before initiating evacuation on these devices.

  This issue is resolved in this release.

- **vSAN precheck for maintenance mode or disk decommission doesn't list objects that might lose accessibility**

  This issue affects objects with resyncing components, and some components reside on a device to be removed or placed into maintenance mode. When you run a precheck with the **No-Action** option, the precheck does not evaluate the object correctly to report it in the `inaccessibleObjects` list.

  This issue is resolved in this release. Precheck includes all affected objects in the `inaccessibleObjects` list.

- **The vSAN iSCSI Target service might fail due to a rare race condition**

  When you run the ESXCLI command `esxcli network firewall load`, the operation does not reload existing dynamic firewall rules and they are lost. When you run the ESXCLI command `esxcli network firewall refresh`, the operation reloads existing dynamic firewall rules, but in some cases a race condition might cause some rules to be lost. The issue occurs only when multiple firewall refresh commands run at the same time, which leads to the race condition. As a result, the vSAN iSCSI Target service might fail.

  This issue is resolved in this release.

- **If the Internet Control Message Protocol (ICMPA) is not active, ESXi host reboot might take long after upgrading to vSphere 8.0 and later**

  If ICMPA is not active on the NFS servers in your environment, after upgrading your system to vSphere 8.0 and later, ESXi hosts reboot might take an hour to complete, because restore operations for NFS datastores fail. NFS uses the `vmkping` utility to identify reachable IPs of the NFS servers before executing a mount operation and when ICMP is not active, mount operations fail.

This issue is resolved in this release. To remove dependency on the ICMP protocol to find reachable IPs, the fix adds socket APIs to ensure that IPs on a given NFS server are available.

- **When you migrate a VM with recently hot-added memory, an ESXi host might repeatedly fail with a purple diagnostic screen**

  Due to a race condition while the memory hotplug module recomputes the NUMA memory layout of a VM on a destination host after migration, an ESXi host might repeatedly fail with a purple diagnostic screen. In the backtrace, you see errors such as:

  ```
  0x452900262cf0:[0x4200138fee8b]PanicvPanicInt@vmkernel#nover+0x327 stack:
  0x452900262dc8, 0x4302f6c06508, 0x4200138fee8b,
  0x420013df1300, 0x452900262cf0  0x452900262dc0:
  [0x4200138ff43d]Panic_WithBacktrace@vmkernel#nover+0x56 stack: 0x452900262e30,
  0x452900262de0, 0x452900262e40, 0x452900262df0, 0x3e7514  0x452900262e30:
  [0x4200138fbb90]NMI_Interrupt@vmkernel#nover+0x561 stack: 0x0, 0xf48, 0x0,
  0x0, 0x0  0x452900262f00:[0x420013953392]IDTNMIWork@vmkernel#nover+0x7f stack:
  0x420049800000, 0x4200139546dd, 0x0, 0x452900262fd0, 0x0  0x452900262f20:
  [0x4200139546dc]Int2_NMI@vmkernel#nover+0x19 stack: 0x0, 0x42001394e068, 0xf50,
  0xf50, 0x0  0x452900262f40:[0x42001394e067]gate_entry@vmkernel#nover+0x68
  stack: 0x0, 0x43207bc02088, 0xd, 0x0, 0x43207bc02088  0x45397b61bd30:
  [0x420013be7514]NUMASched_PageNum2PhysicalDomain@vmkernel#nover+0x58 stack: 0x1,
  0x420013be34c3, 0x45396f79f000, 0x1, 0x100005cf757  0x45397b61bd50:
  [0x420013be34c2]NUMASched_UpdateAllocStats@vmkernel#nover+0x4b stack:
  0x100005cf757, 0x0, 0x0, 0x4200139b36d9, 0x0  0x45397b61bd80:
  [0x4200139b36d8]VmMem_NodeStatsSub@vmkernel#nover+0x59 stack: 0x39,
  0x45396f79f000, 0xbce0dbf, 0x100005cf757, 0x0  0x45397b61bdc0:
  [0x4200139b4372]VmMem_FreePageNoBackmap@vmkernel#nover+0x8b stack: 0x465ec0001be0,
  0xa, 0x465ec18748b0, 0x420014e7685f, 0x465ec14437d0
  ```

  This issue is resolved in this release.

