Release Notes for Cisco HX Data Platform, Release 4.0

First Published: 2019-05-02

Last Modified: 2021-08-25

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 that integrates smoothly with Cisco Application Centric Infrastructure (Cisco ACI). 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 4.0, and describe the features, limitations and caveats for the Cisco HX Data Platform.

Recent Revisions

For the complete revision history, see Revision History, on page 90.

Release	Date	Description
4.0(2f)	August 25, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.1(3e) is qualified for HX 4.0(2x) releases.
4.0(2f)	August 9, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.0(4m), and 4.1(3d) are qualified for HX 4.0(2x) releases.
4.0(2f)	June 21, 2021	Added Single Socket support in New Features, on page 2.
4.0(2f)	June 3, 2021	Created release notes for Cisco HX Data Platform Software, Release 4.0(2f).

Release	Date	Description
4.0(2e)	May 7, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.0(4l) is qualified for HX 4.0(2x) releases.

New Features

New Featues in Release 4.0(2f)

- Single Socket for Stretched Cluster Configurations: This allows users to optimize the hardware configuration cost and licensing cost for certain applications for a stretched cluster configuration. This support was introduced in HXDP 4.0(2f).

Note Limited options for cache drives available with single socket configurations.

Note Single socket stretch cluster nodes are not supported.

New Features in Release 4.0(2e)

- **Cisco HyperFlex HTML5 Plugin for VMware vCenter**—Provides users the ability to manage and monitor your HyperFlex clusters from the VMware vCenter Web UI. Additional functionally in version 2.1.0 includes:
  - Nodes and Disk View
  - Virtual Machines Summary
  - Events and Tasks
  - VLAN Creation
  - Rename Cluster

Note HXDP Release 4.5(1a) is the final release that supports the Cisco HyperFlex Flash Plugin. This change coincides with the end of flash support in popular browsers. It is recommended that users upgrade to the Cisco HyperFlex HTML5 Plugin 2.1.0.

- **New Drive support**—Support for new SED cache drive for Hybrid 240 M5: HX-SD16TBHNK9.

New Features in Release 4.0(2d)

- There are no new software features in this release.

New Features in Release 4.0(2c)

Cisco HX Data Platform, Release 4.0 provides the following features. These features, including the Invisible Cloud Witness for HyperFlex Edge clusters, are supported both on the Intersight Virtual Appliance and on Intersight.com.

- **Cisco HyperFlex CSI Interoperability Metrics Update**—Added support for Kubernetes Version 1.17, and CCP, and Anthos Versions.
- **Cisco HyperFlex HTML5 Plugin for VMware vCenter**—Provides users the ability to manage and monitor your HyperFlex clusters from the VMware vCenter Web UI. Functionally includes:
  - Discover the Registered HX Cluster
  - View HX Cluster Summary
  - Events and Alarms
  - HX Datastore Management
  - Create HX Snapshots and clones at the virtual machine level
  - Manage users and access to HX Clusters
  - Cross launch HX Connect for upgrade
  - Embedded vCenter server actions at Host and Clusters level

  Note Version 2.0.0 replaces the 1.0.1 HTML Plugin version and users need to upgrade.

- **HyperFlex Edge Short Depth Servers**—New, short depth server offerings are now available for HyperFlex Edge. Both All-flash (HXAF240c-M5SD) and Hybrid (HX240c-M5SD) configuration options are available. See the HyperFlex HX*240 SD Short Depth Edge Nodes spec sheet for full details.
- **New Drive support**—New 7.6TB SED SSD capacity point introduced (HX-SD76TBEM2NK9). Also new drives SKUs for existing drive capacity points 3.8TB (HX-SD38TBEM2NK9) and 960G (HX-SD960GBM2NK9).
- **Cluster Scale Limits increase**—Support for maximum scale limit increase on a cluster on 7.6TB SSDs. See Cisco HX Data Platform Compatibility and Scalability Details - 4.0(x) Releases, on page 16
- **Intel Optane DC Persistent Memory Support**—Support on all NVMe (HXAF220-M5SN) and All Flash (HXAF220-M5SX and HXAF240-M5SX) for App-Direct mode (CS-DCPMM-AD).
- **NVIDIA RTX 6000/RTX 8000 Support**—Support on HX240 and HXAF240.
- **Single Socket Configurations Support**—Support on HX240 and HXAF240.

  Note Limited options for cache drives available with single socket configurations.

  Note Single socket stretch cluster nodes are not supported.

New Features in Release 4.0(2b)

- **7.6TB SSD data drive**—Support on HX Edge configurations.
- **Support for UCS Fabric Interconnects with limited cluster**—UCS Fabric Interconnect (FI-64108) with limited scale of 32 nodes now supported.
- **All NVMe and All Flash limits increase**—
  - Maximum cluster size for All NVMe (with 1TB, 4TB, or 8TB data drives) increased to 32 nodes.
  - Maximum cluster size for HX220 All Flash with 7.6TB data drive increased to 32 nodes.
  - HX240 All Flash increased to 16 nodes for full drive population (23 data drives/node) or 32 nodes at up to 12 drives/node.
- **Cluster Scale Limits increase**—Support for maximum scale limit increase on a cluster. See Cisco HX Data Platform Compatibility and Scalability Details - 4.0(x) Releases, on page 16
- **HW Offload option**—Support for Hardware Offload option with Stretched cluster configurations.
- **Cisco Overlay Transport Virtualization for Stretched Cluster**—Support for OTV as an overlay for Stretched Cluster.

New Features in Release 4.0(2a)

- **Boost Mode**—This release introduces Boost Mode for the following configurations: All NVMe, All Flash C240, All Flash C220, and Hypervisor: ESX. Boost Mode allows the Cisco HyperFlex cluster to deliver higher IOPs by increasing the storage controller VM CPU resources by 4 vCPU. For configuration information, see the Cisco HyperFlex Data Platform Administration Guide.
- **Cisco HyperFlex HTML plug-in for VMWare vCenter**—Enables virtualization administrator to manage and monitor the Cisco HyperFlex physical infrastructure by cross launching HyperFlex Connect from the vSphere Client UI and perform management actions in the HyperFlex Connect UI.
- **25GE networking for HX Edge**—Support for 25GE networking for HX Edge.
- **All NVMe with Stretched Cluster**—Support for All NVMe with Stretched Cluster (ESX only).
- **Cluster Upgrade Eligibility Test**—This release adds the capability to perform a pre-upgrade test which checks for cluster readiness before upgrading. Example checks include: validating cluster state, rebalance status, controller VM Free Space, ESXi version, and much more. The Eligibility test is intended to help avoid unexpected problems that may arise during the upgrade process. It is highly recommended to run the test before performing the Hyperflex upgrade.
- **Registering Smart Software Licensing**—This release adds support for software that allows easy tracking of the status of license and software usage trends and simplifies the three core licensing functions: Purchasing, Management, and Reporting.
- **Dynamic self-signed certificate generation enhancements**—This release adds support for Self-signed SSL certificates on the Controller VMs, which were static in prior releases. The static certificates are replaced with dynamically generated self-signed certificates upon upgrading to HXDP 4.0(2a) so that the certificates are unique per cluster. The new clusters installed with HXDP 4.0(2a) have dynamically generated self-signed certificates.
- **Test Upgrade Eligibility**—This release adds support for testing your cluster readiness and infrastructure compatibility for an upgrade. For more information, see the **Test Upgrade Eligibility** sections in the

Cisco HyperFlex Systems Upgrade Guide for VMware ESXi, Release 4.0 or the Cisco HyperFlex Upgrade Guide for Microsoft Hyper-V, Release 4.0.

Disaster Recovery

- **Recovery Settings Configuration**—This release supports configuration of recovery settings to define global recovery parameters and mapping for resources across recovery sites. These parameters are used during recovery, test recovery and migrate operations.
- **HyperFlex DR Powershell Runbook**—PowerShell runbook functionality is extended to support the recovery configuration settings in the runbooks for all recovery scenarios. New-HXrunbook cmdlet can now be used to generate runbook for a single or multiple protection groups. In addition, two new cmdlets, Wait-HXTask and Get-HXTaskStatus, are introduced.
- **Protected Virtual Machine Scalability**—This release adds support for 1500 VMs across both clusters and 750 VMs per cluster in a bi-direction or any split between the two clusters without exceeding the limit of 1500 VMs. For more information, see Cisco HyperFlex Data Platform Administration Guide.
- **System Management REST API enhancements**—Pause data replication actions briefly using REST API to explicitly inform users on the current status of replication actions.

Cisco HyperFlex with Data Platform for Hyper-V

- **Cluster-Aware Updating (CAU)**—This is an automated feature that allows you to perform updates on windows servers in a failover cluster with little or no loss in availability during the upgrade process.

New Features in Release 4.0(1b)

- **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).

New Features in Release 4.0(1a)

The following new features are in Release 4.0(1a).

- **Ultra-Light HyperFlex Edge Clusters**—This release introduces support for two-node HyperFlex Edge clusters, enabling HyperFlex to run in environments requiring a small footprint. Cisco Intersight provides comprehensive lifecycle management and includes remote cloud-based installation, centralized upgrades, and invisible witnessing. Both 1GE and 10GE networking topology options are available.
- **Scaled-Up HyperFlex Edge Clusters**—This release adds support for four node HyperFlex Edge clusters, enabling a full range of size options for remote and branch offices. Size the branch office environments to suit current needs with a two, three, or four node HyperFlex Edge cluster. Cisco Intersight provides full-lifecycle management and 1GE, and 10GE networking options are available.
- **Cisco Intersight Invisible Cloud Witness**—For two node clusters, this feature eliminates the need for witness VMs, the infrastructure to run those VMs, and the management overhead to deploy, scale, and patch witnessing software. The Invisible Cloud Witness is responsible for maintaining cluster HA in the event of failure scenarios. This feature is included at no additional cost and is automatically deployed and managed by Cisco Intersight.
- **Cloud-delivered HyperFlex Edge Upgrades**—Powered by Cisco Intersight, this feature adds support for multi-site orchestrated remote upgrades of the HyperFlex Data Platform. This feature will be enabled with the next 4.0 patch release and will allow HyperFlex Edge clusters deployed via Intersight to perform orchestrated upgrades across one or many sites in parallel.

