Release Notes for Cisco HX Data Platform, Release 3.5

First Published: 2018-10-16

Last Modified: 2021-07-20

Introduction

Cisco HyperFlex Systems unlock the full potential of hyperconvergence. The systems are based on an end-to-end software-defined infrastructure, combining software-defined computing in the form of Cisco Unified Computing System (Cisco UCS) servers, software-defined storage with the powerful Cisco HX Data Platform, and software-defined networking with the Cisco UCS fabric. Together with a single point of connectivity and hardware management, these technologies deliver a pre-integrated and adaptable cluster that is ready to provide a unified pool of resources to power applications as your business needs dictate.

These release notes pertain to the Cisco HX Data Platform, Release 3.5, and describe the features, limitations and caveats for the Cisco HX Data Platform.

Recent Revisions

For the complete revision history, see .

Release	Date	Description
3.5(2i)	May 7, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.0(4l) qualified for HX 3.5(2i) and 3.5(2h).
3.5(2i)	April 29, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.1(3c) qualified for HX 3.5(2i).
3.5(2i)	March 17, 2021	Updated link to indicate UCSM 4.1(2f) is the recommended Host Upgrade Utility (HUU) for M5 for HX 3.5(2x).
3.5(2i)	March 11, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.0(4k) is the recommended release.
3.5(2i)	February 19, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.1(2c) qualified for HX 3.5(2i) and HX 3.5(2h).

New Features

New Features in Release 3.5(2i)

- There are no new software features in this release.

New Features in Release 3.5(2h)

- **Support for a new 375G Intel Optane high performance cache drive (PID: HX-NVMEXPB-I375) has been added**—This drive was introduced in the HXDP 4.0 release.

New Features in Release 3.5(2g)

- **Support for Second-Generation Intel® Xeon® Scalable Processor Refresh**—This release includes support for the Second-Generation Intel® Xeon® Scalable processor refresh (formerly Cascade Lake).

- **Support for new FIPS Compliant SED SSD drive HX-SD800GBHNK9**—This drive replaces the old caching drive (HX-SD800GBENK9) which is no longer orderable. The cluster needs to be upgraded to HX Release 3.5(2g) or higher to use the new drive for expansion or as a replacement.

- **Support for 1.6TB cache drive HX-SD16T123X-EP for Edge (server PID: HXAF-E-220M5SX)**—This drive replaces the old cache drive option HX-SD400G12TX-EP which is no longer orderable. Please note that the HX Edge cluster will need to be installed with HX Release 3.5(2g) to take advantage of the new drive PID. Also for RMA scenarios the HX Edge cluster will need to be upgraded to HX Release 3.5(2g) or higher.

New Features in Release 3.5(2f)

> **Note** Cisco HyperFlex Release 3.5(2f) requires that stretch clusters upgrade the Witness VM to version 1.0.6; For details on how to upgrade your witness VM, see Upgrading a Witness VM.

> **Note** Encryption with stretched cluster and valid configuration for SED drives is not supported.

New Features in Release 3.5(2e)

- **Support for 1.6TB cache drive HX-SD16T123X-EP for the following server configurations HXAF220C-M5SX and HXAF240C-M5SX**—This drive replaces the old cache drive option HX-SD400G12TX-EP which is no longer orderable. Please note that the cluster will need to be installed with 3.5(2g) to take advantage of the new drive PID. Also for expansion of existing clusters (using HX-SD400G12TX-EP cache drive) with nodes containing the new drive OR when replacing failed drives, the cluster will need to be upgraded to 3.5(2g) or higher.

New Features in Release 3.5(2d)

- There are no new software features in this release.

New Features in Release 3.5(2c)

- This release is deferred.

![Warning]

Warning   Cisco HyperFlex Data Platform release 3.5(2c) is no longer available for download. It is recommended that you upgrade to HX 3.5(2g) or the latest suggested release on the Cisco Software Download site. For more information, refer to the Software Deferral Notice for CSCvp90129 and Software Advisory for CSCvp90129: Stretched cluster nodes that experience failures may become unavailable.

New Features in Release 3.5(2b)

- **EMC RecoverPoint**—Starting with Cisco HX Data Platform release 3.5(2b), support for RecoverPoint for VMs was initiated in conjunction with RP4VMs version 5.2.P1. The synchronous replication feature of RecoverPoint for VMs is not supported.

New Features in Release 3.5(2a)

- **HyperFlex Hardware Acceleration Engine**—New purpose-built PCIe hardware acceleration cards (HX-PCIE-OFFLOAD-1) for higher compression efficiency, speed and increased storage capacity. These cards are currently available with HX240 M5(All Flash/Hybrid) nodes and require an HX Data Platform Enterprise license. For more information, see Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 3.5.

- **Microsoft Hyper-V Enhancements**—Starting with Cisco HX Data Platform release 3.5(2a), cluster expansion for Hyper-V compute-only nodes is supported. For more information, see Cisco HyperFlex Systems Installation Guide for Microsoft Hyper-V, Release 3.5.

- **Citrix Cloud Services**— New capability to connect to Citrix Workspaces and associated Citrix Cloud subscription services such as Citrix Virtual Apps and Desktops Services (formerly XenApp and XenDesktop services). For more information, see Cisco HyperFlex Systems Administration Guide for Citrix Workspace Appliance, Release 3.5.

- **Support for VMware ESXi 6.7 U1**—For more information, see Cisco HyperFlex Systems — Networking Topologies, and Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 3.5.

- **Support for UCS 6454 Fabric Interconnects (HX-FI-6454) with HX Clusters**—For more information, see Cisco HyperFlex Systems — Networking Topologies.

New Features in Release 3.5(1a)

- **Native Disaster Recovery Enhancements**— A simple to use Planned Migration workflow for Disaster Recovery, VM migration, and resuming replication. In addition, native support for Replication, Planned Migration and Disaster Recovery in Stretched Cluster deployments. For more information, see Cisco HyperFlex Systems Administration Guide, Release 3.5

- **HX Data Platform Installer Enhancements**—New extended capabilities in hardening and reliability in the HX Data Platform Installer.

  - Cluster Expansion for Hyper-V converged nodes.

  - Cluster expansion for Stretched Cluster compute-only and converged nodes.

- Integrated Hyper-V and Windows Server OS bare metal installation included as part of cluster creation workflow.

- **Networking Enhancements**—Support for multi-VIC network designs and third-party NIC for HX converged and compute-only nodes. For more information, see Cisco HyperFlex Systems — Networking Topologies

- **Upgrade Enhancements**—Support for orchestrated ESXi hypervisor upgrades. Combined with existing support for HX Data Platform and server firmware upgrades, this release provides the ability to perform seamless full stack upgrades all orchestrated through HX Connect.

  Starting with Cisco HX Data Platform release 3.5(1a) and later, all future upgrades can be completed in HX Connect UI. For all future upgrades, this functionality will take affect for all clusters on release 3.5. To upgrade to the 3.5 release from older versions, continue to run the bootstrap script as outlined in the documentation. This new end-to-end UI-based upgrade capability will be utilized on all subsequent upgrades. For more information, see Cisco HyperFlex Systems Upgrade Guide, Release 3.5

- **ESXi Lockdown Mode**—Support for VMware ESXi lockdown mode to increase security of an ESXi host by limiting access allowed for the host. When enabled, the ESXi host can only be accessed through vCenter Server or Direct Console User Interface (DCUI). For more information, Cisco HyperFlex Systems Installation Guide, Release 3.5

- **HX Edge 10GbE Edge Network option**—New 10GbE Edge support provides an additional fully redundant, high speed connectivity option for HyperFlex Edge clusters. For more information, see Cisco HyperFlex Systems Edge Deployment Guide, Release 3.5

- **Cisco Container Platform (CCP) and Open Shift Platform integration (OpenShift)**—Storage integration with Kubernetes that enables dynamic (on-demand) persistent volumes from HyperFlex. This feature is supported with OpenShift (version 3.10) and Cisco Container Platform (CCP). For more information, see Cisco HyperFlex Systems Kubernetes Administration Guide, Release 3.5

- **Artificial intelligence and machine learning (AI/ML) Workloads on HyperFlex with NVIDIA V100 GPUs**—Ability to create applications for AI/ML with NVIDIA Tesla V100 GPUs integration within HyperFlex nodes. For more information, see Cisco HyperFlex HX-Series Spec Sheets

- **Permanent License Reservation (PLR)** – This feature is designed for highly secure intelligence, air-gapped and military environments where external communication may be limited. For more information, see Cisco HyperFlex Systems Ordering and Licensing Guide

- **DISA STIG Automation**—Enhance the security posture of HyperFlex converged and compute-only nodes by automating the implementation of the Defense Information Systems Agency's (DISA) recommended Security Technical Implementation Guides (STIGs), pertaining to VMware vSphere

- **Tech Support Mode** —Enhance the security posture of HyperFlex converged nodes by disabling Tech Support Mode, which disables remote access to Controller VMs over SSH.

Supported Versions and System Requirements

Cisco HX Data Platform requires specific software and hardware versions, and networking settings for successful installation.

For a complete list of requirements, see:

- Cisco HyperFlex Systems Installation Guide for VMware ESXi, or

- Cisco HyperFlex Systems Installation Guide for Microsoft Hyper-V

Hardware and Software Interoperability

For a complete list of hardware and software inter-dependencies, refer to respective Cisco UCS Manager release version of Hardware and Software Interoperability for Cisco HyperFlex HX-Series.

Recommended FI/Server Firmware - 3.5(x) Releases

The HX components—Cisco HX Data Platform Installer, Cisco HX Data Platform, and Cisco UCS firmware—are installed on different servers. Verify that each component on each server used with and within an HX Storage Cluster are compatible.

- HyperFlex does not support UCS Manager and UCS Server Firmware versions 4.0(4a), 4.0(4b), and 4.0(4c).

> ☞
>
> **Important**   Do not upgrade to these versions of firmware.
>
> Do not upgrade to these versions of UCS Manager.

- Verify that the preconfigured HX servers have the same version of Cisco UCS server firmware installed. If the Cisco UCS Fabric Interconnects (FI) firmware versions are different, see the *Cisco HyperFlex Systems Upgrade Guide* for steps to align the firmware versions.

  - **M4**: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M4 or HX220c M4) deployments, verify that Cisco UCS Manager 3.1(3k), 3.2(3i), or 4.0(2d) is installed.

  - **M5**: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M5 or HX220c M5) deployments, verify that the recommended UCS firmware version is installed.

> ☞
>
> **Important**   If you are upgrading Cisco UCS Manager 4.0(2a) or 4.0(2b) in the presence of more than one Nvidia GPUs, please remove the GPUs, perform the upgrade and reinstall. For more details, see CSCvo13678.

> ☞
>
> **Important**   For SED-based HyperFlex systems, ensure that the A (Infrastructure), B (Blade server) and C (Rack server) bundles are at Cisco UCS Manager version 4.0(2b) or later for all SED M4/M5 systems. For more details, see CSCvh04307. For SED-based HyperFlex systems, also ensure that all clusters are at HyperFlex Release 3.5(2b) or later. For more information, see **Field Notice (70234)** and CSCvk17250.

- To reinstall an HX server, download supported and compatible versions of the software. See the Cisco HyperFlex Systems Installation Guide for VMware ESXi, Release 3.5 for the requirements and steps.

Table 1: HyperFlex Software Versions for M4/M5 Servers

HyperFlex Release	M4 Recommended FI/Server Firmware *(be sure to review important notes above)	M5 Recommended FI/Server Firmware *(be sure to review important notes above)	M4/M5 Qualified FI/Server Firmware *(be sure to review important notes above)
3.5(2i)	4.0(4k)	4.0(4k)	4.0(4k), 4.0(4l), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c), 4.1(3b), 4.1(3c)
3.5(2h)	4.0(4k)	4.0(4k)	4.0(4k), 4.0(4l), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c), 4.1(3b)
3.5(2g)	4.0(4k)	4.0(4k)	4.0(4h), 4.1(1d), 4.1(1e), 4.1(3b)
3.5(2f)	4.0(4e)	4.0(4e)	4.0(4d)[1], 4.0(4e)[2], 4.1(3b)
3.5(2e)	4.0(4e)	4.0(4e)	4.0(4g), 4.1(3b)
3.5(2d)	4.0(4e)	4.0(4e)	4.1(3b)
3.5(2c)	Release Deferred		
3.5(2b)	4.0(2d), 3.2(3i), 3.1(3k)	4.0(2d)	4.1(3b)
3.5(2a)	4.0(1c), 3.2(3i), 3.1(3k)	4.0(1c)	4.1(3b)
3.5(1a) - Unsupported	4.0(1b), 3.2(3h), 3.1(3j)	4.0(1a)	

[1] 4.0(4d) qualified only for M5.
[2] 4.0(4e) qualified only for M5.

*UCS Server Firmware 4.1(2a) and 4.1(2b) are not supported on clusters with self-encrypting drives (SED). For more information, see CSCvv69704.

☞

Important If your cluster is connected to a Fabric Interconnect 6400 series using VIC 1455/1457 with SFP-H25G-CU3M or SFP-H25G-CU5M cables, only use UCS Release 4.0(4k) and later, or 4.1(2a) and later. Do not use the any other UCS version listed in the table of qualified releases. Using a UCS Release that is not UCS Release 4.0(4k) and later, or 4.1(2a) and later may cause cluster outages.

Fore more information, see the Release Notes for UCS Manager, Firmware/Drivers, and Blade BIOS for any UCS issues that affect your environment and CSCvu25233.

NOTE: If your current server firmware version is not on the recommendation list above, follow the upgrade procedure in the Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Known Issues chapter.

HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments

Cisco HX Data Platform, Release 3.x based Deployments

Confirm the component firmware on the server meets the minimum versions listed in the following tables.

☞

Important **HyperFlex Edge does not support Cisco IMC versions 4.0(4a), 4.0(4b), 4.0(4c), 4.0(4d), and 4.0(4e).**

Table 2: HX220c M4 / HXAF220c M4 Cluster

Component	Minimum Firmware Version - HXDP 3.x *(be sure to review important note(s) above)*	Recommended Firmware Version - HXDP 3.x *(be sure to review important note(s) above)*
Cisco Integrated Management Controller (CIMC)	3.0(3f)	4.0(2f)
Host Upgrade Utility (HUU) Download Link	3.0(3f) Download Software	4.0(2f) Download Software

Table 3: HX220c M5 / HXAF220c M5 Cluster

Component	Minimum Firmware Version - HXDP 3.x *(be sure to review important note(s) above)*	Recommended Firmware Version - HXDP 3.x *(be sure to review important note(s) above)*
Cisco Integrated Management Controller (CIMC)	3.1(2d)	4.1(2f)
Host Upgrade Utility (HUU) Download Link	3.1(2d) Download Software	4.1(2f) Download Software

HyperFlex Licensing

Beginning with Cisco HyperFlex Release 2.6(1a), HyperFlex supports VMware PAC licensing. Existing VMware embedded licenses will continue to be supported.

Beginning with Cisco HyperFlex Release 2.5(1a), HyperFlex uses a smart licensing mechanism to apply your licenses. See the *Cisco HyperFlex Systems Installation Guide for VMware ESXi* for details and steps.

VMware vSphere Licensing Requirements

How you purchase your vSphere license determines how your license applies to your HyperFlex system.

- **If you purchased your vSphere license with HyperFlex**

  Each HyperFlex server either has the Enterprise or Enterprise Plus edition preinstalled at the factory.