- **After upgrading the ntg3 driver to version 4.1.9.0-4vmw, Broadcom NICs with fiber physical connectivity might lose network**

  Changes in the ntg3 driver version `4.1.9.0-4vmw` might cause link issues for the fiber physical layer and connectivity on some NICs, such as Broadcom 1Gb, fails to come up.

  This issue is resolved in this release.

- **vSAN cluster shutdown fails on a cluster with IPV6 disabled**

  This issue occurs on vSAN hosts running ESXi 8.0 Update 1 with IPv6 disabled. When you use the vSAN cluster shutdown wizard, the workflow fails with the following error message: `'NoneType' object is not iterable`.

  This issue is resolved in this release.

- **You see trap files in a SNMP directory under /var/spool even though SNMP is not enabled**

After the hostd service starts, for example after an ESXi host reboot, it might create a SNMP directory under `/var/spool` and you see many `.trp` files to pile up in this directory.

This issue is resolved in this release. The fix makes sure that the directory `/var/spool/snmp` exists only when SNMP is enabled.

- **The hostd service repeatedly fails and the ESXi host disconnects from the vCenter system**

  If for any reason an ESXi host is temporarily in a state of insufficient memory, the hostd service might repeatedly fail due to a vSphere Replication filter that prevents the allocation of bitmaps. As a result, the ESXi host disconnects from the vCenter system and cannot connect back.

  This issue is resolved in this release.

- **Performance of certain nested virtual machines on AMD CPUs might degrade**

  Nested virtual machines on AMD CPUs with operational systems such as Windows with virtualization-based security (VBS) might experience performance degradation, timeouts, or unresponsiveness due to an issue with the virtualization of AMD's Rapid Virtualization Indexing (RVI), also known as Nested Page Tables (NPT).

  This issue is resolved in this release.

- **Changing the mode of the virtual disk on a running virtual machine might cause the VM to fail**

  If you use the VMware Host Client to edit the disk mode of a running virtual machine, for example from Independent - Nonpersistent to Dependent or Independent - Persistent, the operation fails and might cause the VM to fail. In the vmware.log, you see errors such as:

  ```
  msg.disk.notConfigured2] Failed to configure disk 'scsi0:4'. The virtual machine
  cannot be powered on with an unconfigured disk.
  ```

  ```
  [msg.checkpoint.continuesync.error] An operation required the virtual machine to
  quiesce and the virtual machine was unable to continue running.
  ```

  This issue is resolved in this release. The fix blocks changing the mode of an Independent - Nonpersistent disk on a running virtual machine by using the VMware Host Client. The vSphere Client already blocks such operations.

- **ESXi NVMe/TCP initiator fails to recover paths after target failure recovery**

  When an NVMe/TCP target recovers from a failure, ESXi cannot recover the path.

  This issue is resolved in this release.

- **ESXi host becomes unresponsive and you cannot put the host in Maintenance Mode or migrate VMs from that host**

  Asynchronous reads of metadata on a VMFS volume attached to an ESXi host might cause a race condition with other threads on the host and make the host unresponsive. As a result, you cannot put the host in Maintenance Mode or migrate VMs from that host.

This issue is resolved in this release.

- **In some cases, you might see low throughput in encapsulated traffic with Mellanox NICs**

  In non Enhanced Datapath mode, the throughput of encapsulated traffic of NICs with the Mellanox (nmlx5) driver might be low and you might see uneven traffic flow across RSS queues.

  This issue is resolved in this release.

- **A Logical Volume Manager (LVM) disk goes offline during datastore expansion**

  If during the expansion of a datastore on an ESXi host in a cluster you run a storage refresh, a LVM extent might go offline and virtual machines on this volume become unresponsive. The issue occurs because the storage refresh operation triggers a slow refresh of the volume attributes on all the ESXi hosts in the cluster. As a result, LVM metadata on the disk might not match the cached capacity information in the Pluggable Storage Architecture (PSA) layer and ESXi marks the LVM extent offline for the safety of metadata and data.