- **All-NVMe HyperFlex**—Starting with this release, a new, high-end performance node powered by all NVMe drives is available for HyperFlex clusters. Co-engineered with Intel to support Intel VMD for hot-plug and surprise removal, this offering represents an industry first: an enterprise-ready and fully validated all-NVMe HCI appliance. The all-NVMe offering is available in the 220 form factor (1RU) and is powered by Intel Optane cache drives for maximum performance and highest endurance.
- **VMware Site Recovery Manager (SRM) Integration**—This release brings support for a Cisco developed Storage Replication Adapter (SRA) for SRM. The SRA provides the ability to leverage HyperFlex native async replication with the powerful orchestration and runbook capabilities of SRM. The SRA includes the ability to perform test recoveries, planned migrations, and full disaster recovery.

  Note HX SRA is certified by VMware and is available for download from VMWare SRM site.

- **HyperFlex DR Powershell Runbooks**—New Powershell cmdlets are included for automated runbook generation when using HyperFlex native disaster recovery. The New-HXRunbook cmdlet supports the following workflows: Test Recovery, Planned Migration, and Disaster Recovery. These runbooks can be used to orchestrate DR workflows without the requirement for any third-party software.
- **Windows Server 2019 with Hyper-V**—Support has been added in this release for the Windows Server 2019 operating system for Hyper-V based HyperFlex deployments.
- **Kubernetes CSI Plugin**—This release adds support for the HyperFlex CSI (HX-CSI) plugin based on the Kubernetes Container Storage Interface (CSI) specification. Customers can now use the HX-CSI plugin to provision and manage persistent volumes in Kubernetes version 1.13 and later. Note: Kubernetes 1.13 support for Cisco Container Platform & Openshift Container Platform is forthcoming in their respective future releases.
- **C480 ML Compute only node**—This release introduces support for C480 ML as a new compute-only node for Deep Learning/Machine Learning Workloads. Data scientists can now use the power of up to eight NVidia SXM2 V100 GPUs to accelerate deep learning workloads. VMs running deep-learning workloads will need to use PCIe pass-through for access to GPUs.
- **Higher Capacity Drives**—A new 2.4TB 10k rpm SAS HDD option for SFF Hybrid HyperFlex clusters, and a 12TB 7.2K rpm SAS HDD option for LFF Hybrid are now available. Both HyperFlex and HyperFlex Edge support the 2.4TB capacity point for maximum density in this form factor. Note that HyperFlex Edge does not support LFF-drives. HyperFlex HyperV version does not yet support the new 12TB drive option. See the HyperFlex spec sheets for a full list of configurable options.
- **New Cache and increased scale for Hyper-V** —NVMe & Optane SSDs are now supported as cache drives for Hyper-V deployments. Furthermore, scale limits have been increased to 16+16 (Converged+Compute-only) for both SFF (AF and Hybrid) & LFF (Hybrid) clusters.
- **Centralized Audit Log Export**—This release adds support for audit logging via a remote syslog server. This capability enables customers to retain audit logs from all HyperFlex nodes in a centralized remote syslog server to meet retention and compliance requirements.
- **DISA STIG Compliance**—This release adds new HX REST APIs for setting, removing, and checking status of DISA STIGs for Controller VMs, ESXi hosts and vCenter. These APIs enable customers to meet DISA security requirements by centrally and securely applying STIGs, detecting and correcting for drifts in any STIG settings.

New Supported Drives

New drives are qualified for the 4.0(2a) release. The new drives include new capacity points and new cache drive options. Several of the new drives are alternate drives to already qualified existing drives in function which are qualified in 4.0(2a). These drives are functionally compatible with the existing drives and are available as alternates in case of lack of availability of existing drives. For expansion of existing clusters or general information about interoperability of different drives, see Cisco HyperFlex Drive Compatibility.

Note NVMe Caching SSD's slot information is unavailable from HX-Connect for all AF server PIDs except for the All-NVMe server PIDs. Please refer to UCSM management console for NVMe SSD slot information.

Table 1: Supported Drives

Drive Function	Drive PID	Applicable Platforms	Version
1.6TB SED SSD Cache drive	HX-SD16TBHNK9	HX240C-M5SX	4.0(2e)
7.6TB SED SSD Capacity drive	HX-SD76TBEM2NK9	All existing HX M5 servers except All NVMe	4.0(2c)
SKUs for existing 3.8TB and 960GB Capacity drive capacity	HX-SD38TBEM2NK9, HX-SD960GBM2NK9	All existing HX M5 servers except All NVMe	4.0(2c)
Alternate system (or housekeeping) drive	HX-SD480G6I1X-EV	All existing HX M5 servers except All NVMe	4.0(2b)
Alternate system (or housekeeping) drive	HX-SD480GM1X-EV	All existing HX M5 servers except All NVMe	4.0(2b)
New 960G FIPS compliant SED SSD data drives	HX-SD960G2HTNK9	HXAF220C-M5SX, HXAF240C-M5SX	4.0(2b)
Alternate boot drive	HX-M2-960GB	All existing HX M5 servers	4.0(2a)
All NVMe 4TB Capacity drive	HX-NVME2H-I4000	All NVMe: HXAF220C-M5SN	4.0(2a)
New high density All NVMe 8TB Capacity drive	HX-NVMEHW-I8000	All NVMe: HXAF220C-M5SN	4.0(2a)
New high density All Flash 7.6TB Capacity drive Full drive scale of 23 drives on HX240 for up to 32 converged nodes requires HX release 4.0(2c).	HX-SD76T61X-EV	All Flash Configuration – namely: HXAF220C-M5SX, HXAF240C-M5SX, HXAF-E-220M5SX ESX support only.	4.0(2a)
New 3.8TB FIPS compliant SED SSD data drives	HX-SD38T2HTNK9	HXAF220C-M5SX, HXAF240C-M5SX	4.0(2a)
Alternate drive for 8TB LFF capacity	HX-HD8T7K4KAN	HX240C-M5L	4.0(2a)

Drive Function	Drive PID	Applicable Platforms	Version
800G 12G SAS Cache drive option for All Flash	HX-SD800G123X-EP	The following HX M5 servers: HXAF220C-M5SX, HXAF240C-M5SX, HXAF-E-220M5SX	4.0(2a)
All NVMe 1TB Capacity drive	HX-NVME2H-I1000	All NVMe: HXAF220C-M5SN	4.0(1b)

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 - 4.0(x) Releases

Table 2: HyperFlex Software Versions for M4/M5 Servers

HyperFlex Release	M4/M5 Recommended FI/Server Firmware *(be sure to review important notes above)	M4/M5 Qualified FI/Server Firmware
4.0(2f)	4.0(4k)[1]	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c)*, 4.1(3b), 4.1(3c), 4.1(3d), 4.1(3e)
4.0(2e)	4.0(4k)[2]	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c)*, 4.1(3b), 4.1(3c), 4.1(3d), 4.1(3e)
4.0(2d)	4.0(4k)[3]	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c)*, 4.1(3b), 4.1(3c), 4.1(3d), 4.1(3e)
4.0(2c)	4.0(4k)[4]	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2b)*, 4.1(2c)*, 4.1(3b)
4.0(2b)	4.0(4k)	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2a)*, 4.1(2c)*, 4.1(3b)
4.0(2a)	4.0(4k)	4.0(4i), 4.0(4k), 4.0(4l), 4.0(4m), 4.1(1d), 4.1(1e), 4.1(2c)*, 4.1(3b)

HyperFlex Release	M4/M5 Recommended FI/Server Firmware *(be sure to review important notes above)	M4/M5 Qualified FI/Server Firmware
4.0(1b) - Unsupported	4.0(4i)	4.0(4i)
4.0(1a) - Unsupported	4.0(4i)	-

[1] UCS release 4.1(1c) is no longer recommended due to **Field Notice: FN - 70595**.
[2] UCS release 4.1(1c) is no longer recommended due to **Field Notice: FN - 70595**.
[3] UCS release 4.1(1c) is no longer recommended due to **Field Notice: FN - 70595**.
[4] UCS release 4.1(1c) is no longer recommended due to **Field Notice: FN - 70595**.

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

Important

If your environment (or deployment) is a Fabric Interconnect 6400 connected to VIC 1455/1457 using SFP-H25G-CU3M or SFP-H25G-CU5M cables, only use UCS Release 4.0(4k), 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), or 4.1(2a) and later may cause cluster outages.

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

Use the following upgrade sequence ONLY for Fabric Interconnect 6400 connected to VIC 1455/1457 using SFP-H25G-CU3M or SFP-H25G-CU5M cables:

- Upgrade the UCS server firmware from HX Connect.
- Upgrade the UCS Infrastructure.
- Upgrade HXDP.
- Upgrade ESXi.

If you have the above hardware combination, combined upgrade of UCS server firmware is not supported. However, combined upgrade of HXDP and ESXi is supported after UCS server firmware and UCS infrastructure firmware upgrade is completed.

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 VMware ESXi, Release 4.0 for steps to align the firmware versions.
    - **M4**: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M4 or HX220c M4) 3.1(3k), 3.2(3i), or 4.0(2b) or higher 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 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 for the requirements and steps.