> **Note**
> - HX Nodes have OEM licenses preinstalled. If you delete or overwrite the content of the boot drives after receiving the HX servers, you also delete the factory-installed licenses.
>
> - OEM license keys is a new VMware vCenter 6.0 U1b feature. Earlier versions do not support OEM licenses.
>
> - All factory-installed HX nodes share the same OEM license key. With vSphere OEM keys, the `Usage` count can exceed the `Capacity` value.
>
> - When you add an HX host to vCenter through the **Add Host** wizard, in the **Assign license** section, select the **OEM license**.
>
>   We obfuscate the actual vSphere OEM license key; for example, 0N085-XXXXX-XXXXX-XXXXX-10LHH.
>
> - Standard, Essentials Plus, and ROBO editions are not available preinstalled on HX servers.

- **If you did NOT purchase your vSphere license with HyperFlex**

  The HX nodes have a vSphere Foundation license preinstalled. After initial setup, you can apply the license to a supported version of vSphere.

- **If you purchased a vSphere PAC license**

  Follow the instructions in your PAC license letter from VMware to add the license to your MY VMware account, then follow the instructions to add your HX host to vCenter and assign the PAC license.

HX Data Platform Software Versions for HyperFlex Witness Node

Table 4: HX Data Platform Software Versions for HyperFlex Witness Node

HyperFlex Release	Witness Node Version
3.5(2i)	1.0.9
3.5(2h)	1.0.8
3.5(2g)	1.0.6
3.5(2f)	1.0.6
3.5(2e)	1.0.4
3.5(2d)	1.0.3
3.5(2c)	Release Deferred
3.5(2b)	1.0.3
3.5(2a)	1.0.3

HyperFlex Release	Witness Node Version
3.5(1a) - Unsupported	1.0.2

Note Cisco HyperFlex Release 3.5(2f) requires that stretch clusters upgrade the Witness VM to version 1.0.6. For details on how to upgrade the Witness VM, see Upgrading a Witness VM.

Note Older versions of witness VMs are supported when the cluster is upgraded to the latest HXDP version.

Software Requirements for VMware ESXi - 3.5(x) Releases

The software requirements include verification that you are using compatible versions of Cisco HyperFlex Systems (HX) components and VMware vSphere, VMware vCenter, and VMware ESXi.

- Verify that all HX servers have a compatible version of vSphere preinstalled.

- Verify that the vCenter version is the same or later than the ESXi version.

- Verify that the vCenter and ESXi versions are compatible by consulting the VMware Product Interoperability Matrix. Newer vCenter versions may be used with older ESXi versions, so long as both ESXi and vCenter are supported in the table below.

- Verify that you have a vSphere administrator account with root-level privileges and the associated password.

The following table applies to the following VMware vSphere Editions: Enterprise, Enterprise Plus, Standard, Essentials Plus, ROBO.

Note Any other licensed editions of VMware vSphere not listed above are not supported, including Essentials Edition.

Table 5: Software Requirements for VMware ESXi

HyperFlex Version	VMware ESXi Versions	VMware vCenter Versions[3]
3.5(2i)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 - See limitations:[4]	6.0 U3, 6.5 U3, 6.7 U3
3.5(2h)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 - See limitations:4	6.0 U3, 6.5 U3, 6.7 U3
3.5(2g)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 - See limitations:4	6.0 U3, 6.5 U3, 6.7 U3
3.5(2f)	6.0 U3, 6.5 U2, 6.7 U2[5]	6.0 U3, 6.5 U2, 6.7 U2

HyperFlex Version	VMware ESXi Versions	VMware vCenter Versions[3]
3.5(2e)	6.0 U3, 6.5 U2, 6.7 U2[6]	6.0 U3, 6.5 U2, 6.7 U2
3.5(2d)	6.0 U3, 6.5 U2, 6.7 U2[7]	6.0 U3, 6.5 U2, 6.7 U2
3.5(2c)	Release Deferred	
3.5(2b)	6.0 U3, 6.5 U2, 6.7 U1 [8], 6.7 U2[9] [10]	6.0 U3, 6. 5U2, 6.7 U1, 6.7 U2
3.5(2a)	6.0 U3, 6.5 U2, 6.7 U1[11]	6.0 U3, 6.5 U2, 6.7 U1
3.5(1a) - Unsupported	6.0 U3, 6.5 U1, 6.5 U2	6.0 U3, 6.5 U1, 6.5 U2, 6.7 U1[12]

[3] vCenter Server 7.0 and vCenter Server 7.0 U1 are not supported and contain a software defect that may impact HyperFlex cluster operations. For further information, see Cisco field notice FN70620.

[4] ESXi 6.7 U3 P04 (Build 17167734) or later is not currently supported with HXDP 3.5(2). For further details, see SSH Incompatibility with ESXi 6.7P04 Tech Note.

[5] Use of 6.7 U2 is not recommended, see Software Advisory for HX Release 3.5(2f) for further details.

[6] Use of 6.7 U2 is not recommended, see Software Advisory for HX Release 3.5(2e) for further details.

[7] Use of 6.7 U2 is not recommended, see Software Advisory for HX Release 3.5(2d) for further details.

[8] Use of 6.7 U1 is not recommended, see software advisory for CSCvo56350 for further details.

[9] For the 3.5(2b) release, cluster installation must first be performed with 6.0 U3 or 6.5U2. After cluster deployment, you may upgrade to 6.7 U2 using the Zip bundle available for download on Cisco.com. See the Upgrade Guide for detailed ESXi upgrade procedures. Existing clusters upgraded to 3.5(2b) may upgrade ESXi to 6.7 U2 at any time.

[10] Use of 6.7 U2 is not recommended, see Software Advisory for HX Release 3.5(2b) for further details.

[11] Use of 6.7 U1 is not recommended, see software advisory for CSCvo56350 for further details.

[12] For the 3.5(1a) release, use of vCenter 6.7U1 is supported only with the ESXi 6.0 and 6.5 versions listed.

Note For vSphere 6.0 users. VMware's last day of general support for vSphere 6.0 occurred on March 12, 2020. HXDP will continue to support vSphere 6.0 U3 on both 3.5(2) and 4.0(2) long lived releases. However, no bug or security fixes will be provided by VMware or Cisco for ESXi going forward due to reaching the last day of support. Cisco TAC will continue to support customers to the best of their ability on ESXi 6.0 builds that have already been released. Cisco strongly recommends upgrading as soon as possible to a supported VMware vSphere 6.5 or 6.7 release and follow Cisco's recommendations as outlined in Recommended Cisco HyperFlex HX Data Platform Software Releases - for Cisco HyperFlex HX-Series Systems.

Software Requirements for Microsoft Hyper-V

The software requirements include verification that you are using compatible versions of Cisco HyperFlex Systems (HX) components and Microsoft Hyper-V (Hyper-V) components.

HyperFlex Software versions

The HX components—Cisco HX Data Platform Installer, Cisco HX Data Platform, and Cisco UCS firmware—are installed on different servers. Verify that each component on each server used with and within the HX Storage Cluster are compatible.

- **Cisco HyperFlex M5 Converged nodes**— For Hybrid (Cisco HyperFlex HX240c M5, HX220c M5) and All Flash (Cisco HyperFlex HXAF240c M5, HXAF220c M5) verify that Cisco UCS Manager 4.0(2b) is installed. For detailed information on installation requirements and steps, see the *Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V*.

Table 6: Supported HyperFlex Software versions for M5 Servers on Hyper-V

HyperFlex Release	M5 Recommended Server Firmware
3.5(2i)	4.0(4i)
3.5(2h)	4.0(4i)
3.5(2g)	4.0(4i)
3.5(2f)	4.0(4e)
3.5(2e)	4.0(4e)
3.5(2d)	4.0(4e)
3.5(2c)	Release Deferred
3.5(2b)	4.0(2b)
3.5(2a)	4.0(1d)
3.5(1a) - Unsupported	4.0(1a)

☞

Important   If your configuration is a Fabric Interconnect 6400 connected to VIC 1455/1457 using SFP-H25G-CU3M or SFP-H25G-CU5M cables, then do not use the recommended UCS version of 4.0(4i) release. You must use UCS release 4.1(2a) with a qualified HXDP 3.5 or 4.0 version or the cluster may experience an outage.

Refer to Release Notes for UCS Manager, Firmware/Drivers, and Blade BIOS for any UCS issues that may affect your environment.

Table 7: Supported Microsoft Software versions

Microsoft Component	Version
Windows Operating System (Windows OS)	Windows Server 2016 Core and Desktop Experience
	Note: For Windows Server 2016 Datacenter Core and Desktop Experience, the Windows 2016 ISO image should be Update Build Revision (UBR) 1884 at a minimum.
	OEM activated ISO and Retail ISOs are currently not supported. Earlier versions of Windows Server such as Windows 2012r2 are not supported.
	Non-English versions of the ISO are currently not supported.
Active Directory	A Windows 2012 or later domain and forest functionality level.

Supported Microsoft License Editions

The Microsoft Windows Server version that is installed on one or more HyperFlex hosts must be licensed as per Microsoft licensing requirements listed on Microsoft Licensing.

Browser Recommendations - 3.5(x) Releases

Use one of the following browsers to run the listed HyperFlex components. These browsers have been tested and approved. Other browsers might work, but full functionality has not been tested and confirmed.

Table 8: Supported Browsers

Browser	Cisco UCS Manager	HX Data Platform Installer	HX Connect
Microsoft Internet Explorer	9 or later	11 or later	11 or later
Google Chrome	14 or later	56 or later	56 or later
Mozilla Firefox	7 or later	52 or later	52 or later

Notes

- **Cisco HyperFlex Connect**:

  The minimum recommended resolution is 1024 X 768.

- **Cisco HX Data Platform Plug-In**:

  The **Cisco HX Data Platform Plug-In** runs in vSphere. For VMware Host Client System browser requirements, see the VMware documentation.

  The **Cisco HX Data Platform Plug-In** is not displayed in the vCenter HTML client. You must use the vCenter flash client.

- **Cisco UCS Manager**:

  The browser must support the following:

    - Java Runtime Environment 1.6 or later.

    - Adobe Flash Player 10 or higher is required for some features.

  For the latest browser information about Cisco UCS Manager, refer to the most recent Cisco UCS Manager Getting Started Guide.

Cisco HX Data Platform Compatibility and Scalability Details - 3.5(x) Releases

Cluster Limits

- Cisco HX Data Platform supports up to 100 clusters managed per vCenter as per VMware configuration maximums.

- Cisco HX Data Platform supports any number of clusters on a single FI domain. Each HX converged node must be directly connected to a dedicated FI port on fabric A and fabric B without the use of a FEX. C-series compute only nodes must also directly connect to both FIs. B-series compute only nodes will connect through a chassis I/O module to both fabrics. In the end, the number of physical ports on the FI

will dictate the maximum cluster size and maximum number of individual clusters supported in a UCS domain.

- Using a FEX on uplink ports connecting the Fabric Interconnects to the top of rack (ToR) switches is not supported due to the possibility of network oversubscription leading to the inability to handle HyperFlex storage traffic during failure scenarios.

The following table provides Cisco HX Data Platform Compatibility and Scalability Details.

Table 9: Cisco HX Data Platform Storage Cluster Specifications

Node	VMware ESXi				Microsoft Hyper-V		Stretched Cluster* (Available on ESX Only)		
HX Servers	HX220c M5 HX220c AF M5 HX240c M5 HX240c AF M5 HX220c M4 HX220c AF M4 HX240c M4 HX240c AF M4	HX240c M5L	HX220c M5 Edge HXAF220c M5 Edge HX220c M4 Edge HXAF220c M4 Edge	HXAF240c M5N All NVMe - HXAF220c M5N Not supported with Hyper-V.	HX220c M5 HX220c AF M5 HX240c M5 HX240c AF M5	HX240c M5L	HX220c M5 HX220c AF M5 HX240c M5 HX240c AF M5	HX240c M5L	All NVMe - HXAF220c M5N

Node	VMware ESXi				Microsoft Hyper-V		Stretched Cluster* (Available on ESX Only)		
Compute Only UCS B Series C Series Servers	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5	—	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5	C240 M5, C220 M5, B200 M4, B200 M5	C220 M5,C240 M5, B200 M4, B200 M5	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5,	B200 M5/M4/M3, B260 M4, B420 M4, B460 M4, B480 M5, C220 M5/M4/M3, C240 M5/M4/M3, C460 M4, C480 M5,
Supported Nodes	Converged and Compute only nodes	Converged and Compute only nodes	Converged nodes only	Converged and Compute only nodes	Converged and Compute only nodes	Converged and Compute only nodes	Converged and Compute only nodes	Converged and Compute only nodes	Converged and Compute only nodes
HXDP-S Licensed Node Limits 1:1 ratio of HXDP-S to Compute only nodes (Min—Max)	Converged nodes:3-32 (7.6TB drive configs on HX 240c AF M5 require HX 4.0(2c) release for full scale) Compute only nodes: 0-32 Compute only nodes: 0-32	Converged nodes:3-16 (12TB drive is limited to a maximum of 8 nodes) Compute only nodes: 0-16 (12TB drive is limited to a maximum of 8 nodes)	M4 Converged nodes: 3 M5 Converged nodes: 2,3,or 4 Requires HXDP-E License	N/A (requires Enterprise HXDP-P License)	Converged nodes: 3-16 Compute only nodes: 0-16	Converged nodes: 3-16 (12TB HDD option is not supported for HyperV) Compute only nodes: 0-16	N/A (requires Enterprise HXDP-P License)	N/A (requires Enterprise HXDP-P License)	N/A (requires Enterprise HXDP-P License)

Node	VMware ESXi				Microsoft Hyper-V		Stretched Cluster* (Available on ESX Only)		
HXDP-P Licensed Node Limits **1:2 ratio of HXDP-P to Compute only nodes** **(Min—Max)**	Converged nodes:3-32 (7.6TB drive configs on HX 240c AF M5 require HX 4.0(2c) release for full scale) Compute only nodes: 0-32 (up to max cluster size) Compute only nodes: 0-32 (up to max cluster size)	Converged nodes:3-16 (12TB drive is limited to a maximum of 8 nodes) Compute only nodes: 0-32 (12TB drive is limited to a maximum of 16 nodes)	Converged nodes: 3 (requires HXDP-E License)	Converged nodes: 3-32 Compute Only nodes: 0-32 (up to max cluster size)	Converged nodes: 3-16 Compute only nodes: 0-16	Converged nodes: 3-16 (12TB HDD option is not supported for HyperV) Compute only nodes: 0-16	Converged nodes: 2-16 per Site Compute only nodes: 0-21 per Site (up to max cluster size) 7.6TB drive configs on HX240c AF M5 require HX 4.0(2c) release for full scale) Compute only nodes: 0-32 (up to max cluster size)	Converged nodes: 2-8 per Site Compute only nodes: 0-16 per Site (up to max cluster size)	Converged nodes: 2-16 per Site Compute only nodes: 0-21 per Site (up to max cluster size)
Max Cluster Size	64	48	3	64	32	32	32 per Site/ 64 per cluster	24 per Site/ 48 per cluster	32 per Site/ 64 per cluster
Max Compute to Converged ratio	2:1*	2:1*	—	2:1*	1:1	1:1	2:1*	2:1*	2:1*
Expansion	✓	✓	No	✓	✓	✓	✓**	✓**	✓**

* Requires Enterprise license

** Requires uniform expansion across both sites

Guidelines and Limitations

- **HX REST API Access Token Management** – Applications leveraging HX REST APIs should re-use access tokens when making API calls. Once obtained using the AAA Obtain Access Token API, access tokens are valid for 18 days (1,555,200 seconds). In a 15 minute window, /auth should be invoked (successfully) a maximum of 5 times only and a user should create a maximum of 8 unrevoked tokens. For more information, see Cisco HyperFlex Systems REST API Reference guide.

Upgrade Guidelines

The following list is a highlight of critical criteria for performing an upgrade of your HyperFlex system.