  This issue is resolved in this release.

- **The durable name of a SCSI LUN might not be set**

  The durable name property for a SCSI-3 compliant device comes from pages `80h` and `83h` of the Vital Product Data (VPD) as defined by the T10 and SMI standards. To populate the durable name, ESXi first sends an inquiry command to get a list of VPD pages supported by the device. Then ESXi issues commands to get data for all supported VPD pages. Due to an issue with the target array, the device might fail a command to get VPD page data for a page in the list with a `not supported` error. As a result, ESXi cannot populate the durable name property for the device.

  This issue is resolved in this release. The fix ignores the error on command to get VPD page data, except pages `80h` and `83h`, if that data is not required for the generation of durable name.

- **Auto discovery of NVMe Discovery Service might fail on ESXi hosts with NVMe/TCP configurations**

  vSphere 8.0 adds advanced NVMe-oF Discovery Service support in ESXi that enables the dynamic discovery of standards-compliant NVMe Discovery Service. ESXi uses the mDNS/DNS-SD service to obtain information such as IP address and port number of active NVMe-oF discovery services on the network. However, in ESXi servers with NVMe/TCP enabled, the auto discovery on networks configured to use vSphere Distributed Switch might fail. The issue does not affect NVMe/TCP configurations that use standard switches.

  This issue is resolved in this release.

- **Excessive time for path recovery with ESXi High Performance Plug-in (HPP) after a link up event for NVMe over Fibre Channel**

  In certain scenarios, after a Fibre Channel link up event, it might take as long as 5 minutes for HPP managed NVMe over Fibre Channel paths to recover.

This issue is resolved in this release.

- **NVMe over Fabrics controller might unexpectedly disconnect during discovery**

  If discovery and I/O controllers already exist for a NVMe over Fabrics storage target that supports a persistent discovery controller on an ESXi host, a concurrent NVMe over Fabrics controller discovery operation might cause some I/O controllers to disconnect unexpectedly.

  This issue is resolved in this release.

- **If your ESXi PTP service uses hardware timestamping, enabling IGMP snooping on switches might cause synchronization failures**

  When you enable Internet Group Management Protocol (IGMP) snooping on a connected switch, the Precision Time Protocol (PTP) client needs to send IGMP multicast requests to receive PTP multicast stream from the grandmaster. If the ESXi PTP agent is based on hardware timestamping, the agent might fail to send IGMP join/leave requests to the switch. As a result, the PTP multicast stream cannot go forward to the ESXi host and prevents proper PTP synchronization.

  This issue is resolved in this release. For more details, see VMware knowledge base article 92276.

- **If parallel volume expand and volume refresh operations on the same VMFS volume run on two ESXi hosts in the same cluster, the VMFS volume might go offline**

  While a VMFS volume expand operation is in progress on an ESXi host in a vCenter cluster, if on another host a user or vCenter initiates a refresh of the same VMFS volume capacity, such a volume might go offline. The issue occurs due to a possible mismatch in the device size, which is stamped on the disk in the volume metadata during a device rescan, and the device size value in the Pluggable Storage Architecture (PSA) layer on the host, which might not be updated if the device rescan is not complete.

  This issue is resolved in this release. The fix improves the resiliency of the volume manager code to force a consecutive refresh of the device attributes and comparison of the device sizes again if vCenter reports a mismatch in the device size.

- **Operations with stateless ESXi hosts might not pick the expected remote disk for system cache, which causes remediation or compliance issues**

  Operations with stateless ESXi hosts, such as storage migration, might not pick the expected remote disk for system cache. For example, you want to keep the new boot LUN as LUN 0, but vSphere Auto Deploy picks LUN 1.