- **Important:** For Intersight edge servers running older than 4.0(1a) CIMC version, HUU is the suggested mechanism to update the firmware.

HyperFlex Edge and Firmware Compatibility Matrix for 4.x Deployments

Cisco HX Data Platform, Release 4.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 3: HX220c M4 / HXAF220c M4 Cluster

Component	Recommended Firmware Version - HXDP 4.x *(be sure to review important note(s) above)
Cisco Integrated Management Controller (CIMC)	4.0(2h)
Host Upgrade Utility (HUU) Download Link	4.0(2h) Download Software

Table 4: HX220c M5 / HXAF220c M5 Cluster

Component	Recommended Firmware Version - HXDP 4.x *(be sure to review important notes above)
Cisco Integrated Management Controller (CIMC)	4.1(2f)
Host Upgrade Utility (HUU) Download Link	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 for Stretched Cluster - 4.0(x) Releases

Table 5: Witness Node Versions

HyperFlex Release	Witness Node Version
4.0(2f)	1.1.1
4.0(2e)	1.1.1
4.0(2d)	1.0.10
4.0(2c)	1.0.9
4.0(2b)	1.0.8
4.0(2a)	1.0.8
4.0(1b) - Unsupported	1.0.4
4.0(1a) -Unsupported	1.0.4

Software Requirements for VMware ESXi - 4.0(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. For information on VMware ESXi recommended releases, see Cisco HyperFlex 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 vCenter administrator account with root-level privileges and the associated password.

Note For VIC1457, there is no support for ESXi 6.0.

The below table applies for all of 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 6: Software Requirements for VMware ESXi

HyperFlex Version	VMware ESXi Versions	VMware vCenter Versions
4.0(2f)	6.0 U3, 6.5 U3, and 6.7 U3	6.0 U3, 6.5 U3, 6.7 U3 and 7.0 U2 See limitations: [5]
4.0(2e)	6.0 U3, 6.5 U3, and 6.7 U3	6.0 U3, 6.5 U3, 6.7 U3, 7.0 U1c (build 17327517) through 7.0 U1d (build 17491101), and 7.0 U2 - See limitations: 5
4.0(2d)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 See limitations:[6]	6.0 U3, 6.5 U3, 6.7 U3 and 7.0 U1c (build 17327517) through 7.0 U1d (build 17491101), and 7.0 U2 - See limitations: 5
4.0(2c)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 See limitations: 5	6.0 U3, 6.5 U3, 6.7 U3 and 7.0 U1c (build 17327517) through 7.0 U1d (build 17491101), and 7.0 U2 - See limitations: 5
4.0(2b)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 See limitations: 5	6.0 U3, 6.5 U3, 6.7 U3
4.0(2a)	6.0 U3, 6.5 U3, and 6.7 U3 up to build 17098360 See limitations: 5	6.0 U3, 6.5 U3, 6.7 U3
4.0(1b) - Unsupported	6.0 U3, 6.5 U3, 6.7 U2[7]	6.0 U3, 6.5 U3, 6.7 U2
4.0(1a) - Unsupported	6.0 U3[8], 6.5 U2, 6.7 U2	6.0 U3, 6.5 U2, 6.7 U2

[5] **For HX releases 4.0(2f), 4.0(2e), 4.0(2d), and 4.0(2c) care should be taken to use the minimum vCenter 7.0 version listed in the table. Prior to 7.0 U1 vCenter versions are susceptible to a software interoperability issue (see** Field Notice: FN - 70620**). When using vCenter 7.0 U1 or 7.0 U2 with a 4.0(2a) through 4.0(2d) HXDP cluster, the following limitations apply. These limitations do not apply with 4.0(2e) and later.**

- **Fresh Installation cannot be performed with vCenter 7.0 U1 or 7.0 U2. Clusters may be deployed without vCenter initially and then subsequently registered to vCenter Server. Clusters must be registered to a vCenter server before entering production.**
- **Cluster expansion (converged & compute only) cannot be performed with vCenter 7.0 U1 or 7.0 U2. Reregister the cluster to a vCenter Server 6.x before cluster expansion is attempted.**
- **vCenter Server 7.0 U1 and above utilizes vCLS cluster VMs. These VMs must reside on a shared HX datastore to ensure smooth upgrade operations. If the vCLS VMs reside on local storage, storage vMotion them to a shared HX datastore before attempting upgrade. See** VMware documentation **for full details .**

[6] ESXi 6.7 U3 P04 (Build 17167734) or later is not supported with HXDP 4.0(2a) through 4.0(2d). See SSH Incompatibility with ESXi 6.7P04 Tech Note for further details.

[7] Use of 6.7 U2 for unsupported releases 4.0(1b) and 4.0(1a) is not recommended, see Software Advisory for Cisco HyperFlex Stretched Cluster Operations, Release 4.0(1a) for further details.

[8] Cisco HyperFlex Release 4.0(2) is the last major HyperFlex release that will support vSphere 6.0 (ESXi and vCenter) due to those versions reaching end of VMware general support on March 12, 2020.

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 General Recommendation for New and Existing Deployments.

Software Requirements for Microsoft Hyper-V - 4.0(x) Releases

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, HX240c-M5L) and All Flash (Cisco HyperFlex HXAF240c M5, HXAF220c M5) verify that Cisco UCS Manager 4.0(2b) is installed. HX 4.0 (1a) does not support Hyper-V on the All NVMe (HXAF220C-M5SN) nodes. For detailed information on installation requirements and steps, see the *Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V*.

Table 7: Supported HyperFlex Software versions for M5 Servers on Hyper-V

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

HyperFlex Release	M5 Recommended Server Firmware
4.0(1a) - Unsupported	4.0(4i)

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.

Table 8: Supported Microsoft Software versions

Microsoft Component	Version
Windows Operating System (Windows OS)	Windows Server 2016 Datacenter Core & 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. Windows Server 2019 Datacenter-Desktop Experience is supported starting from HXDP 4.0.1(a) onwards. **Note** For Windows Server 2019 Desktop Experience, the Windows 2019 ISO image should be Update Build Revision (UBR) 107 at a minimum. Windows Server 2019 Datacenter–Core is **not** supported currently. Also note that the following are currently not supported: OEM activated ISOs and Retail ISOs are **not** supported. Earlier versions of Windows Server such as Windows 2012r2 are **not** supported. Non-English versions of the ISO are **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 - 4.0(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 9: Supported Browsers

Browser	Cisco Intersight	Cisco UCS Manager	HX Data Platform Installer	HX Connect
Microsoft Internet Explorer	NA	11 or higher	11 or higher	11 or higher
Google Chrome	62 or higher	57 or higher	70 or higher	70 or higher
Mozilla Firefox	57 or higher	45 or higher	60 or higher	60 or higher
Apple Safari	10 or higher	9 or higher	NA	NA
Opera	NA	35 or higher	NA	NA

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 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 - 4.0(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 10: 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	HX240c M5 Edge Short Depth HXAF240c M5 Edge Short Depth HX220c M5 Edge HXAF220c M5 Edge HX220c M4 Edge HXAF220c M4 Edge	HXAF220c-M5SN All NVMe - HXAF220c-M5SN 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-M5SN

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 Compute-only nodes: 0-16	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 Compute-only nodes: 0-32	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

- Starting with release 4.0(2a), SCVM is no longer needed on a Compute node.
- **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 addition, AAA enforces rate limiting on Obtain Access Token API requests: in a 15 minute window, /auth can be invoked (successfully) a maximum of 5 times. A user is allowed to create a maximum of 8 unrevoked tokens. Subsequent call to /auth will automatically revoke the oldest issued token to make room for the new token. A maximum of 16 unrevoked tokens can be present in system. In order to prevent brute-force attacks, after 10 consecutive failed authentication attempts, a user account is locked for a period of 120 seconds. For more information, see Cisco HyperFlex Systems REST API Reference guide.

  HxConnect makes use of AAA Authentication REST API for login and the above rate limit applies to HxConnect also.
- Single socket stretch cluster nodes are not supported.
- Intersight Managed Mode is not currently supported for HyperFlex.

Upgrade Guidelines

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

- **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 or any other qualified releases. 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.
- **Unsupported HX Data Platform 1.7.x, 1.8.x, 2.0, 2.1x, 2.5x, and 2.6x clusters**—Users from any version prior to 2.6(1a) must step through an intermediate version before upgrading to 4.0 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).
- **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.
- **vSphere 6.7 Software Advisory**—Do not upgrade to Cisco HX Data Platform Release 4.0(1a) when running ESXi 6.7U1 EP06 (build # 11675023). Do not upgrade to 6.7U1 EP06 (build # 11675023) if running Cisco HX Data Platform Release 4.0(1a). See the Software Advisory CSCvo56350 for further details.

  The software build version posted at release will override any other local versions.
- **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 Cisco HX Data Platform Release 3.5 or later. In addition, ESXi should be upgraded as required to meet HX Data Platform compatibility requirements.

- **Minimum HXDP version for upgrade**—HX Data Platform clusters running 2.6(1a) or later may upgrade directly to 4.0 using the HX Connect UI.
- **Cluster Readiness**—Ensure that the cluster is properly bootstrapped and the updated plug-in is loaded before proceeding. Manual cluster bootstrap is required for upgrade from a pre-3.5 release.
- **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 4.0. 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 4.0.

  Manual bootstrap is not supported on Intersight clusters.
- **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.