Guidelines for All HXDP 3.5 Upgrades

- **Upgrade Considerations for configurations using SFP-H25G-CU3M or SFP-H25G-CU5M cables**— If your configuration is a Fabric Interconnect 6400 connected to VIC 1455/1457 using SFP-H25G-CU3M or SFP-H25G-CU5M cables, then do not use the recommended UCS version of 4.0(4i) release. You must use UCS release 4.1(2a) with a qualified HXDP 3.5 or 4.0 version or the cluster may experience an outage. For information on any UCS issues that may affect your environment, see Release Notes for UCS Manager, Firmware/Drivers, and Blade BIOS.

- **Hypercheck Health Check Utility**— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and will help ensure a seamless upgrade experience. For more information see the Hyperflex Health & Pre-Upgrade Check Tool TechNote for full instructions on how to install and run Hypercheck.

- **CPU upgrade from First Generation Intel Xeon Scalable Processors to Second Generation Intel Xeon Scalable Processors is not supported**

- —In place CPU upgrade or swap from First Generation Intel Xeon Scalable Processors to Second Generation Intel Xeon Scalable Processors (for example, HX-CPU-6148 to HX-CPU-I6248) is not supported for HX converged nodes.

- **Cluster Readiness for HX Release 3.5(2a) and later**— Upgrades in HX 3.5(2a) and later are automatically bootstrapped using the HX Connect UI. For more information, see the Cisco HyperFlex Systems Upgrade Guide.

- **Upgrade to the latest Witness VM for Stretch Clusters only**—For customers implementing stretch clusters on Cisco HX Data Platform release 3.5(2f), upgrading to the latest Witness VM - 1.0.6 is mandated.

- **Required vCenter upgrade**—For enhanced security, Cisco HX Data Platform release 3.5(1a) or later requires the use of TLS 1.2. Therefore, vCenter must be upgraded to 6.0 U3f or later before upgrading to HX 3.5. In addition, ESXi should be upgraded as required to meet HX Data Platform compatibility requirements.

- **Complete your Upgrade**—The self-healing (or rebalance) capability is disabled temporarily during the upgrade window; If the upgrade fails, you should complete the upgrade as soon as possible.

- **Unsupported Self-Encrypting Drives (SEDs)**—If adding or replacing self-encrypting drives (SEDs) that have been recently qualified in newer versions of HX Data Platform, insert the new drives only after upgrading HX Data Platform to a compatible version.

- **Maintenance Window**—If upgrading both HX Data Platform and UCS firmware, either a combined or split upgrade can be selected through the vSphere HX Data Platform Plug-in depending on the length of

the maintenance window. Cisco UCS Manager infrastructure upgrade is only supported using AutoInstall and the direct server firmware upgrade should be performed only through the upgrade orchestration framework provided by the HX Data Platform Plug-in.

- **M4 Server Firmware Upgrades**—Server firmware should be upgraded to ensure smooth operation and to correct known issues. Specifically, newer SAS HBA firmware is available in this release and is recommended for long-term stability.

> **Note**
> - Users are encouraged to upgrade to Release 3.1(3c) C-bundle or later whenever possible.
>
> - Users running C-bundle versions before 3.1(2f) must upgrade server firmware by performing a combined upgrade of UCS server firmware (C-bundle) to 3.1(3c) or later and HX Data Platform to 2.5. Do not split the upgrade into two separate operations.
>
> - If the cluster is already on 3.1(2f) C-bundle or later, you may perform an HX Data Platform only or combined upgrade, as required.

- **M5 Server Firmware Upgrades**—M5 generation servers must run firmware version 3.2(2d) or later.

- **Firmware Downgrades** — Downgrading UCSM from the HX-installer is not supported.

- Uplinks from the UCS Fabric Interconnects to all top of rack switch ports must configure spanning tree in **edge trunk** or **portfast edge** mode depending on the vendor and model of the switch. This extra configuration ensures that when links flap or change state, they do not transition through unnecessary spanning tree states and incur an extra delay before traffic forwarding begins. Failure to properly configure FI uplinks in **portfast edge** mode may result in network and cluster outages during failure scenarios and during infrastructure upgrades that leverage the highly available network design native to HyperFlex.

Additional Guidelines for Upgrading an Unsupported HX Release to HXDP 3.5

- **Upgrade Reminder for HyperFlex Clusters Running Versions 3.0(1x) or 3.5(1x)**—HyperFlex versions 3.0(1x) and 3.5(1x) are unsupported and have been declared end-of-life as documented in the End-of-Life notice. For more information see Software Advisory for CSCvt22244.

- **Minimum HXDP version for upgrade**—HX Data Platform clusters running 2.1(1x) or later may upgrade directly to 3.5 using the vCenter plug-in.

- **Initiating Upgrade**—Use the HX Connect UI or CLI `stcli` commands when upgrading from 2.5(1a) or later releases. Use either the CLI `stcli` commands or the HX Data Platform Plug-in to the vSphere Web Client when upgrading from a pre-2.5(1a) release. The vCenter plug-in should not be used for upgrades starting with the 2.5(1a) release.

- **Minimum HXDP version for upgrade**—HX Data Platform clusters running 1.8(1f) or later may upgrade directly to 3.0.

- **HX Data Platform 1.7.x, 1.8.x, 2.0 and 2.1x clusters**—Users from any version prior to 2.6(1a) must step through an intermediate version before upgrading to 3.5x or later releases. If you need to upgrade your environment from a Cisco HyperFlex HX Data Platform software release that is past the last date of support, to the latest suggested release on the Cisco Software Download site, see Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases. For more information, see the Software

Advisory for CSCvq66867: WARNING: Only Use HXDP 2.6(1e) Upgrade Package When Upgrading From HXDP 1.8(1a)-1.8(1e).

- **HX Data Platform 2.6(1x) and higher to 3.5(2x) clusters: Direct upgrade to 3.5(2x) is supported**—Users from any version prior to 2.6(1x) must step through an intermediate version before upgrading to 3.5(2x) or later releases. If you need to upgrade your environment from a Cisco HyperFlex HX Data Platform software release that is past the last date of support, to the latest suggested release on the Cisco Software Download site, see Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases. For more information, see the Software Advisory for CSCvq66867: WARNING: Only Use HXDP 2.6(1e) Upgrade Package When Upgrading From HXDP 1.8(1a)-1.8(1e).

- **Required vCenter upgrade**—For enhanced security, Cisco HX Data Platform release 3.0(1a) and later requires the use of TLS 1.2. Therefore, vCenter must be upgraded to 6.0 U3c or later before upgrading to Cisco HX Data Platform release 3.0. In addition, ESXi should be upgraded as required to meet HX Data Platform compatibility requirements.

- **Cluster Readiness**—Ensure that the cluster is properly bootstrapped and the updated plug-in is loaded before proceeding. Manual cluster bootstrap is required for HX releases earlier than 3.5(1a). For more information, see the Manual Bootstrap Upgrade Process in the Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 3.5. Do not skip this cluster bootstrap step, it is required for all upgrades until HX Release 3.5(1a). Auto bootstrap is supported beginning with HX release 3.5(1a). For more information, see the Auto Bootstrap Upgrade Process from HX Connect UI in the Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 3.5.

- **vSphere 5.5 Upgrades**—Users on vSphere 5.5 must upgrade to 6.0 U3/6.5 U1 before starting HX Data Platform upgrade. vSphere 5.5 support was deprecated with HX Data Platform 2.5(1a) and upgrade fails if attempted.

  - For HX220 users running 5.5, contact TAC for upgrade assistance.

  - For HX240 users running 5.5, upgrade components in the following order.

    1. Upgrade vCenter to 6.0 U3f or later. If upgrading to 6.5, you must upgrade your vCenter in place. Using a new vCenter 6.5 is not supported for users migrating from 5.5.

    2. Upgrade ESXi to 6.0/6.5 using the offline zip bundle.

    Note During upgrade, it might be necessary to reconnect ESXi host manually in vCenter after ESXi upgrade and host reboot.

    3. Upgrade HX Data Platform (and optionally the UCS firmware).

- **If Upgrading to vSphere 6.5**:

  - Certain cluster functions such as native and scheduled snapshots, ReadyClones, and Enter or Exit HX Maintenance Mode will not operate from the time the upgrade is started until the HX Data Platform upgrade to 3.5 or later is complete.

  - After upgrading ESXi using the offline zip bundle, use the ESX Exit Maintenance Mode option. The HX Exit Maintenance Mode option does not operate in the vSphere Web Client until the HX Data Platform upgrade is complete.

- **vSphere 6.0** VMware's last day of general support for vSphere 6.0 occurred on March 12, 2020. HXDP will continue to support vSphere 6.0 U3 on both 3.5(2x) and 4.0(2x) long lived releases. However, no bug or security fixes will be provided by VMware or Cisco for ESXi going forward due to reaching the last day of support. Cisco TAC will continue to support customers to the best of their ability on ESXi 6.0 builds that have already been released. Cisco strongly recommends upgrading as soon as possible to a supported VMware vSphere 6.5 or 6.7 release and follow Cisco's recommendations as outlined in Recommended Cisco HyperFlex HX Data Platform Software Releases - for Cisco HyperFlex HX-Series Systems.

- **vSphere 6.0 Upgrades**—Users on vSphere 6.0 migrating to 6.5, upgrade components in the following order:

  1. Upgrade HX Data Platform and UCS firmware.

  2. Upgrade HX Data Platform and ESXi.

  3. Upgrade HX Data Platform only first, then upgrade ESXi or UCS firmware or both.

- **M4/M5 Mixed Domains**—A mixed domain occurs when a new, separate M5 cluster is installed under the same UCS domain that contains existing M4 cluster(s). Under these conditions, orchestrated UCS server firmware upgrade will not operate until Cisco HX Data Platform, Release 2.6 or later is installed on the M4 clusters. Therefore, it is best practice to first upgrade UCS server firmware to the latest 3.1(3) or 3.2(2) patch release before adding a new M5 cluster to the existing UCS domain. Additionally, any 1.7 HX Data Platform clusters must first be upgraded before adding any new M5 clusters to the same domain.

- **Cisco HX Data Platform 2.1(1b) with Self-Encrypting Drives (SEDs)**—Upgrading SED-ready systems running 2.1 requires UCS infrastructure and server firmware upgrades. For more information, see **Field Notice (70234)** and CSCvk17250.

- **Admin User Account** - Users may need to reset the cluster admin password if upgrading from a cluster initially deployed with Cisco HX Data Platform, Release 1.7, or if the password was manually changed after deployment. For more information, see the Cisco HyperFlex Systems Upgrade Guide.

Mixed Cluster Expansion Guidelines

- Hypercheck Health Check Utility— Cisco recommends running this proactive health check utility on your HyperFlex cluster prior to upgrade. These checks provide early visibility into any areas that may need attention and help ensure a seamless **upgrade** experience. For more information on how to install and run Hypercheck, see the Hypercheck: Hyperflex Health & Pre-Upgrade Check Tool Tech Note.

- Expanding existing M4 cluster with M5 converged nodes is supported.

- Expanding existing M5 cluster with M4 converged nodes is not supported.

- Expanding existing mixed M4/M5 cluster with M4 or M5 converged nodes is supported.

- Adding any supported compute-only nodes is permitted with all M4, M5, and mixed M4/M5 clusters using the HX Data Platform 2.6 or later Installer. Some example combinations are listed here, many other combinations are possible.

```
Example combinations:
Expand mixed M4/M5 cluster with compute-only B200, C220, C240 M4/M5
Expand M4 cluster with compute-only B200 M5, C220 M5, C240M5
```

- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.

- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers. For more information on drive compatibility, refer to the Cisco Hyperflex Drive Compatibility document.

  - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.

- HX Edge, SED, LFF, Hyper-V, and Stretched Clusters do not support mixed M4 and M5 clusters.

Mixed Cluster Expansion Guidelines for Release 3.5

A mixed cluster is defined by having both M4 and M5 HX converged nodes within the same storage cluster. When configuring a mixed cluster, the following guidelines apply:

- Expanding existing M4 cluster with M5 converged nodes is supported.

- Expanding existing M5 cluster with M4 converged nodes is not supported.

- Expanding existing mixed M4/M5 cluster with M4 or M5 converged nodes is supported.

- Adding any supported compute-only nodes is permitted with all M4, M5, and mixed M4/M5 clusters using the HX Data Platform 2.6 or later Installer. Some example combinations are listed here, many other combinations are possible.

```
Example combinations:
Expand mixed M4/M5 cluster with compute-only B200, C220, C240 M4/M5
Expand M4 cluster with compute-only B200 M5, C220 M5, C240M5
```

- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.

- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers.

  - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.

- HyperFlex Edge does not support mixed clusters.

- SED SKUs do not support mixed clusters.

Security Fixes

The following security issues are resolved:

Release	Defect ID	CVE	Description
3.5(2i)	CSCvs41636	CVE-2019-11745	It was discovered that NSS incorrectly handled certain memory operations. A remote attacker could use this issue to cause NSS to crash, resulting in a denial of service, or possibly execute arbitrary code. Note that Tenable Network Security has extracted the preceding description block directly from the Ubuntu security advisory. Tenable has attempted to automatically clean and format it as much as possible without introducing additional issues. Solution: Update the affected libnss3 package. See Also https://usn.ubuntu.com/4203-1/ Output: Installed package : libnss3_2:3.28.4-0ubuntu0.16.04.6 Fixed package : libnss3_2:3.28.4-0ubuntu0.16.04.8
3.5(2i)	CSCvs43371	CVE-2008-5161	Cisco HyperFlex System includes a version of **OpenSSH CBC Mode Ciphers** that are affected by the vulnerabilities identified.
3.5(2i)	CSCvs41636	CVE-2019-11745	Cisco HyperFlex includes a version of **libnss** that is affected by the vulnerabilities identified.

Release	Defect ID	CVE	Description
3.5(2i)	CSCvu99537		This product includes a version of **Third-party Software** that is affected by the vulnerabilities identified.

Release	Defect ID	CVE	Description
		CVE-2017-1110, CVE-2017-11109,	
		CVE-2017-5953, CVE-2017-6349,	
		CVE-2017-6350, CVE-2018-20786,	
		CVE-2018-8740, CVE-2019-12387,	
		CVE-2019-12855, CVE-2019-13734,	
		CVE-2019-13750, CVE-2019-13751,	
		CVE-2019-13752, CVE-2019-13753,	
		CVE-2019-1547, CVE-2019-1549,	
		CVE-2019-1551, CVE-2019-1563,	
		CVE-2019-17023, CVE-2019-19603,	
		CVE-2019-19645, CVE-2019-19880,	
		CVE-2019-19923, CVE-2019-19924,	
		CVE-2019-19925, CVE-2019-19926,	
		CVE-2019-19956, CVE-2019-19959,	
		CVE-2019-20079, CVE-2019-20218,	
		CVE-2019-3689, CVE-2019-8457,	
		CVE-2019-9512, CVE-2019-9514,	
		CVE-2019-9515, CVE-2020-10108,	
		CVE-2020-10109, CVE-2020-10531,	
		CVE-2020-11655, CVE-2020-12049,	

Release	Defect ID	CVE	Description
		CVE-2020-12399, CVE-2020-12762, CVE-2020-13434, CVE-2020-13435, CVE-2020-13630, CVE-2020-13631, CVE-2020-13632, CVE-2020-13790, CVE-2020-7595, CVE-2020-8169, CVE-2020-8177, CVE-2020-8597, CVE-2020-9327	
3.5(2i)	CSCvv15388	CVE-2020-14422	This product includes **Third-party Software** that is affected by the vulnerabilities identified.
3.5(2i)	CSCvr36903	CVE-2015-9383, CVE-2016-3977, CVE-2018-11490, CVE-2018-20406, CVE-2018-20852, CVE-2019-10160, CVE-2019-15133, CVE-2019-15903, CVE-2019-5010, CVE-2019-5481, CVE-2019-5482, CVE-2019-9636, CVE-2019-9740, CVE-2019-9947, CVE-2019-9948	Cisco HyperFlex includes a version of **Ubuntu** that is affected by the vulnerabilities identified.