  This issue is resolved in this release. The fix provides a consistent way to sort the remote disks and always pick the disk with the lowest LUN ID. To make sure you enable the fix, follow these steps:

  a On the Edit host profile page of the Auto Deploy wizard, select **Advanced Configuration Settings** > **System Image Cache Configuration** b In the **System Image Cache Profile Settings** drop-down menu, select **Enable stateless caching on the host**.

c Edit **Arguments for first disk** by replacing **remote** with **sortedremote** and/or **remoteesx** with **sortedremoteesx**.

- **VMware VIB installation might fail during concurrent vendor package installations**

  When you install update packages from several vendors, such as JetStream Software, Microsoft, and VMware, multiple clients call the same PatchManager APIs and might lead to a race condition. As a result, VMware installation packages (VIBs) might fail to install. In the logs, you see an error such as `vim.fault.PlatformConfigFault`, which is a catch-all fault indicating that some error has occurred regarding the configuration of the ESXi host. In the vSphere Client, you see a message such as `An error occurred during host configuration`.

  This issue is resolved in this release. The fix is to return a `TaskInProgress` warning instead of `PlatformConfigFault`, so that you are aware of the actual issue and retry the installation.

- **Certain applications might take too many ESXi file handles and cause performance aggravation**

  In very rare cases, applications such as NVIDIA virtual GPU (vGPU) might consume so many file handles that ESXi fails to process other services or VMs. As a result, you might see GPU on some nodes to disappear, or report zero GPU memory, or performance degradation.

  This issue is resolved in this release. The fix reduces the number of file handles a vGPU VM can consume.

ESXi-8.0U1sc-22082334-standard

Profile Name	ESXi-8.0U1c-22082334-standard
Build	For build information, see Patches Contained in This Release.
Vendor	VMware, Inc.
Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A

Affected VIBs	■ VMware_bootbank_bmcal-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 ■ VMware_bootbank_vsan_8.0.1-0.20.22082334 ■ VMware_bootbank_bmcal_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-base_8.0.1-0.20.22082334 ■ VMware_bootbank_trx_8.0.1-0.20.22082334 ■ VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_gc_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 ■ VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334 ■ VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 ■ VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 ■ VMware_bootbank_clusterstore_8.0.1-0.20.22082334 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-base_8.0.1-0.20.22082334 ■ VMware_bootbank_vdfs_8.0.1-0.20.22082334 ■ VMware_bootbank_crx_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-update_8.0.1-0.20.22082334 ■ VMware_bootbank_loadesx_8.0.1-0.20.22082334 ■ VMware_bootbank_loadesxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-update_8.0.1-0.20.22082334 ■ VMware_locker_tools-light_12.2.5.21855600-22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149, 3186166
Related CVE numbers	N/A

This patch updates the following issues:

- **ESXi 8.0 Update 1c provides the following security updates:**
  - The Envoy proxy is updated to version v1.23.9.
  - The ESXi userworld libxml2 library is updated to version 2.10.4.
  - The cURL library is updated to version 8.0.1.

    - The Go library is updated to version 1.19.9.
    - The etcd package is updated to 3.4.25.

- Updates the `tools-light` VIB.
    - The following VMware Tools ISO images are bundled with ESXi 8.0 Update 1c:
        - **windows.iso**: VMware Tools 12.2.5 supports Windows 7 SP1 or Windows Server 2008 R2 SP1 and later.
        - **linux.iso**: VMware Tools 10.3.25 ISO image for Linux OS with **glibc** 2.11 or later.

      The following VMware Tools ISO images are available for download:

        - VMware Tools 11.0.6:
            - **windows.iso**: for Windows Vista (SP2) and Windows Server 2008 Service Pack 2 (SP2).
        - VMware Tools 10.0.12:
            - **winPreVista.iso**: for Windows 2000, Windows XP, and Windows 2003.
            - **linuxPreGLibc25.iso**: supports Linux guest operating systems earlier than Red Hat Enterprise Linux (RHEL) 5, SUSE Linux Enterprise Server (SLES) 11, Ubuntu 7.04, and other distributions with **glibc** version earlier than 2.5.
        - **solaris.iso**: VMware Tools image 10.3.10 for Solaris.
        - **darwin.iso**: Supports Mac OS X versions 10.11 and later. VMware Tools 12.1.0 was the last regular release for macOS. Refer VMware knowledge base article 88698 for details.