  If the current cluster version is at 3.5(1a) or above, you do not need to use the `stcli` command. Direct upgrade to 4.0 is possible.
- **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.
- **ESXi and HXDP Compatibility—**Ensure your cluster is running a compatible version of ESXi based on the running the HX Data Platform version (see the section Software Requirements for VMware ESXi). ESXi compatibility is determined by the major version and update release of ESXi. It is generally best to upgrade HXDP and ESXi together if combining the upgrade operations into a single optimized reboot. When running a split upgrade, first upgrade the HX Data Platform, then proceed to upgrade ESXi.
- 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.
- **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.
- **If Upgrading to vSphere 6.5:**

Note

- 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 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 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 Server Firmware Upgrades**—Upgrade server firmware 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.
  - Users are encouraged to upgrade to 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.
- **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 clusters. 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.
- **Maintenance Window**—If upgrading both HX Data Platform and UCS firmware, you can select either a combined or split upgrade 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.
- **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. All drives must be SED drives, mixing SED and non-SED is not supported.

- **Enabling External Host Access**—With Cisco HX Data Platform Release 4.0(1a), port 445 on the management network is blocked for enhanced security. Note that prior to 4.0, port 445 port was open enabling external host access. If you are upgrading to 4.0(1a) from a prior release, and would like to continue external host access, you can use a utility to open select hosts. For more information about enabling external host access, see the "Configuring HyperFlex Share to SCVMM" section in the Installation Guide for Microsoft Hyper-V.

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.

Security Fixes

The following security issues are resolved:

Release	Defect ID	CVE	Description
4.0(2f)	CSCvy14839	NA	A vulnerability in the logging subsystem of the Cisco HyperFlex System could allow an authenticated, remote attacker to view sensitive information in a system log file which should be restricted.

Release	Defect ID	CVE	Description
4.0(2e)	CSCvx36028	CVE-2021-1499	A vulnerability in the web-based management interface of Cisco HyperFlex HX Data Platform could allow an unauthenticated, remote attacker to upload files to an affected device. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvx52126	CVE-2021-1499	A vulnerability in the web-based management interface of Cisco HyperFlex HX Data Platform could allow an unauthenticated, remote attacker to upload files to an affected device. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvx37435	CVE-2021-1498	A vulnerability in the web-based management interface of Cisco HyperFlex HX Installer Virtual Machine could allow an unauthenticated, remote attacker to perform a command injection attack against an affected device. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvx36019	CVE-2021-1497	A vulnerability in the web-based management interface of Cisco HyperFlex HX Installer Virtual Machine could allow an unauthenticated, remote attacker to perform a command injection attack against an affected device. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvx36014	CVE-2021-1497	A vulnerability in the web-based management interface of Cisco HyperFlex HX Installer Virtual Machine could allow an unauthenticated, remote attacker to perform a command injection attack against an affected device. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2e)	CSCvv75781	CVE-2017-18269, CVE-2018-11236, CVE-2018-11237, CVE-2018-19591, CVE-2018-6485, CVE-2019-17514, CVE-2019-18348, CVE-2019-18874, CVE-2019-20907, CVE-2019-9169, CVE-2019-9674, CVE-2020-14344, CVE-2020-14422, CVE-2020-14556 ,CVE-2020-14578, CVE-2020-14579, CVE-2020-14583, CVE-2020-14593, CVE-2020-14621, CVE-2020-1751, CVE-2020-2754, CVE-2020-2755, CVE-2020-2756, CVE-2020-2757, CVE-2020-2767, CVE-2020-2773, CVE-2020-2781, CVE-2020-2800, CVE-2020-2803, CVE-2020-2805, CVE-2020-2816, CVE-2020-2830,CVE-2020-8492	Multiple vulnerabilities from multiple TPS components. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvv15388	CVE-2020-14422	Lib/ipaddress.py in Python through 3.8.3 improperly computes hash values in the IPv4Interface and IPv6Interface classes, which might allow a remote attacker to cause a denial of service if an application is affected by the performance of a dictionary containing IPv4Interface or IPv6Interface objects, and this attacker can cause many dictionary entries to be created. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvw50465	NA	Includes updates to address vulnerabilities in multiple third party software packages. For more information, see the related Cisco Security Advisory.
4.0(2e)	CSCvu95813	CVE-2020-12049, CVE-2019-8457, CVE-2020-13790, CVE-2020-12762, CVE-2018-8740, CVE-2019-19603, CVE-2019-19645, CVE-2020-11655, CVE-2020-13434, CVE-2020-13435, CVE-2020-13630, CVE-2020-13631, CVE-2020-13632, CVE-2019-17023, CVE-2020-12399, CVE-2019-3689, CVE-2019-1547, CVE-2019-1549, CVE-2019-1551, CVE-2019-1563	`napi_get_value_string_*()` allows various kinds of memory corruption in node. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2c)	CSCvu33080	CVE-2019-9512, CVE-2019-9514, CVE-2019-9515, CVE-2020-10108, CVE-2020-10109, CVE-2020-8597, CVE-2018-12327, CVE-2017-6350, CVE-2017-6349, CVE-2017-5953, CVE-2019-20079, CVE-2017-11109, CVE-2018-20786, CVE-2017-1110, CVE-2020-10531, CVE-2020-7595, CVE-2019-19956, CVE-2019-19923, CVE-2019-20218, CVE-2019-19925, CVE-2020-9327, CVE-2019-19959, CVE-2019-8457, CVE-2019-13753, CVE-2019-19924, CVE-2019-13734, CVE-2019-13752, CVE-2019-19926, CVE-2019-13751, CVE-2019-13750, CVE-2019-19880, CVE-2019-5188, CVE-2019-5094, CVE-2015-9383, CVE-2019-15796, CVE-2019-15795	Includes updates to address vulnerabilities in multiple third party software packages. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq63138	CVE-2019-13132, CVE-2019-9924	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. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq71240	CVE-2019-11719, CVE-2019-11727, CVE-2019-11729	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. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvr06339	CVE-2019-1125	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. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2a)	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	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. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp65019	CVE-2017-13168, CVE-2017-18174, CVE-2017-18216, CVE-2018-10876, CVE-2018-10877, CVE-2018-10878, CVE-2018-10879, CVE-2018-10880, CVE-2018-10881, CVE-2018-10882, CVE-2018-10902, CVE-2018-10938, CVE-2018-12233, CVE-2018-12896, CVE-2018-13053, CVE-2018-13094, CVE-2018-13096, CVE-2018-13405, CVE-2018-13406, CVE-2018-14609, CVE-2018-14617, CVE-2018-14633, CVE-2018-14734, CVE-2018-15572, CVE-2018-15594, CVE-2018-16276, CVE-2018-16658, CVE-2018-17182, CVE-2018-17972, CVE-2018-18021, CVE-2018-18690, CVE-2018-18710, CVE-2018-6554, CVE-2018-6555, CVE-2018-9363	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. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvo98516	NA	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 client-side browser attacks.
4.0(2a)	CSCvj95584	NA	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.