Release	Defect ID	CVE	Description
3.5(2h)	CSCvs06094	CVE-2015-9383, CVE-2018-14498 CVE-2018-20406, CVE-2018-20852, CVE-2019-10160, CVE-2019-13117, CVE-2019-13118, CVE-2019-14287, CVE-2019-14973, CVE-2019-15903, CVE-2019-17546, CVE-2019-18197, CVE-2019-18218, CVE-2019-5010, CVE-2019-5094, CVE-2019-5481, CVE-2019-5482, CVE-2019-9636, CVE-2019-9740, CVE-2019-9947, CVE-2019-9948	This bug is to address Tenable Scan Vulnerabilities.

3.5(2h)	CSCvp71632	CVE-2018-12127, CVE-2018-12126,
		
		CVE-2018-12130, CVE-2019-11091
			HyperV - Evaluation of HyperFlex for Intel 2019.1 QSR - MDS. The following Microsoft patches were tested:
			
			`May 14, 2019—KB4494440 (OS Build 14393.2969)`
			
			Applies to: Windows 10, version 1607 Windows Server 2016
			
			`May 14, 2019—KB4494441 (OS Build 17763.503)`
			
			Applies to: Windows 10, version 1809 Windows Server 2019, all versions
			
			For more information on the hypervisor-specific mitigation for HyperFlex systems running Hyper-V, see Microsoft KB4494440 (Windows Server 2016) and KB4494441 (Windows Server 2019).

Release	Defect ID	CVE	Description
3.5(2h)	CSCvp71634	CVE-2018-12127, CVE-2018-12126, CVE-2018-12130, CVE-2019-11091	The Sequential-context attack vector is mitigated by a hypervisor update to the product versions listed in VMSA-2019-0008. This mitigation is enabled by default in ESXi and does not impose a significant performance impact according to VMware. The Concurrent-context attack vector is mitigated through enablement of the ESXi Side-Channel-Aware Scheduler Version 1 or Version 2. These options may impose a non-trivial performance impact and are not enabled by default by ESXi. For more information on the hypervisor-specific mitigation for HyperFlex systems running ESXi, see VMware KB67577.
3.5(2g)	CSCvr20154	CVE-2019-10086	Multiple Vulnerabilities in commons beanutils commons-beanutils.
3.5(2g)	CSCvj95584	CVE-2019-12620	A vulnerability in the statistics collection service of Cisco HyperFlex Software could allow an unauthenticated, remote attacker to inject arbitrary values on an affected device. The vulnerability is due to insufficient authentication for the statistics collection service. An attacker could exploit this vulnerability by sending properly formatted data values to the statistics collection service of an affected device. A successful exploit could allow the attacker to cause the web interface statistics view to present invalid data to users. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
3.5(2g)	CSCvo98516	CVE-2019-1975	A vulnerability in the web-based interface of Cisco HyperFlex Software could allow an unauthenticated, remote attacker to execute a cross-frame scripting (XFS) attack on an affected device.
			This vulnerability is due to insufficient HTML iframe protection. An attacker could exploit this vulnerability by directing a user to an attacker-controlled web page that contains a malicious HTML iframe. A successful exploit could allow the attacker to conduct clickjacking or other clientside browser attacks.
			For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
3.5(2g)	CSCvr03322		Vulnerabilities in open source software components identified by routine scans.

Release	Defect ID	CVE	Description
		CVE-2014-9092, CVE-2015-9262,	
		CVE-2016-10087, CVE-2016-10165,	
		CVE-2016-10708, CVE-2016-10713,	
		CVE-2016-3616, CVE-2016-9318,	
		CVE-2017-10053, CVE-2017-10067,	
		CVE-2017-10074, CVE-2017-10078,	
		CVE-2017-10081, CVE-2017-10087,	
		CVE-2017-10089, CVE-2017-10090,	
		CVE-2017-10096, CVE-2017-10101,	
		CVE-2017-10102, CVE-2017-10107,	
		CVE-2017-10108, CVE-2017-10109,	
		CVE-2017-10110, CVE-2017-10111,	
		CVE-2017-10115, CVE-2017-10116,	
		CVE-2017-10118, CVE-2017-10135,	
		CVE-2017-10176, CVE-2017-10193,	
		CVE-2017-10198, CVE-2017-10243,	
		CVE-2017-10274, CVE-2017-10281,	
		CVE-2017-10285, CVE-2017-10295,	
		CVE-2017-10345, CVE-2017-10346,	
		CVE-2017-10347, CVE-2017-10348,	

Release	Defect ID	CVE	Description
		CVE-2017-10349, CVE-2017-10350,	
		CVE-2017-10355, CVE-2017-10388,	
		CVE-2017-15232, CVE-2017-16932,	
		CVE-2017-17512, CVE-2017-18258,	
		CVE-2017-3509, CVE-2017-3511,	
		CVE-2017-3526, CVE-2017-3533,	
		CVE-2017-3544, CVE-2018-0734,	
		CVE-2018-0735, CVE-2018-0737,	
		CVE-2018-1000030, CVE-2018-1000156,	
		CVE-2018-1000802, CVE-2018-1000807,	
		CVE-2018-1000808, CVE-2018-1060,	
		CVE-2018-1061, CVE-2018-10916,	
		CVE-2018-10963, CVE-2018-11212,	
		CVE-2018-11214, CVE-2018-1152,	
		CVE-2018-11574, CVE-2018-12384,	
		CVE-2018-12404, CVE-2018-13785,	
		CVE-2018-14404, CVE-2018-14567,	
		CVE-2018-14647, CVE-2018-15473,	
		CVE-2018-16428, CVE-2018-16429,	
		CVE-2018-16435, CVE-2018-16890,	

Release	Defect ID	CVE	Description
		CVE-2018-17100, CVE-2018-17101,	
		CVE-2018-18311, CVE-2018-18312,	
		CVE-2018-18313, CVE-2018-18314,	
		CVE-2018-18557, CVE-2018-18585,	
		CVE-2018-18661, CVE-2018-2579,	
		CVE-2018-2588, CVE-2018-2599,	
		CVE-2018-2602, CVE-2018-2603,	
		CVE-2018-2618, CVE-2018-2633,	
		CVE-2018-2634, CVE-2018-2637,	
		CVE-2018-2641, CVE-2018-2663,	
		CVE-2018-2677, CVE-2018-2678,	
		CVE-2018-2783, CVE-2018-2794,	
		CVE-2018-2795, CVE-2018-2796,	
		CVE-2018-2797, CVE-2018-2798,	
		CVE-2018-2799, CVE-2018-2800,	
		CVE-2018-2814, CVE-2018-2815,	
		CVE-2018-2952, CVE-2018-3149,	
		CVE-2018-3150, CVE-2018-3169,	
		CVE-2018-3180, CVE-2018-3183,	
		CVE-2018-3214, CVE-2018-6594,	

Release	Defect ID	CVE	Description
		CVE-2018-6951, CVE-2018-7456, CVE-2018-8905, CVE-2019-2422, CVE-2019-3462, CVE-2019-3822, CVE-2019-3823, CVE-2019-6109, CVE-2019-6111	
3.5(2g)	CSCvq24176	CVE-2018-15380	A vulnerability in the cluster service manager of Cisco HyperFlex could allow an unauthenticated, adjacent attacker to execute commands as the root user. The vulnerability is due to insufficient input validation. An attacker could exploit this vulnerability by connecting to the cluster service manager and injecting commands into the bound process. A successful exploit could allow the attacker to run commands on the affected host as the root user. Cisco has released software updates that address this vulnerability. There are workarounds that address this vulnerability. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
3.5(2g)	CSCvo88997		The vulnerabilities identified with JVM 1.8U121 results in a memory leak getting VC alarms (concurrent 40 calls using REST API).

Release	Defect ID	CVE	Description
		CVE-2017-10053, CVE-2017-10067,	
		CVE-2017-10074, CVE-2017-10078,	
		CVE-2017-10081, CVE-2017-10087,	
		CVE-2017-10089, CVE-2017-10090,	
		CVE-2017-10096, CVE-2017-10101,	
		CVE-2017-10102, CVE-2017-10107,	
		CVE-2017-10108, CVE-2017-10109,	
		CVE-2017-10110, CVE-2017-10111,	
		CVE-2017-10115, CVE-2017-10116,	
		CVE-2017-10118, CVE-2017-10135,	
		CVE-2017-10176, CVE-2017-10193,	
		CVE-2017-10198, CVE-2017-10243,	
		CVE-2017-10274, CVE-2017-10281,	
		CVE-2017-10285, CVE-2017-10295,	
		CVE-2017-10345, CVE-2017-10346,	
		CVE-2017-10347, CVE-2017-10348,	
		CVE-2017-10349, CVE-2017-10350,	
		CVE-2017-10355, CVE-2017-10356,	
		CVE-2017-10357, CVE-2017-10388,	
		CVE-2017-3509, CVE-2017-3511,	

Release	Defect ID	CVE	Description
		CVE-2017-3526, CVE-2017-3533,	
		CVE-2017-3539, CVE-2017-3544,	
		CVE-2018-2579, CVE-2018-2582,	
		CVE-2018-2588, CVE-2018-2599,	
		CVE-2018-2602, CVE-2018-2603,	
		CVE-2018-2618, CVE-2018-2629,	
		CVE-2018-2633, CVE-2018-2634,	
		CVE-2018-2637, CVE-2018-2641,	
		CVE-2018-2663, CVE-2018-2677,	
		CVE-2018-2678, CVE-2018-2783,	
		CVE-2018-2790, CVE-2018-2794,	
		CVE-2018-2795, CVE-2018-2796,	
		CVE-2018-2797, CVE-2018-2798,	
		CVE-2018-2799, CVE-2018-2800,	
		CVE-2018-2814, CVE-2018-2815,	
		CVE-2018-2952, CVE-2018-3136,	
		CVE-2018-3139, CVE-2018-3149,	
		CVE-2018-3150, CVE-2018-3169,	
		CVE-2018-3180, CVE-2018-3183,	
		CVE-2018-3214, CVE-2019-2422	

Release	Defect ID	CVE	Description
3.5(2a)	CSCvj95606	CVE-2018-15380	A vulnerability in the cluster service manager of Cisco HyperFlex could allow an unauthenticated, adjacent attacker to perform a command injection as the root user. The vulnerability is due to an unprotected listening interface. An attacker could exploit this vulnerability by connecting to the listening interface and injecting commands to the bound process. An exploit could allow the attacker to run commands on the affected host as the root user. For more information, see the related Cisco Security Advisory.
3.5(2a)	CSCvn22303	CVE-2016-1000031	The vulnerabilities associated with the third-party software included in **Apache Struts Commons FileUpload RCE**.
3.5(2a)	CSCvm93059	CVE-2018-18074	The vulnerabilities associated with the software version included in the **Python package**, in the **Ubuntu** kernel.
3.5(2a)	CSCvm53142	CVE-2018-14598 CVE-2018-14599 CVE-2018-14600	The vulnerabilities associated with **libx11**.
3.5(2a)	CSCvm53132	CVE-2018-14622	The vulnerabilities associated with **libtirpc**.
3.5(2a)	CSCvm34693	CVE-2018-1060	The vulnerabilities associated with the software version included in **Python pop3lib apop() Method Denial of Service**.
3.5(2a)	CSCvm02920	CVE-2018-3615 CVE-2018-3620 CVE-2018-3646	The vulnerabilities associated with **August CPU Side-Channel Information Disclosure**.
3.5(2a)	CSCvk31047	CVE-2019-1664	The vulnerabilities associated with the **hxterm** service included in Cisco HX Data Platform.
3.5(2a)	CSCvj08921	CVE-2018-7750	The vulnerabilities associated with software version included in **Paramiko transport.py Authentication Bypass**.
3.5(2a)	CSCvj95590	CVE-2019-1667	The vulnerabilities associated with the **Graphite** interface included in Cisco HX Data Platform.

Release	Defect ID	CVE	Description
3.5(2a)	CSCvj95580	CVE-2019-1666	The vulnerabilities associated with the **Graphite** service included in Cisco HX Data Platform.
3.5(1a)	CSCvk59165	CVE-2019-1665	The vulnerabilities associated with the web-based management interface of Cisco HyperFlex software that may allow an unauthenticated, remote attacker to conduct a cross-site scripting (XSS) attack.
3.5(1a)	CSCvm29418	CVE-2017-18342	The vulnerabilities associated with the third-party software included in **PyYAML**.
3.5(1a)	CSCvm20484	CVE-2018-14682 CVE-2018-14679 CVE-2018-14680 CVE-2018-14681	The vulnerabilities associated with the **libmspack** software package version included in Cisco HX Data Platform.
3.5(1a)	CSCvm15800	CVE-2018-5391	Cisco HyperFlex System includes a version of the Linux kernel that is affected by the IP Fragment Reassembly Denial of Service Vulnerability
3.5(1a)	CSCvm10159	CVE-2015-9262	The vulnerabilities associated with **libXcursor** included in Cisco HX Data Platform.
3.5(1a)	CSCvk59406	CVE-2018-15407	A vulnerability in the installation process of Cisco HyperFlex Software could allow an authenticated, local attacker to read sensitive information.
3.5(1a)	CSCvk32464	No CVE ID has been assigned to this issue.	A vulnerability in file permissions of Cisco HyperFlex Software could allow an authenticated, local attacker to read sensitive files.
3.5(1a)	CSCvk22858	CVE-2018-15382	A vulnerability in Cisco HyperFlex Software could allow an unauthenticated, remote attacker to generate signed session tokens.
3.5(1a)	CSCvk09234	CVE-2018-1092, CVE-2018-7492 CVE-2018-8087, CVE-2018-1068 CVE-2018-8781	The vulnerabilities associated with the Ubuntu Linux kernel version included in Cisco HX Data Platform.
3.5(1a)	CSCvk05700	CVE-2018-12015	The vulnerabilities associated with the **perl** software package version included in Cisco HX Data Platform.

Release	Defect ID	CVE	Description
3.5(1a)	CSCvk05679	CVE-2014-9620 CVE-2014-9653 CVE-2015-8865 CVE-2018-10360 CVE-2014-9621	The vulnerabilities associated with the **file** software package version included in Cisco HX Data Platform.
3.5(1a)	CSCvk00405	CVE-2016-6796 CVE-2017-12615 CVE-2017-7674 CVE-2016-0762 CVE-2016-6797 CVE-2017-12616 CVE-2018-1304 CVE-2016-5018 CVE-2016-6816 CVE-2017-12617 CVE-2017-5647 CVE-2016-6794 CVE-2016-8735	The vulnerabilities associated with the **Apache Tomcat** software package version included in Cisco HX Data Platform.
3.5(1a)	CSCvj95644	CVE-2018-15423	A vulnerability in the web interface of the Cisco HyperFlex Software could allow an unauthenticated remote attacker to affect the integrity of the device via a clickjacking attack.

Release	Defect ID	CVE	Description
3.5(1a)	CSCvj81584		The vulnerabilities associated with the **Apache Tomcat 7.x** version included in Cisco HX Data Platform.