      Follow the procedures listed in the following documents to download VMware Tools for platforms not bundled with ESXi:

        - Updating VMware Tools
        - VMware Tools for hosts provisioned with Auto Deploy
        - What Every vSphere Admin Must Know About VMware Tools
        - Earlier versions of VMware Tools
        - VMware Tools 12.2.5 Release Notes

ESXi-8.0U1sc-22082334-no-tools

Profile Name	ESXi-8.0U1c-22082334-no-tools
Build	For build information, see Patches Contained in This Release.
Vendor	VMware, Inc.

Release Date	July 27, 2023
Acceptance Level	Partner Supported
Affected Hardware	N/A
Affected Software	N/A
Affected VIBs	■ VMware_bootbank_bmcal-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_native-misc-drivers_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-dvfilter-generic-fastpath_8.0.1-0.20.22082334 ■ VMware_bootbank_vsan_8.0.1-0.20.22082334 ■ VMware_bootbank_bmcal_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-base_8.0.1-0.20.22082334 ■ VMware_bootbank_trx_8.0.1-0.20.22082334 ■ VMware_bootbank_gc-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_gc_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-xserver_8.0.1-0.20.22082334 ■ VMware_bootbank_cpu-microcode_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-combiner-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-dvfilter-generic-fastpath_8.0.1-0.20.22082334 ■ VMware_bootbank_vsanhealth_8.0.1-0.20.22082334 ■ VMware_bootbank_vds-vsip_8.0.1-0.20.22082334 ■ VMware_bootbank_clusterstore_8.0.1-0.20.22082334 ■ VMware_bootbank_native-misc-drivers-esxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-combiner_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-base_8.0.1-0.20.22082334 ■ VMware_bootbank_vdfs_8.0.1-0.20.22082334 ■ VMware_bootbank_crx_8.0.1-0.20.22082334 ■ VMware_bootbank_esx-update_8.0.1-0.20.22082334 ■ VMware_bootbank_loadesx_8.0.1-0.20.22082334 ■ VMware_bootbank_loadesxio_8.0.1-0.20.22082334 ■ VMware_bootbank_esxio-update_8.0.1-0.20.22082334
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149
Related CVE numbers	N/A

This patch updates the following issues:

- **ESXi 8.0 Update 1c provides the following security updates:**
  - The Envoy proxy is updated to version v1.23.9.
  - The ESXi userworld libxml2 library is updated to version 2.10.4.
  - The cURL library is updated to version 8.0.1.
  - The Go library is updated to version 1.19.9.
  - The etcd package is updated to 3.4.25.

ESXi 8.0 U1c - 22088125

Name	ESXi
Version	ESXi80U1c-22088125
Release Date	July 27, 2023
Category	Bugfix
Affected Components	■ ESXi Component ■ ESXi Install/Upgrade Component ■ Mellanox 5th generation NICs (ConnectX and BlueField DPU series) core Ethernet and RoCE Drivers for VMware ESXi ■ Broadcom NetXtreme I ESX VMKAPI ethernet driver ■ VMware NVMe over TCP Driver
PRs Fixed	3239946, 3239870, 3210610, 3228383, 3248279, 3179236, 3183209, 3219196, 3229111, 3221890, 3219889, 3223909, 3225464, 3228391, 3158175, 3217410, 3211624, 3223427, 3223420, 3181600, 3212384, 3221902, 3219121, 3219196, 3222717, 3221598, 3210610, 3181603, 3213042, 3221593, 3210931, 3210931, 3221549, 3161147, 3213110, 3219262, 3118977, 3217167, 3210610, 3210610, 3219971, 3112043, 3218145, 3218218, 3217477, 3214491, 3166665, 3210840, 3210837, 3210956, 3213914, 3212431, 3187725, 3213177, 3185230, 3213207, 3187539, 2625439, 3122074, 3197383, 3187416, 3187420, 3118241, 3176359, 3184331, 3182257, 3187875, 3187547, 3210192, 3180881, 3163270, 3179236, 3157222, 3187709, 3187716, 3187494, 3186022, 3188105, 3183522, 3166280, 3183038, 3183531, 3183526, 3184327, 3166566, 3171992, 3172063, 3159074, 3181553, 3183529, 3146205, 3038908, 3038908, 3153396, 3038908, 3038908, 3038908, 3179111, 3152476, 3187446, 3219191
Related CVE numbers	N/A