Release	Defect ID	CVE	Description
4.0(2a)	CSCvp24343	CVE-2018-15380	The vulnerability is due to insufficient CSRF protections for the web UI on an affected device. An attacker could exploit this vulnerability by persuading a user of the interface to follow a malicious link. A successful exploit could allow the attacker to perform arbitrary actions with the privilege level of the affected user. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq19949	CVE-2019-11834, CVE-2019-11835	Vulnerabilities with a version of cJSON identified by CVE-2019-11834 and CVE-2019-11835. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvr54398	CVE-2018-12207, CVE-2019-11135	HX ESXi image patches for MCEPSC and TAA vulnerabilities with VMware ESXi identified by CVE-2018-12207 and CVE-20190-11135. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq19546	CVE-2019-11477, CVE-2019-11478, CVE-2019-11479	CP networking vulnerabilities affecting the Linux kernel identified by CVE-2019-11477, CVE-2019-11478, CVE-2019-11479. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvr54399	CVE-2018-12207, CVE-2019-11135	Qualification of Microsoft security patches for vulnerabilities with Microsoft Hyper-V hypervisor identified with CVE-2018-12207, CVE-2019-11135. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp76463	CVE-2016-10708, CVE-2018-20685, CVE-2019-6109, CVE-2019-6111	OpenSSH vulnerabilities identified by CVE-2019-6111. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2a)	CSCvp66679	CVE-2016-2105, CVE-2016-2106, CVE-2016-2107, CVE-2016-2109, CVE-2016-2176, CVE-2016-2177, CVE-2016-2178, CVE-2016-2179, CVE-2016-2180, CVE-2016-2181, CVE-2016-2182, CVE-2016-2183, CVE-2016-6302, CVE-2016-6303, CVE-2016-6304, CVE-2016-6305, CVE-2016-6306, CVE-2016-6307, CVE-2016-6308, CVE-2016-7055, CVE-2016-8610, CVE-2017-3731, CVE-2017-3732, CVE-2017-3735, CVE-2017-3736, CVE-2017-3737, CVE-2017-3738, CVE-2018-0495, CVE-2018-0732, CVE-2018-0734, CVE-2018-0735, CVE-2018-0737, CVE-2018-0739, CVE-2018-12384, CVE-2018-12404, CVE-2018-5407, CVE-2019-1559	Multiple vulnerabilities associated with OpenSSL and LibNSS. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp66555	CVE-2016-10087	libpng vulnerability identified by CVE-2016-10087. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp34586	CVE-2014-9092, CVE-2016-3616, CVE-2017-15232, CVE-2018-11212, CVE-2018-11213, CVE-2018-11214, CVE-2018-1152, CVE-2018-13785	Vulnerabilities associated with libjpg and libpng. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp31207	CVE-2018-16428, CVE-2018-16429	Vulnerabilities associated with Glib identified by CVE-2018-16428 and CVE-2018-16429. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvr36903	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	Multiple vulnerabilities associated with curl, expat, python 2.7, python 3.5, 3.6 and 3.7, freetype and giflib. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq92032	CVE-2019-14379, CVE-2019-12384, CVE-2019-14439	Multiple third party software vulnerabilities. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2a)	CSCvq43250	CVE-2018-16062, CVE-2018-16402, CVE-2018-16403, CVE-2018-18310, CVE-2018-18520, CVE-2018-18521, CVE-2019-7149, CVE-2019-7150, CVE-2019-7665	Vulnerabilities associated with elfutils. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq43230	CVE-2017-5953, CVE-2019-12735	Vulnerabilities associated with vim. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq43213	CVE-2016-6153, CVE-2017-10989, CVE-2017-13685, CVE-2017-2518, CVE-2017-2519, CVE-2017-2520, CVE-2018-20346, CVE-2018-20505, CVE-2018-20506, CVE-2019-8457, CVE-2019-9936, CVE-2019-9937	Vulnerabilities associated with sqlite3. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq43209	CVE-2018-20843	Vulnerabilities associated with glib2.0. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq43205	CVE-2018-20843	Vulnerabilities associated with expat. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq43194	CVE-2016-3189, CVE-2019-12900	Vulnerabilities associated with bzip2. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq10694	CVE-2018-12115, CVE-2018-0734, CVE-2018-5407, CVE-2018-12120, CVE-2018-12121, CVE-2018-12122, CVE-2018-12123, CVE-2018-12116, CVE-2019-5737, CVE-2019-5739	Multiple Vulnerabilities associated with NodeJS. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq10388	CVE-2019-10906, CVE-2016-10745	Vulnerabilities associated with Pallets Jinja str.format_map and identified by CVE-2019-10906 and CVE-2016-10745. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2a)	CSCvq09178	CVE-2014-8501, CVE-2014-9939, CVE-2015-9262, CVE-2016-10087, CVE-2016-2226, CVE-2016-4487, CVE-2016-4488, CVE-2016-4489, CVE-2016-4490, CVE-2016-4491, CVE-2016-4492, CVE-2016-4493, CVE-2016-6131, CVE-2018-10963, CVE-2018-13785, CVE-2018-17100, CVE-2018-17101, CVE-2018-18557, CVE-2018-18661, CVE-2018-7456, CVE-2018-8905	Multiple Security Vulnerabilities. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq07568	CVE-2019-9893	Vulnerabilities associated with libseccomp. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvq06755	CVE-2017-12447	Vulnerabilities associated with gdk-pixbuf. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp93817	CVE-2018-6594	Vulnerabilities associated with python-crypto. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp86721	CVE-2018-20483, CVE-2019-5953	Vulnerabilities associated with wget. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp66748	CVE-2018-6594	Multiple vulnerabilities associated with Python Crypto. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp66734	CVE-2017-17512	Vulnerabilities associated with sensible-utils. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp66689	CVE-2018-1000030	Multiple vulnerabilities associated with Python. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(2a)	CSCvp66672	CVE-2016-10713, CVE-2018-1000156, CVE-2018-6951	Multiple patch vulnerabilities identified by CVE-2016-10713. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp66664	CVE-2016-10165, CVE-2018-16435	Vulnerabilities associated with Little CMS identified by CVE-2016-10165. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp64746	CVE-2011-5325, CVE-2016-7076, CVE-2017-1000368, CVE-2019-11068	Vulnerabilities associated with Tenable Scan identified by CVE-2011-5325, CVE-2019-11068, CVE-2016-7076, CVE-2017-1000368. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp34792	CVE-2016-9318, CVE-2017-16932, CVE-2017-18258, CVE-2018-14404, CVE-2018-14567	Vulnerabilities associated with libxml2 2.9.4 and earlier, as used in XMLSec 1.2.23 and earlier and other products and identified by CVE-2016-9318. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvp29266	CVE-2018-10916	Vulnerabilities associated with LFTP Remote File Names Unauthorized Access Vulnerability and identified by CVE-2018-10916. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvo34097	CVE-2018-7750	Vulnerabilities associated with CryptographyDeprecationWarning: signer and verifier have been deprecated. For more information, see the related Cisco Security Advisory.
4.0(2a)	CSCvr39793	CVE-2019-16056	Multiple Vulnerabilities in python 2.7, 3.5. For more information, see the related Cisco Security Advisory.

Release	Defect ID	CVE	Description
4.0(1b), 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.
4.0(1b), 3.5(2g)	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.

Release	Defect ID	CVE	Description
4.0(1b)	CSCvo88997	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, 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	The vulnerabilities associated with the JVM 1.8U121 memory leak found during VC alarms (concurrent 40 calls using REST API).
4.0(1b)	CSCvm58031	NA	Tomcat and Nginx logs are not being collected in the support bundle generated through HX Connect in Release 3.5.
4.0(1a)	CSCvn35119	CVE-2018-18584 CVE-2018-18585	The vulnerabilities associated with the **libmspack** software package version included in Cisco HX Data Platform.

Release	Defect ID	CVE	Description
4.0(1a)	CSCvn82282	CVE-2018-14719 CVE-2018-14720 CVE-2018-1000873 CVE-2018-14721 CVE-2018-19360 CVE-2018-19362 CVE-2018-19361 CVE-2018-14718	The vulnerabilities associated with FasterXML Jackson-Databind Time Value Field Denial of Service.
4.0(1a)	CSCvo05054	CVE-2013-3587	The vulnerabilities associated with the **OpenSSL Protocol** software package version included in Cisco HX Data Platform.
4.0(1a)	CSCvo27818	CVE-2018-16487 CVE-2018-19361	The vulnerabilities associated with Third-Party Software Denial of Service.
3.5(2a)	CSCvm53149	CVE-2018-1092 CVE-2018-7492 CVE-2018-8087 CVE-2018-1068 CVE-2018-8781	The vulnerabilities associated with **Linux kernel for Ubuntu 17.10**.

Resolved Caveats in Release 4.0(2f)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvy07797	HX upgrade to release 4.0(2e) fails with `Checking if pre-upgrade has been completed on the nodes. Next on empty iterator.`	4.0(2e)	4.0(2f)
CSCvy03362	HyperFlex release 4.0(2b) to 4.0(2e) vib upgrade fails due to special characters in hxuser password.	4.0(2e)	4.0(2f)

Resolved Caveats in Release 4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvx01406	Frequent APD hits.	4.0(2d)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvv88204	An ESXi OpenSSH Interoperability Issue exists with HXDP 3.5(2) and 4.0(2a-d). Starting with 6.7 P04 and later, the following functional areas of HX will be impacted including: • Fresh cluster creation (may fail with Algorithm negotiation fail) • Cluster expansion (may fail with Algorithm negotiation fail) • Cluster reregistaration (stcli cluster reregister may fail with "Algorithm negotiation fail") • System information page in HX Connect • Upgrades may fail with "Failed to Establish SSH Connection to host" or "Errors found during upgrade"	3.5(2e)	4.0(2e)
CSCvr31746	This defect tracks the condition where bank or rank level ADDDC/VLS Sparing copy causes a temporary stall of HX Controller VM on the impacted node to trigger one or more of the following failure symptoms: **1.** If the impacted node had the Zookeeper Leader process running, it can potentially terminate multiple Zookeeper sessions leading to storfs restarts on multiple nodes and eventually an APD. **2.** The stalling may cause Zookeeper client running on the impacted node to timeout and the session could expire leading to storfs process on that node to restart. This will result in a temporary unhealthy event. **3.** The stalling may cause storfs process to observe a high IO latency on one or more drives with active IO requests pending on those drives. This could lead to drives being marked as blacklisted and the cluster would become unhealthy until the drives are auto-repaired.	3.5(2e)	4.0(2e)
CSCvv19737	On some Hyperflex Edge clusters, when registering with Smart Licensing they will consume the "Cisco SP HyperFlex HX Data Platform SW v2.0" license instead of the "HyperFlex Data Platform Edge Edition Subscription"	4.0(2c)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvv59521	You may see the following error message during a HyperFlex install or expand using the local OVA installer: `Installing Software Packages on Storage Controller VMfailed in Task: 'Initializing Storage Controller VM for Installation' with Error: 'The conditional check '(not packagesinstalled.stat.exists) or (not existingBuildManifest.stat.exists) or (not targetBuildManifest.stat.exists) or (targetBuildManifest.stat.md5 != existingBuildManifest.stat.md5)' failed. The error was: error while evaluating conditional ((not packagesinstalled.stat.exists) or (not existingBuildManifest.stat.exists) or (not targetBuildManifest.stat.exists) or (targetBuildManifest.stat.md5 != existingBuildManifest.stat.md5)): 'dict object' has no attribute 'md5`	4.0(2c)	4.0(2e)
CSCvt35006	HyperFlex datastores may report high IO latency during CRM Master failover. If current CRM Master node reboots, the new CRM Master initialization can take more time and results in IO latency.	3.5(2h)	4.0(2e)
CSCvv21905	If RO user is not created at the install time (this happens if the installation done using Intersight), then when the user goes to the encryption page, an **Authenticate UCSM** button is shown. When user clicks that button, it fails with invalid CSRF token.	4.0(2a)	4.0(2e)
CSCvv09832	Replication not working after Cluster Expansion between new node and remote site. Network Tests fails.	4.0(2b)	4.0(2e)
CSCvt16158	Please see below KB and blog from VMware. 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, 6.7 and 7.0 versions. If the build number is older, a two step upgrade will be required to get to the latest 6.5, 6.7 and 7.0 ESXi builds. See the following VMware KB: https://kb.vmware.com/s/article/76555 See the following VMware blog: https://blogs.vmware.com/vsphere/2020/01/vsphere-signing-certificate-expiry-what-you-need-to-know.html	4.0(2a)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvq15478	VM network performance degraded and/or Poor HyperFlex storage performance. Significant and incrementing rx_no_buf errors seen on HyperFlex Storage Data VNIC's which correlate to the above statement.	3.5(2a)	4.0(2e)
CSCvq59999	Call-home notifications generated when a node is placed into MM during known maintenance activity and/or during an upgrade.	3.5(2f)	4.0(2e)
CSCvp79511	HX Upgrade to release 3.0(1i) was allowed while vCenter and ESXi were both on version 6.0u2 while Release Notes and version check require version 6.0u3.	4.0(1a)	4.0(2e)
CSCvv13773	During upgrade, the disks get marked as **IGNORED**.	4.0(2b)	4.0(2e)
CSCvx09397	Following a full cluster power outage, in rare situations the cluster may not recover on its own.	4.0(2a)	4.0(2e)
CSCvp09978	`stcli` cluster information shows Smart call home enabled, even though disabled.	3.5(2b)	4.0(2e)
CSCvw84976	When Replication is configured on the Cluster, replication network tests(inter and intra cluster) fails due to missing replIpSettings in the nodes inventory. Datastores cannot be mapped from the UI. replIpSettings values can be listed by stcli node list | grep -i -C 7 replIpSettings	4.0(2c)	4.0(2e)
CSCvv03450	HX Installer Password textbox validation logic flags incorrect when password entered is too complex.	3.5(2h)	4.0(2e)