Release	Defect ID	CVE	Description
		CVE-2010-4172, CVE-2011-1088	
		CVE-2011-1582, CVE-2009-0783	
		CVE-2011-5062, CVE-2013-4590	
		CVE-2016-0762, CVE-2017-5648	
		CVE-2012-2733, CVE-2014-0099	
		CVE-2016-5018, CVE-2018-1304	
		CVE-2014-0230, CVE-2016-6816	
		CVE-2011-2729, CVE-2013-2071	
		CVE-2015-5346, CVE-2017-12616	
		CVE-2009-3555, CVE-2010-4476	
		CVE-2011-1183, CVE-2011-2204	
		CVE-2011-3190, CVE-2013-4286	
		CVE-2015-5351, CVE-2017-12617	
		CVE-2011-5063, CVE-2014-0050	
		CVE-2016-0763, CVE-2017-5664	
		CVE-2012-3439, CVE-2014-0119	
		CVE-2016-6794, CVE-2018-1305	
		CVE-2012-4534, CVE-2014-7810	
		CVE-2016-8735, CVE-2011-2481	
		CVE-2010-2227, CVE-2011-0013	
		CVE-2011-1184, CVE-2012-5568	
		CVE-2015-5174, CVE-2016-8745	
		CVE-2011-3375, CVE-2013-4322	
		CVE-2016-0706, CVE-2017-15706	
		CVE-2011-5064, CVE-2014-0075	
		CVE-2016-3092, CVE-2017-6056	
		CVE-2012-3544, CVE-2014-0160	
		CVE-2016-6796, CVE-2018-8014	
		CVE-2011-1475, CVE-2005-2090	
		CVE-2012-4431, CVE-2010-3718	
		CVE-2011-0534, CVE-2012-3546	

Release	Defect ID	CVE	Description
		CVE-2014-0227, CVE-2016-6797	
		CVE-2011-2526, CVE-2013-2067	
		CVE-2015-5345, CVE-2017-12615	
		CVE-2011-3376, CVE-2013-4444	
		CVE-2016-0714, CVE-2017-5647	
		CVE-2012-0022, CVE-2014-0096	
		CVE-2016-3427, CVE-2017-7674	
3.5(1a)	CSCvj99081	CVE-2017-16612 CVE-2018-8012 CVE-2018-10237	The vulnerabilities associated with the **Apache zookeeper** version included in Cisco HX Data Platform.
3.5(1a)	CSCvj95632	CVE-2018-15382	A vulnerability in Cisco HyperFlex could allow an unauthenticated, remote attacker to generate signed session tokens.
3.5(1a)	CSCvj08923	CVE-2018-6594	The vulnerabilities associated with the **El Gamal implementation in PyCrypto** included in Cisco HX Data Platform.
3.5(1a)	CSCvj08160	CVE-2017-5715 CVE-2017-5753 CVE-2017-5754	The vulnerabilities associated with the HyperFlex Controller VM software for Windows Server with Microsoft Hyper-V.
3.5(1a)	CSCvj63266	CVE-2018-1000300 CVE-2018-1000301 CVE-2018-1000303	The vulnerabilities associated with **CURL** software package included in Cisco HX Data Platform.
3.5(1a)	CSCvj61269	CVE-2018-0494	The vulnerabilities associated with **GNU Wget** version included in Cisco HX Data Platform.
3.5(1a)	CSCvj55521	CVE-2018-8897, CVE-2018-1087 CVE-2018-1000199	The vulnerabilities associated with Linux kernel version included in Cisco HX Data Platform.
3.5(1a)	CSCvj59134	CVE-2015-9262 CVE-2018-3639 CVE-2017-3640	The vulnerabilities associated with **May CPU Side-Channel** affecting Cisco HX Data Platform.
3.5(1a)	CSCvj42966	—	The vulnerability associated with the Data Protection clone api return value.

Release	Defect ID	CVE	Description
3.5(1a)	CSCvi88567	CVE-2017-16995 CVE-2017-0861 CVE-2017-1000407 CVE-2017-11472 CVE-2017-15129 CVE-2017-16528	The vulnerabilities associated with **Linux kernel** version included in Cisco HX Data Platform.
3.5(1a)	CSCvi60720	—	This is a modification on the product to adopt new secure code best practices to enhance the security posture and resiliency of the Cisco HyperFlex HX Data Platform.
3.5(1a)	CSCvi48372	CVE-2018-15429	A vulnerability in the web-based UI of Cisco HyperFlex HX Data Platform software could allow an unauthenticated, remote attacker to access sensitive information on an affected system.
3.5(1a)	CSCvi46951	CVE-2017-7529	The vulnerabilities associated with the **nginx** software package version included in Cisco HX Data Platform.
3.5(1a)	CSCvi47250	CVE-2011-3389	The vulnerabilities associated with the **OpenSSL Protocol** software package version included in Cisco HX Data Platform.
3.5(1a)	CSCvi26246	CVE-2016-3092 CVE-2013-0248 CVE-2014-0050	The vulnerabilities associated with the **Apache commons-file upload** version included in Cisco UCS.
3.5(1a)	CSCvi50910	CVE-2016-2183	The vulnerabilities associated with the DES and Triple DES ciphers software package version included in Cisco HX Data Platform.

Resolved Caveats in Release 3.5(2i)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvu69826	stcli cluster reregister command fails with the error: ```Storage cluster reregistration with a new vCenter failed. java.rmi.RemoteException: VI SDK invoke exception:; nested exception is: java.rmi.RemoteException: Exception in WSClient.invoke :; nested exception is: java.lang.NullPointerException```	3.5(2h)	3.5(2i)
CSCvr89066	Old files in /var/support/ZKTxnlog are not purged with the daily zklog-cleanup cron job.	3.5(2c)	3.5(2i)
CSCvr89510	A HyperV cluster with RF2 has 2 simultaneous drives failures on 2 different nodes. This makes the cluster critical. Even after fixing the 2 failed drives, some VMs fail to power on.	4.0(2a)	3.5(2i)
CSCvp97422	When network partition heals, the cluster becomes healthy but the datastores on one of the nodes remains unavailable for some time.	3.5(2g)	3.5(2i)
CSCvr16760	The HyperFlex cluster healing state is stuck at 87% and is unable to progress.	3.5(2f)	3.5(2i)
CSCvs22477	Sometimes it is observed that storfs was killed and restarted due to Out of Memory (OOM) situation.	3.5(2e)	3.5(2i)
CSCvs54285	A cluster node may hang in the Linux kernel. This is classified as an oops and a deviation from the expected behavior.	3.5(2h)	3.5(2i)
CSCvo21125	During ESXi upgrade, it waits forever and then times out with error "Waiting for vCenter to connect to cluster node". On the vCenter we see that the host has already joined back.	3.5(1a), 3.5(2h)	3.5(2i)
CSCvq45087	During deployment phase, HX-Installer successfully deploys SCVM and creates 'admin' username with preprovisioned password. In case, when password contains special symbol(s) like ":", subsequent attempts to login into the same VM using admin credentials fails.	3.5(2a)	3.5(2i)
CSCvs55422	Reduction of System log folder to alarming rate due to tomcat logs.	2.5(1b) 3.5(1a) 4.0(2a	3.5(2i)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvt16158	Starting with new 6.5 and 6.7 ESXi releases, HX clusters with certain ESXI build numbers won't be able to do a direct ESXi upgrade to 6.5 and 6.7 versions. If the build number is older, a two step upgrade is required to get to the latest 6.5 and 6.7 ESXi builds.	4.0(2a)	3.5(2i)
CSCvr97089	High percentage of packet loss on the HyperFlex Storage Controller.	3.5(2g)	3.5(2i)
CSCvu58631	HX 3.5(2h) SED stretch cluster expansion fails as converged node is not listed in server selection page.	3.5(2h)	3.5(2i)

Resolved Caveats in Release 3.5(2h)

Defect ID	Symptom	First Release Affected	Resolved in Release
Management			
CSCvs47419	Regenerating new SSL cert and performing re-register with vCenter results in fingerprint mismatch.	3.5(2g)	3.5(2h)
CSCvr15546	On older HX software, the self-signed certificate shows expiry as Dec 2019 with SHA1 as the certificate being used. This does not have any functional impact to the cluster.	3.5(2g)	3.5(2h)
CSCvr92004	Unable to login to HXconnect. User sees authentication failure error message during login.	3.5(2g)	3.5(2h)
CSCvs18117	HX Connect login reports "access token invalid" when fetching replication info.	3.5(2g)	3.5(2h)
CSCvq65056	After expanding an upgraded cluster from a 2.6 cluster to HX 4.0(1b), vMotion fails when tried to and from the expanded node.	3.5(2g)	3.5(2h)
CSCvr37488	Pre-upgrade fails on the last node with **Downgrade error**.	3.5(2g)	3.5(2h)
CSCvr75305	Node add hangs intermittently forever.	3.5(2e)	3.5(2h)
CSCvr96151	Error in server selection page for ESXI on branch 3.5.2i-32123.	3.5(2i)	3.5(2h)
CSCvs02458	Intermittent Upgrade Failure when upgrading to HX 3.5.(2g).	3.5(2g)	3.5(2h)

Defect ID	Symptom	First Release Affected	Resolved in Release

CSCvs04926	On an HX 3.5(2g) cluster, when upgrading UCS Server firmware via HX Connect, it shows one of the following errors:
	
	`getUcsAvailablePackagesLocalizable`
	`Message(Operation did not complete in expected`
	`time and maybe executing in the background.,None`
	`,None,Operation did not complete in expected time`
	
	`and maybe executing in the`
	`background.,ArrayBuffer())`
	
	The following error is more probable with >4 node clusters.
	
	`getUcsHfpVersionsLocalizableMessage`
	`(Operation did not complete in expected time and`
	
	`maybe executing in the`
	`background.,None,None,Operation`
	`did not complete in expected time and maybe`
	`executing in the background.,ArrayBuffer()`
		3.5(2g)	3.5(2h)

CSCvr44548	Expansion workflow is not setting hfp for the M3 compute node.	3.5(2g)	3.5(2h)
CSCvr67532	The HX240C-M5L specification sheet states that PID HX-HD12T7KL4KN is supported starting with HXDP 4.0(1a). No validation error is triggered. The deployment fails at the "Cluster Creation" step with the error: "Disk prepare failed magnetic disk /dev/sdX", and the disks are blacklisted.	3.5(2e)	3.5(2h)

Resolved Caveats in Release 3.5(2g)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvk25616	Datastores not mounted in Vcenter/ESXi after host reboot.	3.0(1d)	3.5(2g)
CSCvk37044	Set computer account fails due to the use of "." in the username.	3.0(1d)	3.5(2g)
CSCvm77294	While upgrading a cluster, the following error appears: "Failed upgrade validations: Checking vCenter configuration. Reason: Upgrade validations failed. DRS Fault: Insufficient resources to satisfy configured failover".	3.5(1a)	3.5(2g)
CSCvn21067	Support bundle does not connect files from /var/support/ZKTxnlog	3.0(1i)3.5(2a)	3.5(2g)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvn23123	Storfs node PANIC with segmentation fault at replication stat update	3.0(1b)	3.5(2g)
CSCvn23812	HyperFlex's controller VM can be configured to have a CA signed certificate for HTTPS communication instead of a default shipped certificate. However, during an upgrade, the CA signed certificate gets overridden by the default certificate prompting the user to configure the CA signed certificate again.	3.5(2a)	3.5(2g)
CSCvn51578	After performing procedure to remove node, stcli cluster reregister for vCenter fails.	3.5(2a) 3.5(2h)	3.5(2g)
CSCvn63359	After modifying the timezone using stcli command, the date always comes in UTC (default).	3.5(2a)	3.5(2g)
CSCvr92004	Unable to log in to HX-Connect. User sees authentication failure error message during log in.	3.5(2g)	3.5(2g)
CSCvp17427 CSCvm47257	Stcli cluster storage-summary takes a long time to return on 16+ node cluster when one node reboots.	3.5(1a)	3.5(2g)
CSCvp58739	HX Install/Upgrade validation fails due to SAS controller being on firmware 09.00.00.06.	3.5(2c)	3.5(2g)
CSCvp63958	HX replication cleanup failing with error: "INFO:DR state is NOT clean".	3.5(2a)	3.5(2g)
CSCvp64140	During cluster creation, the following error occurs on installer: "Failure occurred during Cluster Creation process: Unable to post the content to the down-stream".	3.5(2f)	3.5(2g)
CSCvp64572	ICMP redirection enabled.	3.5(1a)	3.5(2g)
CSCvp65666	In HX Connect under Activity, Job Type: encryptionLocal* will have status: Success but will show the job as if it was still running. Spinning in progress icon is visible with RunStep: next to it.	3.5(2a)	3.5(2g)
CSCvp89523	When a HyperFlex Cluster is deployed on an ACI Fabric, with HX MGMT and HX Storage in the same Bridge Domain, we see ARP responses from wrong HX Storage Controller interface, causing incorrect ARP learning to occur. On the ACI Fabric, we see MAC moves for HX IPs.	3.5(2b)	3.5(2g)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvp95655	When installing Hyperflex from Intersight for 3Node and FI Managed UCS/HX Cluster. One of the steps which shows running is "Witness Node IP Reachability Check". Ideally this shouldn't show unless if it is Stretch Cluster. It appears all the tasks for HX installation are run even though some might be skipped.	3.5(2g)	3.5(2g)
CSCvq01506	During the online upgrade process from releases prior to 2.5 to a post 2.5 release, there is a change in the metadata format. Metadata for files that were truncated during the upgrade window, are not set correctly and this may cause the cluster to go down.	2.1(1b)	3.5(2g)
CSCvq04252	HX Release 3.5(2b) installer fails on the hypervisor configuration step with no visible error. UCSM configuration completes, but Hypervisor Configuration seems to not start.	3.5(2a)	3.5(2g)
CSCvq06952	Snapshot creation on CBT enabled VM fails with error "Failed in vmreparent vmkfstools clone1".	3.5(2c)	3.5(2g)
CSCvq15987	Cluster expansion fails with error: "SP template org-root/org-hx-cluster/ls-compute-nodes-m5 does not exist".	3.5(2b)	3.5(2g)
CSCvq33538	"Internal Error" while running commands 'asupcli --all post --type alert' and 'asupcli --all ping'	3.0(1b)	3.5(2g)
CSCvq34873	Memory usage by carbon cache.	3.5(2b)	3.5(2g)
CSCvq60925	If HX Cluster is upgraded from prior to Release 3.5(1a), and cluster expansion is done in that sequence, then the existing UCS Servers will go into "Pending Reboot", which is seen in UCS Manager.	3.5(2d)	3.5(2g)
CSCvq64208	When Witness Node is powered off, HX Connect still shows witness node online.	3.5(2f)	3.5(2g)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvq70832	User has created macpool with b1 in 5th octet for external storage like iscsi and assigned vmnic9 to b1 which is not a suggested macpool configuration as per the documentation and that caused the logic to pick A1 and B1 macs for mgmt which resulted in unsupported configuration while going through the kernel migration step as part of upgrade. This caused the HX upgrade to break. As per the documentation link, for external storage user is suggested to create a macpool with 00:25:B5:XX:01:01-63 which would have avoided the problem.	3.5(2d)	3.5(2g)
CSCvq77016	LAZ was disabled previously but got enabled again after recovery from outage.	3.5(2d), 4.0(1b)	3.5(2g)
CSCvq86205	During cluster expansion, validation may fail with the following error: "hw_data_disk_same_size. hw_data_disk_same_size_fail". This is caused by UCS discovering disk with slight variance in size, 1 to 500 Mb for example. This variance exists between disks from different vendors or even same vendors.	3.5(2e)	3.5(2g)
CSCvq84609	On a large scale cluster, HX Connect failed with error message "*server call failure.*".	3.5(2g)	3.5(2g)
CSCvq90989	The failure to upgrade or operate in HX maintenance mode due to this defect is extremely rare. This defect is exposed if the witness VM is inadvertently shared across different stretched clusters or incorrectly configured (for example, IP reuse/conflict).	3.5(2f)	3.5(2g)
CSCvq94105	Stretched Cluster status is lost when the cluster is shutdown. The cluster state shows normal until the cluster is back online.	3.5(2f)	3.5(2g)
CSCvq96085	After cluster expansion, some of the datastores failed to mount, and had to manually be mounted using HX Connect..	3.5(2d)	3.5(2g)
CSCvq99358	Unable to expand Stretched Cluster when nodes have SED drives.	3.5(2f)	3.5(2g)
CSCvr06432	On Stretched Cluster, after Cluster Management IP goes down, takes 10+ minutes for new Cluster Management IP to be accessible	3.5(2f)	3.5(2g)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvr08735	Out of memory on 2 nodes during maintenance on a 3rd node.	3.5(2a) 3.5(2b) 3.5(2d)	3.5(2g)
CSCvr11632	Unable to access Stretched Clusters via HX Connect when the vCenter server is offline.	3.5(2e)	3.5(2g)
CSCvr22391	lsass error joining controller to domain.	3.5(2e)	3.5(2g)
CSCvr28095	Tomcat / HTTP 500 errors when viewing jobs ([Activity] within HX Connect).	3.5(2b)	3.5(2g)
CSCvr44129	Storfs not started during upgrade from pre-3.5(2d) to post-3.5(2d) in one or more nodes	3.5(2f)	3.5(2g)