ESXi 8.0 U1sc - 22082334

Name	ESXi
Version	ESXi-8.0U1sc-22082334
Release Date	July 27, 2023
Category	Security
Affected Components	■ ESXi Component ■ ESXi Install/Upgrade Component ■ ESXi Tools Component
PRs Fixed	3229052, 3222888, 3184515, 3184505, 3210921, 3219294, 3217139, 3215295, 3184512, 3184517, 3213042, 3184513, 3184506, 3186149, 3186166
Related CVE numbers	N/A

Known Issues

6

Read the following topics next:

- Virtual Machine Management

Virtual Machine Management

- **Setting the screen resolution of a virtual machine does not always work when manually edited into a VMX file**

  If you manually specify the screen resolution of a virtual machine by editing the VMX file, thechange might not take effect.

  Workaround: Use the vSphere Client or the VMware Host Client to edit VM properties. Avoid manual edits of VMX files.

Known Issues from Previous Releases

7

Read the following topics next:

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

Installation, Upgrade, and Migration Issues

- **If you update your vCenter to 8.0 Update 1, but ESXi hosts remain on an earlier version, vSphere Virtual Volumes datastores on such hosts might become inaccessible**

  Self-signed VASA provider certificates are no longer supported in vSphere 8.0 and the configuration option `Config.HostAgent.ssl.keyStore.allowSelfSigned` is set to `false` by default. If you update a vCenter instance to 8.0 Update 1 that introduces vSphere APIs for Storage Awareness (VASA) version 5.0, and ESXi hosts remain on an earlier vSphere and VASA version, hosts that use self-signed certificates might not be able to access vSphere Virtual Volumes datastores or cannot refresh the CA certificate.

  Workaround: Update hosts to ESXi 8.0 Update 1. If you do not update to ESXi 8.0 Update 1, see VMware knowledge base article 91387.

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

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

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

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

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

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

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

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

  Workaround: Use either of these 2 options:

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

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

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

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

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

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

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

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

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

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

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

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

Miscellaneous Issues

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

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

  Workaround: None

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

  Workaround: None.

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

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

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

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

  See this script as an example:

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

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

  ```
  ~# reboot
  ```

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

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

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

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

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

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

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

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

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

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

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

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

  Workaround: See VMware knowledge base article 89638.

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

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

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

Networking Issues

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

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

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

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

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

Workaround: None.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

  Workaround: None.

Storage Issues

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

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

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

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

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

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

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

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

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

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

  Workaround: Retry the snapshot creation task.

vCenter Server and vSphere Client Issues

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

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

  Workaround: Click the refresh icon.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Virtual Machine Management Issues

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

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

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

vSphere Lifecycle Manager Issues

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

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

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

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

  ESXi 8.0 introduces the option to explicitly stage desired state images, which is the process of downloading depot components from the vSphere Lifecycle Manager depot to the ESXi hosts without applying the software and firmware updates immediately. However, staging of images is only supported on an ESXi 8.0 or later hosts. Attempting to stage a vSphere

Lifecycle Manager image on ESXi hosts of version earlier than 8.0 results in messages that the staging of such hosts fails, and the hosts are skipped. This is expected behavior and does not indicate any failed functionality as all ESXi 8.0 or later hosts are staged with the specified desired image.

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

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

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

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

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

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

VMware Host Client Issues

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

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

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

Security Features Issues

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

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

  Workaround: Use RSA key size larger than 2048.