CSCvu69826	`stcli` cluster reregister command fails with the following error when an HX host has a vVol datastore mounted:
	`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)	4.0(2e)

CSCvi35116	Monitoring process for any VC alarms such as APD event, host disconnection, vm memory usages will not be started.	3.0(1a)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvv09614	After the SCVM reboots, during the bootup process you will see errors such as: Waiting for rootfs to become rw: /usr/share/springpath/storfs-appliance/monitor-bootdev.sh: line 49: cannot create temp file for here-document: Read-only file system /usr/share/springpath/storfs-appliance/monitor-bootdev.sh: line 20: /var/old-log/stv-bootdev.log: Read-only file system The SCVM will reboot after completing the boot process (within a few mins).	3.5(2a)	4.0(2e)
CSCvt03880	Panic hits when previous repl CIP is retained changing replication configuration.	3.5(2e)	4.0(2e)
CSCvv16609	Storfs PANIC is observed on a node which has a failed disk that is still discoverable to HX Controller VM but IO requests to the disk do not get any response within 60 seconds.	3.5(2e)	4.0(2e)
CSCvt23930	You might see all servers online in the HX Connect dashboard, however one or more servers can show Hypervisor state as 'Offline' in the Nodes page of HX Connect. Furthermore, the 'hxcli node list' command will report the Hypervisor offline as well and requests to enter HX Maintenance Mode will fail. /var/log/springpath/stMgr.log from the cluster lead SCVM will show WBEM failing to connect: "[WBEMException: CIM_ERR_FAILED (Unable to connect)]"	4.0(2a)	4.0(2e)
CSCvv23981	Replication network cleanup triggers panic.	4.0(2b)	4.0(2e)
CSCvv09225	After adding a DNS entry to a HyperFlex system, it is not added to the output of `stcli cluster info`.	3.5(2g)	4.0(2e)
CSCvw39220	We will see high CPU usage in stctlvm causing storfs performance to be impacted.	4.0(2c)	4.0(2e)
CSCvx01406	Frequent APD hits.	4.0(2d)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvx17718	HyperFlex cluster expansion will fail validation step of **vCenter and ESXi uniform version check**, Cluster 'XXXX' not found in datacenter. Please create the cluster on the targeted datacenter in vCenter" Installer expects HX cluster object to be in root of datacenter.	4.0(2d)	4.0(2e)
CSCvg81805	DS shows unmounted on the vCenter for one HX host, but ESX CLI shows mounted. Also this one host is not in the list of hosts on VC web client for the Datastore.	3.5(2d)	4.0(2e)
CSCvq77158	HX DS could become inaccessible due to clock skew when you shutdown and re-start 1 node. Running the command: **ntpd -gq on controller** returns with **no ntp servers found**.	3.0(1a)	4.0(2e)
CSCvt43958	As OOM killer kills the main storfs process on a given controller, the resiliency state of cluster turns to WARNING, but eventually (and automatically) is restored to HEALTHY state. Under extreme condition, if 2 or more nodes fault simultaneously, the cluster may shutdown, and may have to be restored manually using CLI. There is no data loss, but workload VMs may suffer storage outage (APD - All paths down) for the duration of cluster downtime.	3.5(2g)	4.0(2e)
CSCvs53555	You may see this error message after a failed upgrade or other task such as attempting to enter a node into HX maintenance mode: getClusterLocalizableMessage(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(2e)	4.0(2e)
CSCvq94466	Node expansion fails due to timeout.	3.5(2d)	4.0(2e)
CSCvs08218	Virtual machine disk consolidation required error after native snapshot removal. Virtual machines flagged with "Virtual Machine disks consolidation is needed".	3.5(2b)	4.0(2e)
CSCvt75978	4-node HXDP version 3.5(2g) loop in snapshot tree causes host to crash.	3.5(2g)	4.0(2e)
CSCvt46022	Cluster stuck in an unhealthy state due to vNode resync being stuck.	3.5(2a)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvt49323	For 10-node HX cluster with LAZ configured, HX Connect shows that the zones are unevenly distributed.	3.5(2g)	4.0(2e)
CSCvu52699	User can observe following symptoms after replacing Hyperflex server systemboard. • Intersight UI - Node is not listed in HyperFlex cluster detailed inventory view page. • Change of License tier for HyperFlex cluster will fail and reverts back to old value (example - If changing from Base to Essentials, it will fail and remain at Base). On the intersight UI, user can notice that the hyperflex.Node (server) object has old server serial number and PhysicalServer object has null value. UCSM, UCSM inventory in Intersight has updated new server details and issue is only with HX inventory in Intersight.	3.5(2h)	4.0(2e)
CSCvr52949	All the ESXi nodes lost connectivity to datastores. They were mounted but unavailable.	3.5(2a)	4.0(2e)
CSCvw69697	Cluster Upgrade validation fails with; Upgrade validations failed. Hosts <node-ip> are not in connected state in the vCenter server.Please make sure all hosts are in connected state.	3.5(2b)	4.0(2e)
CSCvs27184	Upgrade from ESXi versions 6.0 or 6.5 to 6.5 or 6.7 manually, get the error "'Could not find a trusted signer".	3.5(2h)	4.0(2e)

CSCvq38092	When a single node is offline in cluster, 'stcli cluster storage-summary' shows two nodes unavailable.
	`root@SCVM:~# stcli cluster storage-summary`
	`...`
	`    messages:`
	`        ----------------------------------------`
	`          Storage cluster is unhealthy.`
	`        ----------------------------------------`
	`          Storage nodes 192.168.1.4, 192.168.1.1`
	`are unavailable.`
		3.5(2a)	4.0(2e)

CSCvq65830	VM corrupted after migrating from one host to another.	3.5(2a)	4.0(2e)
CSCvr83056	HyperFlex Datastore NFS Queue Depth shows as 256, which can lead to performance (including latency) issues.	3.5(2e)	4.0(2e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs95434	When the Witness VM is deployed through the ovf template and the password has " (double quotes) in it, the witness vm fails to retain its network config after a reboot or two. Modified /etc/network/interfaces to include the network config and also copied the same over to /etc/network/eth0.interface After a reboot, the config in /etc/network/eth0.interface goes back to dhcp.	3.5(2h)	4.0(2e)
CSCvu62527	`service_status.sh` shows scvmclient status as Running though scvmclient on ESXi is stopped.	3.5(2a)	4.0(2e)

Resolved Caveats in Release 4.0(2d)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvw39220	We will see high CPU usage in stctlvm causing storfs performance to be impacted. For workaround details, see Cisco HyperFlex Software Advisory for HX Release 4.0(2c).	4.0(2c)	4.0(2d)