Resolved Caveats in Release 3.5(2f)

Defect ID	Symptom	First Release Affected	Resolved in Release
HX Connect			
CSCvq89852	HX-Connect reports Stretched Cluster as standard cluster after transition from offline to online state.	3.5(2a) 3.5(2e) 3.5(2f)	3.5(2f)
Stretched Cluster			
CSCvq53058	When the Witness VM has a high RTT times (>50ms) from any of the two Stretched Cluster sites, there is a possibility under heavy transaction load that the cluster failover or failback times may be very high.	3.5(1a)	3.5(2f)
CSCvq58829	During site failover, cluster hits an APD due to slow response from cluster metadata replication.	3.5(1a)	3.5(2f)
CSCvq17778	After back-to-back failovers on one site, the cluster encountered a failure that lead to zookeeper going offline.	3.5(2d)	3.5(2f)
CSCvq89852	HX Connect reports Stretched Cluster as standard cluster after transition from offline to online state.	3.5(2a)	3.5(2f)
Management			

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs28167		3.5(1a)	3.5(2f)

Defect ID	Symptom	First Release Affected	Resolved in Release
	In order to install or complete a node replacement on Cisco HyperFlex, customers need to download an HX Installer OVA (Open Virtual Appliance) file; include to deploy a stretched cluster, customers additionally need to download a Witness OVA. All of the code posted on CCO prior to the posting of release HX 3.5(2g) was discovered to have expired certificates as of 11/26/19. Cisco has re-signed and re-posted OVA files associated with HX releases 3.5(2e), 3.5.2(f), 3.5.2(g), 4.0(1a) and 4.0(1b) with updated certificates. For other releases, attempts to deploy an OVF template with an expired OVA will fail with the following error message: "The OVF package is signed with an invalid certificate". **Conditions:** If customers are deploying HX 3.5(2e), 3.5.2(f), 3.5.2(g), 4.0(1a) or 4.0(1b), Cisco has re-signed and re-posted OVA files and customers will not experience the problem if they use the patched OVA files. Look for a "p1" suffix in the OVA filenames, which indicates that OVA file has been fixed: File Name Examples: HX 3.5(2f) patched OVA file for Cisco HyperFlex Data Platform Installer for VMware ESXi: `Cisco-HX-Data-Platform-Installer-v3.5.2f-31787p1-esx.ova` Cisco HyperFlex Data Platform Stretched Cluster Witness: `HyperFlex-Witness-1.0.6p1.ova` Customers using the OVA files for other HX releases, refer to the following workaround. **Workaround** There are two options to move forward after failing to deploy with an OVA file that is affected (applies to the installer and witness OVA files). Option A - Remove the local manifest file. The manifest file can be deleted so vCenter does not check the validity of the certificate. 1. Download and extract the OVA file to a local directory. 2. Remove the .mf file 3. Add the remaining files to a new archive and change the file extension from '.tar' to '.ova'		

Defect ID	Symptom	First Release Affected	Resolved in Release
	4. Proceed to deploy that newly created OVA file using "Deploy by OVF Template" in vCenter. vCenter will show the file as not having a certificate. This is expected and the deployment should continue without issue. Option B - Remove the local manifest file. Manually deploy with ovftool – Use VMware's ovftool to deploy the OVA while bypassing the certificate check. The ovftool can be downloaded and run on customer's computer. The ovftool also comes pre-installed on HX Controller VMs. This is helpful for node replacements and cluster expansions.		

Defect ID	Symptom	First Release Affected	Resolved in Release

	1. Use ovftool to deploy the OVA file to a datastore while raising the --skipManifestcheck switch. For example,
	
	`root@SpringpathControllerABCDEFGH:~# ovftool --skipManifestCheck -ds=datastore http://<path to`
	
	`ova>/Cisco-HX-Data-Platform-Installer-v3.5.2c-31725-esx.ova vi://root@<IP of management ESX host>/`
	
	2. The OVA should be deployed and present in vCenter on the ESXi host previously specified.
	
	3. Power on the VM and console into it
	
	4. Login to the VM with the default username/password combination of root / Cisco123
	
	5. Set the IP of the VM statically by issuing: **vi /etc/network/eth0.interface**
	
	6. Change 'iface eth0 inet dhcp' to 'iface eth0 inet static'. Each of the following needs to be on their own line and tab indented
	
	`address <desired ip address of installer>`
	
	`netmask X.X.X.X`
	
	`gateway X.X.X.X`
	
	`<esc> :wq`
	
	7. After the file is reviewed and saved, restart the VM. The VM should now boot with the desired IP address
	
	8. The first login via the WebGUI (still using default username/password combination) will have the user change the password.
	
	9. After the password change the user can begin the desired install/expand/node replacement activity.		

CSCvo39912	UCSM Only upgrade stalled as part of HX upgrade.	3.5(2d)	3.5(2f)
CSCvq91142	Hyper-V: Cluster shows offline/APD. Seg faults during NS master ownership transfer.	3.5(2a)	3.5(2f)
CSCvn67512	SED Data SSDs may power down after firmware upgrade.	3.0(1d)	3.5(2f)
CSCvq17778	After back-to-back failovers on one site, the cluster encountered a failure that lead to zookeeper going down.	3.5(2d)	3.5(2f)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvq89852	HX Connect reports stretch cluster as standard cluster after transition from offline to online state.	3.5(2a)	3.5(2f)
CSCvq53058, CSCvq58829	For Witness VMs with high RTT times (>50ms) to any of the stretch cluster sites, there is a possibility under heavy transaction load for failover or failback times to be impacted.	3.0(1i) 3.5(1a)	3.5(2f)

Resolved Caveats in Release 3.5(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
Stretched Cluster			
CSCvq18919	Sometimes during Stretched Cluster failover, the cluster resource manager component goes down due to file write error.	3.5(2d)	3.5(2e)
Management			

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs28167		3.5(1a)	3.5(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
	In order to install or complete a node replacement on Cisco HyperFlex, customers need to download an HX Installer OVA (Open Virtual Appliance) file; include to deploy a stretched cluster, customers additionally need to download a Witness OVA. All of the code posted on CCO prior to the posting of release HX 3.5(2g) was discovered to have expired certificates as of 11/26/19. Cisco has re-signed and re-posted OVA files associated with HX releases 3.5(2e), 3.5.2(f), 3.5.2(g), 4.0(1a) and 4.0(1b) with updated certificates. For other releases, attempts to deploy an OVF template with an expired OVA will fail with the following error message: "The OVF package is signed with an invalid certificate". **Conditions:** If customers are deploying HX 3.5(2e), 3.5.2(f), 3.5.2(g), 4.0(1a) or 4.0(1b), Cisco has re-signed and re-posted OVA files and customers will not experience the problem if they use the patched OVA files. Look for a "p1" suffix in the OVA filenames, which indicates that OVA file has been fixed: File Name Examples: HX 3.5(2e) patched OVA file for Cisco HyperFlex Data Platform Installer for VMware ESXi: `Cisco-HX-Data-Platform-Installer-v3.5.2e-31762p1-esx.ova` Cisco HyperFlex Data Platform Stretched Cluster Witness: `HyperFlex-Witness-1.0.4p1.ova` Customers using the OVA files for other HX releases, refer to the following workaround. **Workaround** There are two options to move forward after failing to deploy with an OVA file that is affected (applies to the installer and witness OVA files). Option A - Remove the local manifest file. The manifest file can be deleted so vCenter does not check the validity of the certificate. 1. Download and extract the OVA file to a local directory. 2. Remove the .mf file 3. Add the remaining files to a new archive and change the file extension from '.tar' to '.ova'		

Defect ID	Symptom	First Release Affected	Resolved in Release
	4. Proceed to deploy that newly created OVA file using "Deploy by OVF Template" in vCenter. vCenter will show the file as not having a certificate. This is expected and the deployment should continue without issue. Option B - Remove the local manifest file. Manually deploy with ovftool – Use VMware's ovftool to deploy the OVA while bypassing the certificate check. The ovftool can be downloaded and run on customer's computer. The ovftool also comes pre-installed on HX Controller VMs. This is helpful for node replacements and cluster expansions.		

Defect ID	Symptom	First Release Affected	Resolved in Release

	1. Use ovftool to deploy the OVA file to a datastore while raising the --skipManifestcheck switch. For example,
	
	`root@SpringpathControllerABCDEFGH`
	`:~# ovftool --skipManifestCheck -ds=datastore`
	
	`http://<path to`
	
	`ova>/Cisco-HX-Data-Platform-Installer-`
	`v3.5.2c-31725-esx.ova vi://root@<IP of`
	`management ESX`
	`        host>/`
	
	2. The OVA should be deployed and present in vCenter on the ESXi host previously specified.
	
	3. Power on the VM and console into it
	
	4. Login to the VM with the default username/password combination of root / Cisco123
	
	5. Set the IP of the VM statically by issuing: **vi /etc/network/eth0.interface**
	
	6. Change 'iface eth0 inet dhcp' to 'iface eth0 inet static'. Each of the following needs to be on their own line and tab indented
	
	`address <desired ip address of installer>`
	
	`netmask X.X.X.X`
	
	`gateway X.X.X.X`
	
	`<esc> :wq`
	
	7. After the file is reviewed and saved, restart the VM. The VM should now boot with the desired IP address
	
	8. The first login via the WebGUI (still using default username/password combination) will have the user change the password.
	
	9. After the password change the user can begin the desired install/expand/node replacement activity.		

CSCvq18771 CSCvq02860	HX-SD16T123X-EP caching drive is assigned persistence role while the HX-SD960G61X-EV gets recognized as a cache drive.	3.5(2c) 4.0(1a)	3.5(2e)
CSCvc38351	Cluster Expansion failed because incorrect MTU size was selected in Installer.	1.8(1c) 3.5(1a)	3.5(2e)
CSCvq18919	During stretch cluster failover, zookeeper goes down due to epoch file write error.	3.5(2d)	3.5(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvp29431	HX clusters running Release 3.0.(1c) hit an All Paths Down (APD) condition, and VMs became inaccessible	3.0(1c)	3.5(2e)
CSCvp37536	HX Stretch Cluster Witness VM reverts to DHCP on reboot.	3.5(1b)	3.5(2e)
CSCvq13099	HyperFlex post_install.py script generates incorrect message on run health check step.	3.5(2d)	3.5(2e)
CSCvq18919	Sometimes during Stretched Cluster failover, the cluster resource manager component goes offline due to file write error.	3.5(2d)	3.5(2e)
CSCvc38351 CSCvp55109	MTU changed to 1500 on vm-network after node expansion, which was set as 9000 earlier.	3.5(1a)	3.5(2c)

Resolved Caveats in Release 3.5(2d)

Defect ID	Symptom	First Release Affected	Resolved in Release
Management			
CSCvm58031	Tomcat and Nginx logs are not being collected in the support bundle generated through HX Connect in Release 3.5.	3.5(1a)	3.5(2d)
CSCvp40474	Multiple Hyper-V Hyperflex hosts are unable to access datastores, even when cluster is healthy.	3.5(2a)	3.5(2d)
CSCvp90129	On stretch clusters, when the cluster experiences a failure or a maintenance window that results in rebalance, some nodes may experience panics.	3.5(2c)	3.5(2d)

Resolved Caveats in Release 3.5(2c)

Please note that HX 3.5(2c) is deferred.

Defect ID	Symptom	First Release Affected	Resolved in Release
Management			
CSCvo36198	When retrieving a list of alarms, Virtual Machines or events, message appears intermittently - "Virtual Center unreachable" or "Resource information cannot be updated".	3.5(1a)	3.5(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvp58804	Generating support bundle(s) from HX Connect may leak memory and cause cluster health to diminish due to Out of Memory issues.	3.5(2a)	3.5(2c)
CSCvo75522	Losing data network on one site causes whole cluster to not be accessible. For more information, see the "*Preinstallation Checklist*" and "*Troubleshooting for Site-to-Site Failover*" section in the Cisco HyperFlex Systems Stretched Cluster Guide, Release 3.5.	3.0(1c) 3.5(1a)	3.5(2c)
CSCvp32000	On power outage of nodes, cache re-distribution may be done sub-optimally.	3.5(1a)	3.5(2c)
CSCvk46364	A node shuts down when two disks are replaced (a caching disk and another capacity disk), where the capacity disk is inserted first and then the caching disk is inserted.	2.6(1b)	3.5(2c)
CSCvo67207	HA for some VMs may be sub-optimal in a site-to-site network condition.	3.5(2a)	3.5(2c)
CSCvp42925	Any modification from the HX installer is causing LLDP to be disabled on all HX vNICs.	3.5(1a)	3.5(2c)
CSCvp41404	Incorrect space reporting on node failures and site-failure in stretch-cluster deployments.	3.0(1a) 3.5(1a)	3.5(2c)

Resolved Caveats in Release 3.5(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release
ESXi, Installation, Upgrade, Expansion, Management			
CSCvh08977	When an HX Snapshot is taken with the Quiesced Option, through HX Connect or through an external Backup Vendor, the Virtual Center VM Snapshot Manager will not list the snapshot as being quiesced. The Hyperflex implementation uses out of band quiescing that is not known to the VMware API. When an HX Snapshot API request succeeds, it is reasonable to expect that the snapshot quiesced correctly. Check with your Backup vendor if they rely on HX Quiesced Snapshots.	3.0(1c)	3.5(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvk17250	HX Drive PID HX-SD38TBE1NK9 has two different models with different sector sizes (8K and 4K). When the different drive versions are used together in releases 3.5.2a or below, Hyperflex node may experience panics. The following error message is seen in /var/log/springpath/debug-storfs.log: `storfs[1437:1538]: SUPPORT: PANIC: PANIC ON CONDITION (trustSegmentMetaData == true && CError_Ok(err) == false): PackingError at file /opt/git/cypress/src/modules/kvstore/kvindex.c line 1265` This can cause the cluster to become unrecoverable. The minimum HXDP version for configurations using the above drive is 3.5(2b) or above. Existing clusters running with this drive need to be upgraded to HXDP version 3.5(2b) or above before adding new nodes to the cluster or new drives to existing clusters.	3.0(1d)	3.5(2b) and later
CSCvn73127	Kernel migration fails when a local datastore is searched for in the ESXi host.	3.0(1d)	3.5(2b)
CSCvn37805	HX Replication clean up fails to complete and Zookeeper has stale replication configuration.	2.5(1a)	3.5(2b)

CSCvn17787	The cluster creation/cluster expansion workflow stops with the following error message at the validation step.
	