Resolved Caveats in Release 4.0(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvv05705	HyperFlex Installer VM deployment fails during the task "Installing Software Packages on Storage Controller VM" with the error: "non-zero return code' Retrying the workflow will not fix the issue and continues to halt with the same error.	4.0(2c)	4.0(2c)
CSCvs86562	On a cluster where VMware EAM manages the controller, VMs upgrade fails with exit maintenance mode step failing. You will see more than 3 attempts to power on CVM fail with the error "No host is compatible with the virtual machine", and CVM gets powered on more than 30 seconds after exit maintenance mode.	4.0(2a)	4.0(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvt73521	When doing a deploy or expansion of Hyperflex from a HX installer, and the setup has more than the default 8 NICs created in UCSM. The HX **deploy** stage may hang at the following step: Deploying Storage Controller VM on ESXi Host - Configuring Network (Port Groups) for ESXi and Storage Controller VM. ESXi host may not be reachable on its assigned management address.	3.5(2d)	4.0(2c)
CSCvn89717	During an upgrade from HX release 3.5.1a to HX release 3.5.2a on a 3 node cluster with M5 hardware, the USB interface was deleted on one of the controller VMs. This resulted in the SED drives locking.	3.5(2a)	4.0(2c)
CSCvu73740	Case generated via Smart Call Home which attaches a SCH CLI Output that only contains the cluster_info.	3.5(2h)	4.0(2c)
CSCvu36042	The Storfs process on Springpath Controller VM will panic in inconsistent Network condition (such as disconnects, varying bandwidth or latency) when replication is forced to reconnect to the destination cluster.	4.0(2b)	
CSCvt61403	For HX release 3.5(2g), 5 nodes hit APD with a bad disk in the cluster.	3.5(2g)	4.0(2c)
CSCvr37846	A node in the cluster stopped processing I/Os from clients and other nodes. This caused an All Paths Down timeout in ESX NFS hosts.	3.5(2a)	4.0(2c)
CSCvt22494	An Error occurs while expanding the cluster through classic installer: - The time zone name '<NAME>' was not found on the computer. Applicable to Hyper V environment only.	4.0(2a)	4.0(2c)
CSCvu58631	HX release 3.5(2h) SED stretch cluster expansion failed as the converged node is not listed in server selection page.	3.5(2h)	4.0(2c)
CSCvu58785	Performing an API call for a token refresh fails on HyperFlex. The clusters encounters a failure response.	4.0(2a)	4.0(2c)
CSCvt21961	NTP FQDN resolution fails due to search domain being appended to NTP FQDN.	3.5(2h)	4.0(2c)
CSCvt51128	When eth1 is shut on ZK follower (or leader) on 2N Edge, the cluster reports healthy on it.	4.0(2a)	4.0(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvu08247	Repeated High Memory usage alarms may be seen in HyperFlex Connect after deployment or upgrade on to HX release 3.5(2g)/4.0(2a) where enhanced monitoring was first implemented.	3.5(2g)	4.0(2c)
CSCvu27654	Sometimes, the Witness VM fills up the volume containing Zookeeper logs and transactions. This may lead to Zookeeper service misbehaving within the Witness VM and could also potentially result in an unresponsive Zookeeper service. Also, filling up the folder prevents Zookeeper from logging any further.	3.5(2h)	4.0(2c)
CSCvu50471	vCenter URL is missing in ZK randomly during cluster creation. Day2 operations like upgrade is impacted when vCenter is empty.	3.5(2g)	4.0(2c)
CSCvt43958	As OOM killer kills the main storfs process on a given controller, the resiliency state of cluster turns to WARNING, but eventually (& automatically) is restored to HEALTHY state. Under extreme conditions, if two or more nodes fault simultaneously, the cluster may shutdown, and may have to be restored manually using CLI. There is no data loss, but workload VMs may suffer storage outage (APD - All paths down) for the duration of cluster downtime.	3.5(2g)	4.0(2c)
CSCvt41200	When using Mgmt IP Address change, we may hit - Unsupported KEX algorithm "diffie-hellman-group1-sha1".	4.0(1b)	4.0(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvu07906	Upgrade fails in post upgrade task with error: post upgrade task failed: Creating dynamic Certificate The stMgr.log file has the following exception: ERROR c.s.s.c.http.HttpDownStreamService - Unable to put the content to the downstream, url: /securityservice/v1/certificate?option=dynamic Error Response: com.twitter.util.TimeoutException: 35.seconds DEBUG c.s.s.c.h.HttpDownStreamService$HttpDownStreamServiceUtil$ - putCall operation timed out) com.twitter.util.TimeoutException: 35.seconds at com.twitter.util.Future.$anonfun$within$1(Future.scala:1638) ~[util-core_2.12-17.10.0.jar:17.10.0] at com.twitter.util.Future$$anon$2.apply$mcV$sp(Future.scala:1686) ~[util-core_2.12-17.10.0.jar:17.10.0] at com.twitter.util.Monitor.apply(Monitor.scala:46) ~[util-core_2.12-17.10.0.jar:17.10.0] at com.twitter.util.Monitor.apply$(Monitor.scala:41) ~[util-core_2.12-17.10.0.jar:17.10.0] at com.twitter.util.Future$MonitoredPromise.apply(Future.scala:175) ~[util-core_2.12-17.10.0.jar:17.1	4.0(2a)	4.0(2c)
CSCvs08218	Virtual machine disk consolidation required error after native snapshot removal. Virtual machines flagged with Virtual Machine disks consolidation is needed.	3.5(2b)	4.0(2c)
CSCvt27376	BackUp Vendor REST API session timeouts on upgrade to HX 4.0(2a), due to rate limit of Authentication requests.	4.0(2a)	4.0(2c)
CSCvt90065	Running the custom workflow that configures IPs in the background (after encountering errors earlier which left the system in a partially configured state), the HX installer may raise an Error - the existing IP addresses already provisioned on the servers as Duplicate IPs.	4.0(2b)	4.0(2c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs97460	When the cross data-center replication link bandwidth varies from value set by the user in the HX UI pane for replication during pairing, Springpath controllers do not auto-tune the rate of transmission. This leads to missing heartbeats and failure to replicate the data across the cluster. Replication failures are seen at the UI layer. In addition, in low bandwidth and high latency networks, large number of failures occur due to the non-adaptive nature of the replication rate. This enhancement supports varying link bandwidth and a bandwidth drop in link of up to 50% of configured replication bandwidths in HX by automatically controlling the rate of transmission.	4.0(2a)	4.0(2c)

Resolved Caveats in Release 4.0(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvt36374	*/var/stv* folder in Controller VM may become full.	3.5(2g)	4.0(2b)
CSCvs70967	**stcli services dns remove** should remove the DNS server info from the interface files.	2.5(1d)	4.0(2b)
CSCvt10522	HX 4.0(1b), new deployment from Intersight production cloud. User receives following error while enabling CSI (Kubernetes integration). **Failure occurred during volume_access_enable, Error was: Nonzero exit code 1**	4.0(1b)	4.0(2b)
CSCvt13947	Receive the following alert/event in HX Connect: HX Controller VM {HOSTNAME} one or more configured DNS servers not responding.	4.0(2a)	4.0(2b)
CSCvt14914	API, UI, CLI show few Drive slots Empty after upgrade to 4.0.2a. The drives are seen in the sysmtool, lsscsi and stcli command outputs. The cluster is healthy with no errors.	4.0(2a)	4.0(2b)
CSCvs69154	After successfully changing the DNS server on the HX controller, we still can see the original DNS entry that was added during the deployment.	3.5(2d)	4.0(2b)
CSCvs30080	HyperFlex Connect and VCenter show APD alarms for HX and non HX datastores. This implies there is an issue with HyperFlex when there may not be.	3.5(2e)	4.0(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs74286	An issue where all the disks in the node get locked after rebooting the node. We successfully unlocked using **sed-client.sh -U** command, but wanted to test with another reboot, and drives locked again.	4.0(2a)	4.0(2b)
CSCvt13929	When running "stcli license..." commands on HyperFlex, the following error is seen: root@SpringpathController:/tmp# stcli license show all Show smart licensing failed: Smart Agent is not ready, please wait a minute and try again	4.0(2a)	4.0(2b)
CSCvt20144	After upgrading to HX 4.0(2a) the first Robo/Edge node that gets rebooted, then ignores persistent drives PID:HX-HD24TB10K4KN Cluster will be in degraded state and have reduced capacity.	4.0(2a)	4.0(2b)
CSCvs21562	Zookeeper fails to start while Exhibitor is running, however echo srvr returns nothing.	3.5(2b)	4.0(2b)
CSCvs91787	When performing an HyperFlex upgrade a validation warning may occur due to the host not having Enterprise Plus or Enterprise hypervisor licensing. Upgrade Validation Warning: ESXi host esx1.lab.test should be configured with VMware Enterprise license for upgrade to continue.	4.0(2a)	4.0(2b)
CSCvs69317	Cluster expansion fails at Config Installer stage when the root and admin password for the storage controller (SCVM) are different.	3.5(2g)	4.0(2b)
CSCvt06983	Panic while upgrading ESXi.	3.5(2g)	4.0(2b)
CSCvs54285	A cluster node running may hang in the Linux kernel. This is classified as an oops and a deviation from the expected behavior.	4.0(2a)	4.0(2b)
CSCvs69007	Rebalance failing on 10+10 node HX 3.5(2g) stretch cluster.	3.5(2g)	4.0(2b)
CSCvr54687	The cluster becomes inaccessible to the IOVisor.	3.5(2d)	4.0(2b)
CSCvt61297	Panic on the storage controller.	4.0(2a)	4.0(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvt63306	The size of support bundle is very large when collected with storfs-support command.	4.0(2a)	4.0(2b)
CSCvt72807	storfs crash associated with **FileSystemUsageWarningEvent or FileSystemUsageAlertEvent due to high usage of /var/stv above 80%**.	4.0(2a)	4.0(2b)

Resolved Caveats in Release 4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
ESXi, Installation, Upgrade, Expansion, Management			
CSCvo39912	In normal scenarios, HX upgrade gets pending ACK after updating service profile. In this case, UCS only upgrade was stalled because after updating service profile HX upgrade was not getting pending ACK, and it was always waiting state.	3.5(2d)	4.0(2a)
CSCvh80044	Hyper-V: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)
CSCvq39471	When using motherBoardReplace-1.2 to clean ZK from old stNode/pNodes resulted in unmounted datastore and making the size of the Hyperflex Datastores 0 resulting in all VMs in the cluster going offline.	3.5(1a)	4.0(2a)
CSCvm77294	Upgrading cluster and getting error: Failed upgrade validations : Checking vCenter configuration. Reason: Upgrade validations failed. DRS Fault: Insufficient resources to satisfy configured failover	2.6(1e)	4.0(2a)
CSCvo70650	The cluster expand fails on a node with DR replication configured. When a 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 , 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.	3.5(2a)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvo91624	Customer reported that, server didn't complete the firmware upgrade automatically one by one. The user manually needs to put the host into maintenance mode and they manually acknowledged the pending requests to finish the UCS firmware upgrade. Maintenance policy is set to default (User-ack) as per the design.	3.5(1a)	4.0(2a)
CSCvo93017	When the cluster is in "Failed" state and stcli node remove is attempted, the output shows successful though the node failed to get removed from the cluster.	3.5(1i)	4.0(2a)
CSCvp31021	HyperFlex cluster upgrade may fail during validation with a DRS Validation failed error.	3.5(1a)	4.0(2a)
CSCvp58318	HX cluster expansion will fail with the following error message: MAC address pool configuration failure 150[ErrorDescription]: bad address block range definition collision.	3.5(2b)	4.0(2a)
CSCvp36220	If you perform an "stcli node add" and it exceeds 15 mins, you get a message indicating "Failed to add nodes", and "time out".	3.5(2a)	4.0(2a)
CSCvq34873	Memory usage by carbon cache.	3.5(2b)	4.0(2a)
CSCvr03240	Upgrading ESXi cluster fails with error "Node maintenance mode failed".	3.5(2g)	4.0(2a)
CSCvr88978	storfs process does not automatically start on Exit HX Maintenance Mode or other tasks which restart Storage Controller.	3.5(2e)	4.0(2a)
CSCvp10707	Post install script fails with message: Failed to execute ipmitool on HX node.	3.5(2a)	4.0(2a)
CSCvp46539	HyperFlex expansion workflow doesn't pull VLAN name correctly.	4.0(1a)	4.0(2a)