	`FIRMWARE-Check UCSC-SAS-M5HD`
	
	`FIRMWARE-Check UCSC-SAS-M5HD : Required: 00.00.00.29,00.00.00.32,00.00.00.35,00.00.00.50, Found: 00.00.00.58;`
	
	`Action Needed: Update the Controller Firmware to Required Version`
		3.5(2a)	3.5(2b)

Hyper-V			
CSCvh80044	HX Connect UI allows creation of a datastore by duplicating an existing datastore name that differs only in case. For example, Ds3, ds3, dS3 are allowed as valid datastore.	3.0(1a)	3.5(2b)
CSCvn60486	While upgrading a Hyper-V cluster, on account of a rare race condition between the **stUpgradeService** and Zookeeper servers, the upgrade orchestration throws an upgrade validation error, and the upgrade process is aborted.	3.5(2a)	3.5(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvn54300	During upgrade remove the VLAN on team created for user vSwitch. During fresh installation, only one VLAN tag is set to the vSwitch and team, although multiple VLANs were entered.	3.5(2a)	3.5(2b)

Resolved Caveats in Release 3.5(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: `There are locked drives on the system, unlock them and retry deployment.` In addition, after upgrade from Release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on checking `cluster readiness state` for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	3.0(1a)	3.5(2a)
CSCvk09073	During upgrade from Release 2.x to 3.x, if the upgrade encounters a failure, the cluster management IP address may no longer be reachable (not present on any of the controller VMs)	3.0(1a)	3.5(2a)
CSCvk46179	In a large cluster, "Server Call Failed" error may be seen in HX Connect.	3.5(1a)	3.5(2g)
CSCvm53972	During automatic software repair of a failing disk, a write command to the disk hung in the I/O subsystem. This prevented other nodes from communicating with this node for several minutes.	2.6(1b), 3.5(1a)	3.5(2a)
CSCvm66552	Multiple simultaneous 3.8TB SED SSD drive failures due to a drive firmware bug may cause the HX cluster to go offline. For more details, see the related Software Advisory.	3.0(1c)	3.5(2a)
CSCvm97558	Controller VM restarts due to Out-of-memory condition.	3.0(1c)	3.5(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvn07634	After attempting a Release 3.5(1a) node expansion on a cluster that was initially deployed prior to Release 3.5(1a), a vCON policy reference fault is triggered. In addition, any pre-existing service profiles created prior to expansion may be placed in a pending acknowledgment state.	3.5(1a)	3.5(2a)
CSCvn52412	When deploying Kubernetes for HX, in HX Connect UI when a user enables all nodes (selects **Settings** > **Integrations** > **Kubernetes** > **Enable All Nodes**), ESXi hosts cannot access datastores under the following conditions: • The HX cluster was deployed via Intersight, or, • The HX cluster is manually deployed with storage CMIP value configured on all ESXi hosts iscsi vmk.	3.5(1a)	3.5(2a)
CSCvn59508	During upgrade from earlier releases to Release 3.5(1a), cluster expansion may fail due to missing CRM entries related to node site map.	3.5(1a)	3.5(2a)
CSCvn59619	When an HX Quiesced Snapshot is taken, the ID of the snapshot returned may not be correct under certain circumstances and may impact your Backup vendors integration with HX Snapshots. Check with your Backup vendor if they rely on HX Quiesced Snapshots.	3.5(2a)	3.5(2a)
CSCvp52171 CSCvm60845	Using **stcli node remove** with node-ID option succeeded, but did not remove the node.	2.6(1d), 3.0(1c), 3.5(1a)	3.5(2a)
Hyper-V			
CSCvm59573	In some cases, Hyper-V OS installation fails during hypervisor configuration in the installation process.	3.5(1a)	3.5(2a)

Resolved Caveats in Release 3.5(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
ESXi, Installation, Upgrade, Expansion, Management			
CSCvk62990	HX deployments on M5 servers with ESXi version 6.0 may experience a PSOD during install or upgrade workflows.	2.6(1a)	3.5(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvk39622	HX Connect displays alarms with a "Lockdown mode enabled on one or more nodes in the cluster" message. In addition, the alarms are manually reset to green.	3.5(1a)	3.5(1a)
CSCvj90575	Smartmons tool reports **Reallocated_Sector_Ct** values for Samsung SATA drive (disk model **MZ7LM480HMHQ**)	3.0(1a)	3.5(1a)
CSCvi34303	HX Connect UI displays an error when any table is exported in `.csv` format and opened in excel.	3.0(1a)	3.5(1a)
Hyper-V			
CSCvm53679	HX Installer fails and **HXBootstrap.log** contains the following message: **"Unable to find a default server with Active Directory Web Services running."**	3.0(1e)	3.5(1a)
CSCvm42278	Datastore access leads to frequent alerts. In addition, the SMB SCVM client log file (`/var/log/springpath/debug-smbscvmclient.log`) shows a message similar to the following for host data IP address of the host on which that controller is hosted.	3.0(1e)	3.5(1a)
CSCvk18743	The Storage Controller VM is down for an extended period of time that may cause the VMs to power off.	3.0(1a)	3.5(1a)

Open Caveats in Release 3.5(2i)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvw21968	Support Bundle page in cross launched HX Connect and stand-alone HX Connect shows a blank page.	• Delete the Connected TAC generated support bundles in the path `/var/support/asup_default` for each controller VM in the cluster. These can be identified with "_intersight_" text in the filename. • Generate support bundles only from Standalone HX Connect (i.e logging to HX Connect on cip controller VM)	3.5(2i)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvu48941	HX cluster outages when there is a CATERR, if no-drop class is configured for storage traffic.	N/A	3.5(2h)

Open Caveats in Release 3.5(2h)

There are no open caveats in this release.

Open Caveats in Release 3.5(2g)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs35307	Bootstrap from HX 2.1x or lower to HX 3.5(2g) fails, incompatible Java version.	Please open a service request to have corrective action performed, then upgrade to HX 3.5(2f).	3.5(2g)
CSCvs02466	M.2 boot disk is missing from server inventory after upgrade to server firmware 4.0(4e). As a result server fails to boot to OS installed in the M.2 disk. The issue persists after re-acknowledgement as well as de-commission and re-acknowledgement of the server.	1. De-commission the server 2. Power drain the server - REMOVE BOTH POWER CORDS ON THE BACK OF THE SERVER FOR 10 SECONDS, THEN REINSERT POWER CORDS. 3. Re-commission the server.	4.0(4e)
CSCvq38279	Hyper-V: When replicated DC was used, then during install time, Windows failover cluster was not created successfully.	Clean up the fail over cluster and then recreate the fail over cluster. No need to touch the HX storage cluster.	3.5(2e)
CSCvq65830	Hyper-V: VM corrupted after migrating from one host to another.	NA	3.5(2e)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvp98910	After 2 node network partition heals, datastore on the isolated node is not available for 15 minutes.	There are two options: 1. Wait 15-20 minutes for the ARP entry to time out, then the datastore will become mounted again. 2. Manually reset the ARP entry by following these steps: Check the current eth1:0 mac address from both HyperFlex Controller VMs and node the IP/Mac. Run **ifconfig eth1:0** from the shell of the controller VM. On the node where the datastore is unavailable, check the mac address entries for the above IP addresses in the ESXi ARP cache using the following command: **esxcli network ip neighbor list**. If the IP address is assigned to the incorrect mac address, purge the entry from the ARP table as shown below. `esxcli network ip neighbor remove -a <IP_address> -v 4`	3.5(2g) 4.0(1a)

Open Caveats in Release 3.5(2f)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvq94462	Multiple DNS servers causing expansion validation error; stcli to validate.	NA	3.5(2d)

Open Caveats in Release 3.5(2e)

Defect ID	Symptom	Workaround	Defect Found in Release

CSCvv21905	UCSM Read-Only user missing error in the encryption page on HX Connect UI, later while authenticate UCSM with credentials it throws an error of invalid CSRF token.	Run the following commands:
		
		`stcli security encryption ucsm-ro-user create --hostname <FI-IP> --username <FI-user-name> --password <FI-password>`
		
		`stcli security encryption ucsm-ro-user show`
			3.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CCvq53058 CSCvq49412	Cluster may observe a temporary All Paths Down (APD) condition during site failover depending upon the used storage capacity and the latency on the link to the Witness node.	None in Release 3.5(2e). However, if an APD occurs, the condition will clear up on its own after some time and IO will resume. The fix to avoid the APD will be available in an upcoming release.	3.5(2a)
CSCvq532080 CSCvq54992	HXDP support for EMC Recover Point – Manual workaround required.	Manually install/upgrade scvmclient VIB and make sure that it matches the current HX Data Platform version.	3.5(2b)
CSCvm99150	Changing MTU from 9000 to 1500 on one ESX node causes storfs process to restart on all nodes.	1. Do not change MTU at ESX level on running cluster. 2. Reset back to the original value.	3.5(1a) 3.5(2a)
CSCvn73383	A cluster is shutdown for maintenance. However, due to an existing defect, disks on one node were not published when that node was rebooted again. The cluster failed to come back online.	This workaround is to fix the disk failures on that node, so that all disks are published and re-start the cluster again.	3.0(1c)
CSCvo70650	The cluster expand fails on a node with DR replication configured. When an HX cluster which has DR replication configured is expanded we see the installer UI pulling in the replication VLAN information instead of the management VLAN information. Even if we change that information to the correct mgmt VLAN id and name it does not seem to work as the node is configured with the VLAN of the replication VLAN in ESXi. This leads to the failure of the node add with the host unreachable error.	Have to identify that the replication VLAN is being used. Then log to the KVM and update the MGMT VLAN to the correct VLAN ID and retry the cluster expand.	3.5(2a)
CSCvq11456	HyperFlex stcli cluster info does not show UCSM VIP address.	NA	3.5(2d)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvq22844	Add message to HX Connect Progress flow to not Acknowledge Pending Activities and reboot the servers on UCSM. HX Connect is doing a controlled rolling server upgrade in the background.	NA	3.5(2d)
CSCvn11045	HX node keeps crashing after node restarted.	Verify if the interface is up and that you can ping the loopback interface by running the following commands: **ifconfig -a** **ping 127.0.0.1** - bring up the loopback interface **ip link set lo up** - check service is running **status scvmclient** **status storfs** - start below services **start scvmclient** **start storfs**	3.5(1a) 3.0(1e)

Open Caveats in Release 3.5(2d)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvp93442	storfs crash with (tune.fileAIOPanicOnStuckIO): IO request timer expired.	NA	3.5(2b)
CSCvp96650	Sub-org field entry is not being fetched during cluster expansion in the UCSM configuration step. The input field for sub-org in the UCSM page shows up empty.	In the JSON configuration file, enter the sub-org name and then export that configuration. Then the validation will not fail. Use the import json configuration file feature to import the json with the correct sub-org name detail so that the value gets picked up during validation and does not result in failure.	3.5(2b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvo86431	When a node is in Maintenance Mode, any disk removal or replacement will be reflected in UI only after the node is brought back from maintenance mode. This is because storfs is not running on the node in maintenance, and will not be able to detect disk activities until it is brought out of MM.	Bring the node out of Maintenance Mode.	3.5(2a)
CSCvp79511	HX Upgrade Release 2.1(1c) > 3.0(1i) was allowed when ESXi/vCenter was not on 6.0u3 or later.	Upgrade ESXi and vCenter to 6.0u3 before attempting upgrade.	3.0(1i)

Open Caveats in Release 3.5(2c)

Defect ID	Symptom	Workaround	Defect Found in Release
Install, Upgrade, Expand			

CSCvp88990	HyperFlex upgrade fails due to SCVM unable to login to ESXi.	Copy the SSH public key to the authorized key file on the ESXi host
		
		On SCVM:
		
		`cat /etc/ssh/ssh_host_rsa_key.pub`
		
		**Copy all output that is a single line, do not copy anything past the end of the key even if your terminal window does not put the hostname on a new line
		
		Input the key from the previous output into ESXi's authorized key file:
		
		`vi /etc/ssh/keys-root/authorized_keys`
		
		*after entering only the above output into file, exit and save the file
		
		Now when you try to SSH from the SCVM to ESXi it should be permitted without asking for password. Try to continue upgrade. If upgrade continues to fail, contact Cisco TAC
			3.0(1i)
			
			3.5(2c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvj22992 CSCvs26028	VM appears on multiple nodes.	To recover the VM, copy over the data disks and attach them to the new VM.	3.0(1b)
CSCvn99088	Shavlik snapshot causes protected VM to become unprotected.	Delete the Shavlik snapshot.	3.0(1a)
CSCvo19250	No warnings/alerts generated when cluster capacity is above 70% utilization.	Increase storage capacity (new node or disks) or reduce storage usage (delete unused VMs and snapshots).	3.0(1i)
CSCvo56350	When upgrading to ESXi 6.7 EP06 a PSOD is encountered after a reboot.	Upgrade to ESXi 6.7 Express Patch 08.	4.0(1a)
CSCvo83276	VM powers off during backup VM snapshot.	Retake the snapshot.	3.5(1a)
CSCvp19670	Cluster MGMT IP does not come back up after cluster shutdown.	Start cip-monitor on each storage controller VM. Manually mount datastores.	3.5(2a)
CSCvp26319	Release 3.5(2b) FlexVol to 4.0 CSI upgrade does not work. FlexVol continues to work.	1. Update the conf file manually to change to the link local address to *. 2. Restart scvmclient on the controller and ESX.	4.0(1a)
CSCvp49720	APD when putting zk leader to MM on a 2 node ROBO setup.	Stop zookeeper on the node, make a copy of the zookeeper database, delete the files in the database directory and restart zookeeper.	4.0(1a)
CSCvp60476	After node expansion, zookeeper database not updated with node information, and as a result 'Node IP Settings' in 'stcli cluster info' does not show information for expanded node.	Perform `# stcli node add --node-ips 1x.x.x.x --esx-username root -f` from SSH to CIP with root username.	3.5(2a)
CSCvp42679	HyperFlex: UCSM Upgrade in Queued State when current and desired firmware are identical.	Upgrade needs to be forcefully cleaned. Run the following command: `# stcli cluster upgrade --clean --components ucs-fw --ucsm-host 14.39.51.225`	3.5(2b)
CSCvp62167	After power outage cluster hits panic	NA	3.5(2a)
CSCvp63635 CSCvm36103	Unable to re-enable SSH after upgrade.	NA	3.5(2a)
CSCvp65649	Limited access as Admin user.	NA	3.5(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvp65824	Incorrect nodes status during upgrade from Release 3.5(2b) to 4.0(1a).	NA	4.0(1a)
Management			
CSCvp09978	Cluster info shows Smart call home is enabled, even though it is disabled.	Use the command **stcli services sch show** instead.	3.5(2b)

Open Caveats in Release 3.5(2b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs08218	When Commvault backs up virtual machines running within a HyperFlex cluster, native HyperFlex snapshots are taken and removed upon completion. However, vCenter reports "VM disk consolidation is required" within vSphere. This alert is thrown for the existing sentinel snapshot that is not updated regularly. Commvault needs sentinel snapshot to be present for native HX snapshots to be taken. Commvault also reports error as part of the data aging job: HX Failed to delete Snap. Error Message - Delete snapshot snapshot-100 failed for entity vm-xx. A previously issued Snapshot Now task is in progress on this entity. Fix for this issue is available in Hotfix Pack SP16.36.	1. Take another snapshot via Commvault. The error message goes away. 2. Delete the sentinel snapshot and create again (this might stun the VM). 3. Ignore/suppress alert on vCenter.	3.5(2b)
CSCvp78288	When adding disks to cluster I/O freezes for 96 seconds.	No intervention required, cluster recovers on its own.	3.5(2a)
CSCvp89523	When a HyperFlex Cluster is deployed on an ACI Fabric, with HX MGMT and HX Storage in the same Bridge Domain, we see ARP responses from wrong HX Storage Controller interface, causing incorrect ARP learning to occur. On the ACI Fabric, we see MAC moves for HX IPs.	HX MGMT and HX Storage should be in separate Bridge Domains. Reference Table 21 in the following CVD:	3.0(1i)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvp93442	storfs crash with (tune.fileAIOPanicOnStuckIO): IO request timer expired.	Contact TAC.	3.5(2b)
CSCvp96650	Sub-org field entry is not being fetched during cluster expansion in the UCSM configuration step. The input field for sub-org in the UCSM page shows up empty.	In the JSON configuration file, enter the sub-org name and then export that configuration. Then the validation will not fail. Use the import json configuration file feature to import the json with the correct sub-org name detail so that the value gets picked up during validation and does not result in failure.	3.5(2b)
CSCvo86431	When a node is in Maintenance Mode, any disk removal or replacement will be reflected in UI only after the node is brought back from maintenance mode. This is because storfs is not running on the node in maintenance, and will not be able to detect disk activities until it is brought out of MM.	Bring node out of Maintenance Mode.	3.5(2a)
CSCvp79511	HX upgrade to Release 3.0(1i) was allowed while vCenter and ESXi were both on version 6.0u2 although version check requires version 6.0u3.	Upgrade ESXi and vCenter to 6.0u3 before attempting upgrade	3.0(1i)
CSCvp86483	Cannot reset HyperFlex root or admin passwords without knowing existing password.	Contact TAC.	4.0(1a)

Open Caveats in Release 3.5(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
Install, Upgrade, Expand			
CSCvp78288	When adding disks to cluster I/O freezes for 96 seconds.	No intervention required, cluster recovers on its own.	3.5(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP address) should occur before adding the node to the cluster.	Ensure there are sufficient replication IP addresses available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: `# CIMC-soft-rest`	2.5(1a)
CSCvb94112 CSCvb91660	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	1. Check logs to verify that the expansion workflow is hung. 2. In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
Hyper-V			
CSCvm05523	In rare cases, live migration of VMs fails with error code 0x138D.	Refer to the following workarounds from Microsoft and retry the operation: • Server 2016 S2D Cluster unable to Drain Role • Live migrations fail during drain from Cluster-Aware Updating • Draining Nodes for Planned Maintenance with Windows Server 2012	3.0(1e)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, "`The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvh25238	Hyper-V: During HX Data Platform deployment, the controller VM may be assigned with only one DNS address if you add more than one IP address for DNS.	Usually the primary DNS is sufficient for the HX Controller VM to work. If you need additional DNS, edit the eth0 interface file in the controller VM to add the additional DNS.	3.0(1a)
Management			
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
Encryption			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf17183 CSCvd000362	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCS Manager might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)

Open Caveats in Release 3.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
Install, Upgrade, Expand			
CSCvx37435	A vulnerability in the web-based management interface of Cisco HyperFlex HX Data Platform could allow an unauthenticated, remote attacker to execute arbitrary commands on an affected device.	Please refer to the Security Advisory.	3.5(1a)
CSCvx36014	A Unauthorized Remote Code Execution (login request) could allow an unauthenticated, remote attacker to perform a command injection attack against an affected device.	Please refer to the Security Advisory.	3.5(1a)
CSCvx36019	A Unauthorized Remote Code Execution (login request) could allow an unauthenticated, remote attacker to perform a command injection attack against an affected device.	Please refer to the Security Advisory.	3.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvx36028	A Unauthorized File Upload Vulnerability could allow an unauthenticated, remote attacker to upload files to an affected device.	Please refer to the Security Advisory.	
CSCvx52126	A Unauthorized File Upload Vulnerability could allow an unauthenticated, remote attacker to upload files to an affected device.	Please refer to the Security Advisory.	
CSCvo69067	Adding drive to a cluster fails to increase cluster capacity.	Contact Cisco TAC for resolution.	3.5(2b) 3.5(2h)
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: "There are locked drives on the system, unlock them and retry deployment." In addition, after upgrade from Release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more Release 3.0(1c) information to recover from this issue.	3.0(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: `# CIMC-soft-rest`	2.5(1a)
CSCvb94112 CSCvb91660	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	1. Check logs to verify that the expansion workflow is hung. 2. In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
Hyper-V			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvm05523	In rare cases, live migration of VMs fails with error code 0x138D.	Refer to the following workarounds from Microsoft and retry the operation: • Server 2016 S2D Cluster unable to Drain Role • Live migrations fail during drain from Cluster-Aware Updating • Draining Nodes for Planned Maintenance with Windows Server 2012	3.0(1e)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, `"The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process."`	Wait for a minutes after datastore creation and retry the command.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvh25238	During HX Data Platform deployment, the controller VM may be assigned with only one DNS address if you add more than one IP address for DNS.	Usually the primary DNS is sufficient for the HX Controller VM to work. If you need additional DNS, edit the eth0 interface file in the controller VM to add the additional DNS.	3.0(1a)
Management			
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the 3.0(1b) following: 1. Delete the dummy stCtlVMsfrom the vCenter. 2. Cleanup the old extensions. 3. Re-register to the original vCenter.	3.0(1e)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvg47332	Using the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated` Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
Encryption			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCS Manager might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)

Related Caveats

Defect ID	Symptom	Defect Found in Release	Resolved in Release
CSCvq41985	When attempting to install ESXi 6.5 or 6.7 from a CIMC mounted ISO with an embedded kickstart file, the installation may fail when reading the embedded KS.CFG file. In the ESXi installer, a popup error will state: Could not open file <path>/KS.CFG	Cisco IMC 4.0(1a)	Open

Revision History

Release	Date	Description
3.5(2i)	May 7, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.0(4l) qualified for HX 3.5(2i) and 3.5(2h).
3.5(2i)	April 29, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.1(3c) qualified for HX 3.5(2i).
3.5(2i)	March 17, 2021	Updated link to indicate UCSM 4.1(2f) is the recommended Host Upgrade Utility (HUU) for M5 for HX 3.5(2x).

Release	Date	Description
3.5(2i)	March 11, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.0(4k) is the recommended release.
3.5(2i)	February 19, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.1(2c) qualified for HX 3.5(2i) and HX 3.5(2h).
3.5(2i)	February 10, 2021	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate new new UCSM versions qualified for HX 3.5(2x) release.
3.5(2i)	January 7, 2021	Updated Software Requirements for VMware ESXi to indicate limitations using ESXi 6.7 U3.
3.5(2i)	December 22, 2020	Updated Software Requirements for VMware ESXi to indicate limitations using ESXi 6.7 U3.
3.5(2i)	December 7, 2020	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate new UCSM versions qualified for HX 3.5(2i), HX 3.5(2h), HX 3.5(2g), HX 3.5(2f) and HX 3.5(2e) releases.
3.5(2i)	October 22, 2020	Created release notes for Cisco HX Data Platform Software, Release 3.5(2i).
3.5(2h)	September 24, 2020	Updated Recommended FI/Server Firmware - 3.5(x) Releases to indicate UCSM 4.1(1e) qualified for HX 3.5(2g) and HX 3.5(2h) releases.
3.5(2h)	September 14, 2020	Updated Recommended FI/Server Firmware - 3.5(x) Releases Recommended Firmware Version to UCSM 4.1(1h).
3.5(2h)	September 4, 2020	Updated Recommended FI/Server Firmware - 3.5(x) Releases Recommended FI/Server Firmware to UCSM 4.0(4i) for HX 3.5(2g) and HX 3.5(2h) releases.
3.5(2h)	August 11, 2020	• Added column for M4/M5 Qualified FI/Server Firmware. Listed USCM 4.1(2a) as qualified for HX 3.5(2h). • Added CSCvv21905 to the list of Open Caveats for HX 3.5(2h), 3.5(2g), 3.5(2f), 3.5(2e), 3.5(2d), 3.5(2c), 3.5(2b), 3.5(2a), 3.5(1a).
3.5(2g)	July 23, 2020	Updated Recommended FI/Server Firmware - 3.5(x) Releases. Added qualification for Cisco UCS Manager 4.0(4i) and 4.1(1d).
3.5(2h)	July 16, 2020	Added Resolved Caveats for HX 3.5(2h), HX 3.5(2g), HX 3.5(2f) and HX 3.5(2e).

Release	Date	Description
3.5(1a)	May 15, 2020	HX 3.5(1a) - End of Life
3.5(2h)	May 4, 2020	• Updated Host Upgrade Utility (HUU) for M5 to UCS 4.0(4k) for HX 3.5(2h). • Added description of HX REST API Access Token Management in Guidelines and Limitations section.
3.5(2h)	March 30, 2020	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4h) for HX 3.5(2g).
3.5(2h)	March 24, 2020	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4h) for HX 3.5(2h). Identified HX 3.5(1a) as Unsupported. Added CSCvs08218 to the list of Open Caveats in HX 3.5(2b).
3.5(2h)	March 5, 2020	Added upgrade reminder for HyperFlex versions 3.0(1x) and 3.5(1x) as these versions are unsupported and have been declared end-of-life. Removed references to HX 3.0 requirements.
3.5(2h)	January 28, 2020	Added CSCvs47419 to the list of Resolved Caveats in HX 3.5(2h).
3.5(2h)	January 21, 2020	Created release notes for Cisco HX Data Platform Software, Release 3.5(2h).
3.5(2c)	January 14, 2020	Updated release notes for deferred Cisco HyperFlex Release HX 3.5(2c).
3.5(2g)	January 10, 2020	Added CSCvs35307 to the list of Open Caveats for HX 3.5(2g).
3.5(2g)	December 23, 2019	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4e) for HX 3.5(2f), 3.5(2e), and 3.5(2d).
3.5(2g)	December 13, 2019	Added CSCvs28167 to the list of Resolved Caveats for HX 3.5(2g), 3.5(2f), 3.5(2e). Added CSCvs28167 to the list of Open Caveats for HX 3.5(2d), 3.5(2c), 3.5(2b), 3.5(2a), 3.5(1a).
3.5(2g)	November 25, 2019	Added CSCvs02466 to the list of Open Caveats.
3.5(2g)	November 19, 2019	Updated info in the Upgrade Guidelines section.
3.5(2g)	November 7, 2019	Updated info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments section.
3.5(2g)	October 25, 2019	Added CSCvj95606 and CSCvq24176 to the list of Security Fixes.

Release	Date	Description
3.5(2g)	October 21, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2g).
3.5(2f)	September 17, 2019	Added CSCvq41985 to new section for "Related Caveats".
3.5(2f)	September 10, 2019	Updated HUU/CIMC info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments.
3.5(2f)	September 6, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2f).
3.5(2f)	September 6, 2019	Updated HUU/CIMC recommended firmware versions for HyperFlex Releases 3.5(2e) and 3.5(2d).
3.5(2e)	August 23, 2019	Updated Recommended FI/Server Firmware versions for HyperFlex Releases 3.5(2e) and 3.5(2d).
3.5(2e)	August 21, 2019	Added Cisco IMC version support info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments.
3.5(2e)	August 8, 2019	Added bullet describing the "Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases" in the Upgrade Guidelines section.
3.5(2e)	August 2, 2019	Added important note indicating HyperFlex does not support UCS server firmware 4.0(4a), 4.0(4b), and 4.0(4c).
3.5(2e)	July 22, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2e).
3.5(2d)	July 15, 2019	• Updated UCS Manager interoperability for SED-based HyperFlex systems in "Supported Versions and System Requirements" section. • Updated Release 3.5(2b) support for ESXi 6.7 U2 in "Supported VMware vSphere Versions and Editions".
3.5(2d)	June 20, 2019	Added HyperFlex Edge and Firmware Compatibility Matrix tables.
3.5(2d)	June 19, 2019	• Added CSCvp40474 to the list of "Resolved Caveats in Release 3.5(2d)". • Updated workaround information for CSCvp40474 in list of "Open Caveats in Release 3.5(2b)".
3.5(2d)	June 11, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2d).

Release	Date	Description
3.5(2c)	May 31, 2019	Moved CSCvo36198 to the list of "Resolved Caveats in Release 3.5(2c)".
3.5(2c)	May 29, 2019	• Added important note about not using Release 3.5(2c) with Stretch Clusters. • Added CSCvp90129 to the list of "Open Caveats in Release 3.5(2c)" section.
3.5(2c)	May 21, 2019	• Added note describing UCS/UCSM Install issue CSCvo13678. • Added reference to Stretched Cluster Guide for CSCvo75522 in the "Resolved Caveats in Release 3.5(2c)" section. • Updated link to Hypercheck TechNote article in "Upgrade Guidelines" section. • Updated CSCvo70723 in "Open Caveats in Release 3.5(2b)" section.
3.5(2c)	May 20, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2c).
3.5(2b)	May 16, 2019	Added CSCvp58804 to the list of "Open Caveats in Release 3.5(2b)".
3.5(2b)	May 2, 2019	Added CSCvk38003 to the list of "Resolved Caveats in Release 3.5(2b)".
3.5(2b)	March 22, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2b).
3.5(2a)	January 8, 2019	Created release notes for Cisco HX Data Platform Software, Release 3.5(2a).
3.5(1a)	November 21, 2018	Revised format for the "Cisco HX Data Platform Storage Cluster Specifications" section.
3.5(1a)	November 6, 2018	• Added vCenter version 6.7U1 in the "Supported vSphere Versions" section. • Removed deprecated releases - 1.8, 2.0, 2.1, 2.5 from the "Supported vSphere Versions" table. • Added CSCvn07634 to the list of "Open Caveats in Release 3.5(1a)".
3.5(1a)	October 16, 2018	Created release notes for Cisco HX Data Platform Software, Release 3.5(1a).

Related Documentation

Document	Description
Preinstallation Checklist	Provides an editable file for gathering **required** configuration information prior to starting an installation. This checklist must be filled out and returned to a Cisco account team.
Installation Guide for VMware ESXi	Provides detailed information about Day 0 configuration of HyperFlex Systems and related post cluster configuration tasks. It also describes how to set up multiple HX clusters, expand an HX cluster, set up a mixed HX cluster, and attach external storage.
Stretched Cluster Guide	Provides installation and configuration procedures for HyperFlex Stretched cluster, enabling you to deploy an Active-Active disaster avoidance solution for mission critical workloads.
Installation Guide on Microsoft Hyper-V	Provides installation and configuration procedures on how to install and configure Cisco HyperFlex Systems on Microsoft Hyper-V.
Edge Deployment Guide	Provides deployment procedures for HyperFlex Edge, designed to bring hyperconvergence to remote and branch office (ROBO) and edge environments.
Administration Guide	Provides information about how to manage and monitor the cluster, encryption, data protection (replication and recovery), ReadyClones, Native snapshots, and user management. Interfaces include HX Connect, HX Data Platform Plug-in, and the `stcli` commands.
HyperFlex Intersight Installation Guide	Provides installation, configuration, and deployment procedures for HyperFlex Intersight, designed to deliver secure infrastructure management anywhere from the cloud.
Upgrade Guide	Provides information on how to upgrade an existing installation of Cisco HX Data Platform, upgrade guidelines, and information about various upgrade tasks.
Network and External Storage Management Guide	Provides information about HyperFlex Systems specific network and external storage management tasks.
Command Line Interface (CLI) Guide	Provides CLI reference information for HX Data Platform `stcli` commands.
Cisco HyperFlex PowerShell Cmdlets for Disaster Recovery	Provides information on how to use the Cisco PowerShell Cisco HXPowerCLI cmdlets for Data Protection.
REST API Getting Started Guide REST API Reference	Provides information related to REST APIs that enable external applications to interface directly with the Cisco HyperFlex management plane.
Troubleshooting Guide	Provides troubleshooting for installation, configuration, to configuration, and to configuration. In addition, this guide provides information about understanding system events, errors, Smart Call Home, and Cisco support.

Document	Description
TechNotes	Provides independent knowledge base articles.