CSCvq45087	During HX cluster deployment - Validate Cluster Creation phase - HX Installer might fail with the error:
	`*** from`
	`\var\log\springpath\stDeploy.log`
	`***`
		3.5(2a)	4.0(2a)

CSCvq91380	HX Installer fails to login to the SOL for configuring ESX hosts.	3.5(2b)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvr44222	If DRS is enabled with configuration parameters other then fully automated i.e. manual or partially automated than this will change to fully automated during cluster expansion.	3.5(2d)	4.0(2a)
CSCvp82175	One of the validation task name has a typo. The task name is "Validating node for mixed of SED and non-SED disks". It should be "Validating node for mix of SED and non-SED disks".	3.5(1a)	4.0(2a)
CSCvo12359	This combination of using smaller NVMe drives is not supported in expand operation. For example, if the drives are only 375GB and they cannot be added to an existing cluster with larger caching SSD.	3.0(1i)	4.0(2a)
CSCvp66679	HyperFlex includes a version of OpenSSL that is affected by the vulnerabilities identified by the following Common Vulnerability and Exposures: CVE-2018-0495	3.5(1a)	4.0(2a)
CSCvo00511	SDK invoke exception when taking a native snapshot on a VM with GPU pci passthrough attached.	3.0(1e)	4.0(2a)
CSCvo62867	Node replacement script fails due to EAM error.	3.0(1i)	4.0(2a)
CSCvo79760	While pairing a cluster (HX release >= 3.5) and cluster (HX release < 3.5), remote replication network test fails on the cluster (HX release < 3.5).	3.5(1a)	4.0(2a)
CSCvo87061	Fixed the issue in the latest support-workflow bundle. This should not be hit anymore.	3.0(1b)	4.0(2a)
CSCvo87080	On a 3 node cluster the MbReplace script fails as the script looks for "HEALTHY" cluster status.	3.5(2a)	4.0(2a)
CSCvp12359	MbReplace script (tar) hangs when running on 3.5.2b or 4.0.	3.5(2a)	4.0(2a)
CSCvp63958	HX replication cleanup failing with error NA 3.5(2a) - INFO:DR state is not clean.	3.5(2a)	4.0(2a)
CSCvp66277	The check_stig_parameters API incorrectly shows the compliance state of the cluster as not STIG compliant.	4.0(1a)	4.0(2a)
CSCvq06952	The snapshot creation on CBT enabled VM fails with error "Failed in vmreparent vmkfstools clone1".	3.5(2c)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvo60587	HyperFlex GUI, when setting up remote DR-replication partner/peer, should show remote replication partner TCP port reachability test results. This will provide quick upfront connectivity results, that frequently takes customer a while to manually check and confirm. If reachability tests fail, an info window and HX alert should trigger that list the destination peer IP and the specific port(s) that could not be reached.	3.5(2a)	4.0(2a)
CSCvr67130	If you have more than one IP pool configured for a replication network and local replication network tests are failing, you might be running into this issue.	3.5(2b)	4.0(2a)
CSCvp05204	Skip task option in MBreplace and Software redeploy scripts. The scripts should have the option to skip tasks as user inputs and not by modifying the json file.	3.5(2a)	4.0(2a)
CSCvq34357	Error noticed on SCVM Console "print_req_error: I/O error, dev fd0, sector 0" on HyperV. It doesn't come up on SSH session.	4.0(1a)	4.0(2a)
CSCvk36222	The VM Network switch is already created as part of the install process. Ideally the ip address assignment should be done during the initial install process as well. No post install steps should be required, alternatively it could be done as part of a post install script.	3.0(1d)	4.0(2a)
CSCvg53223	storfs service impacted during HX upgrade from HX 2.1(1) to 2.5(1c) HX node which is not being upgraded storfs service stopped causing outage.	2.5(1c)	4.0(2a)
CSCvp37536	HX Stretch Cluster Witness VM reverts to DHCP on reboot.	3.5(1b)	4.0(2a)
CSCvp46578	HyperFlex Stretch Cluster Witness does not get programmed with an NTP server.	3.5(2b)	4.0(2a)
CSCvr18528	HyperFlex cluster does not heal after maintenance on node/server.	3.5(2b)	4.0(2a)
CSCvr97089	High percentage of packet loss on the eth1 data interface 9K MTU.	4.0(2a)	4.0(2a)
CSCvr16760	The HyperFlex cluster healing state is stuck at 87% and is unable to progress.	3.5(2f)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvq22898	In some circumstances, two replications can occur for the same VM. From the GUI, the replication status for this VM will show stuck in an "In Progress" state, but there are no jobs stuck in replication layer. Subsequent Replications should succeed.	4.0(2a)	4.0(2a)
CSCvp13990	After an unexpected power outage the hyperflex cluster is unable to come up correctly. All processes are running, time is synced, and all nodes can vmkping each other as required.	3.5(2a)	4.0(2a)
CSCvp33657	When specifying IP address, instead of node ID, **stcli node maintenanceMode** commands fail.	4.0(1a)	4.0(2a)
CSCvr31573	2 node edge cluster, if experiencing Intersight/arbitrator connectivity issues, should NOT allow HX Maintenance mode to initiate, as this will bring the cluster down.	4.0(1a)	4.0(2a)
CSCvr89066	Old files in /var/support/ZKTxnlog are not purged with the daily zklog-cleanup cron job.	3.5(2c)	4.0(2a)
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 Maintenance Mode.	3.5(2a)	4.0(2a)

CSCvr01645	Cluster might go offline and stop serving IO requests:
	`root@cvm:~# stcli cluster storage-summary --detail`
	`Get cluster storage summary failed:`
	`java.net.ConnectException: Connection refused: /192.168.142.100:10207`
		3.5(2c)	4.0(2a)

CSCvq80340	A stretch cluster deployment fails with a message 'Formatting nodes Node in Use' on the Hyperflex installer and crmZoneType value is seen as 1 in storfs.cfg or /opt/springpath/config/stretch.tunes.	3.5(2d)	4.0(2a)
CSCvq63888	During upgrade we upload the upgrade package downloaded from Cisco.com. Upgrade packages are of file type .tgz hxconnect accepts storfs-packages-4.0.1a-33028.tar file, a different type of compression, and allows upgrade to start and fail, when it cannot decompress the file for bootstrap and scvmclient upgrade on ESXi.	3.0(1c)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvr52098	While attempting to install HyperFlex via Intersight the following error is seen: Failed in Task: 'Add hosts to vCenter Cluster' with Error: 'Try adding hosts manually to vCenter and retry. Failed to get address:Failed to host by name:lookup Hostname.company.com on 0.0.0.0:53: read udp 127.0.0.1:45430->127.0.0.1:53: read: connection refused '.	4.0(2a)	4.0(2a)
CSCvo92952	CoreAPI call may time out for cluster management API while doing cluster create validation.	3.5(2a)	4.0(2a)
CSCvq95460	Validation fails with "Mixed mode expansion check" due to Enhanced vMotion Compatibility Incompatibility when the ESXi version is different.	3.5(2d)	4.0(2a)
CSCvr66309	This failure is seen in installer when custom workflow is used "hypervisor configuration" + "deploy" (clean disk partition=no) at "Installing software packages on storage controller VM".	3.0(1i)	4.0(2a)
CSCvq93831	The HX Installer apparently replaced a comma by semicolon on the VLAN ID range while exporting the configuration to a JSON file.	3.5(2d)	4.0(2a)
CSCvp22693	When using unsupported browsers, HX Connect users will not be able to login. These users see an error indicating unauthorized user.	4.0(1a)	4.0(2a)
CSCvp24343	A vulnerability in the web-based management interface of Cisco HyperFlex Software could allow an unauthenticated, remote attacker to conduct a cross-site request forgery (CSRF) attack on an affected system. The vulnerability is due to insufficient CSRF protections for the web UI on an affected device. An attacker could exploit this vulnerability by persuading a user of the interface to follow a malicious link. A successful exploit could allow the attacker to perform arbitrary actions with the privilege level of the affected user. Cisco has released software updates that address this vulnerability. There are no workarounds that address this vulnerability. This advisory is available at the following link: https://tools.cisco.com/security/center/content/CiscoSecurityAdvisory/cisco-sa-20190807-hyperflex-csrf	4.0(1a)	4.0(2a)

Defect ID	Symptom	First Release Affected	Resolved 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.	3.5(2d)	4.0(2a)
CSCvr43786	Compute Node expand operation fails due to incorrect cluster name.	4.0(1b)	4.0(2a)
CSCvp41241	RF-2 cluster shutdown during data resynchronization; after a node failure (non storfs); followed by multiple disk read failures; and a HardBlacklist of a disk.	2.6(1e)	4.0(2a)
CSCvo13143	HyperFlex Edge nodes do not properly set the ESXi hostname during deployment.	4.0(1a)	4.0(2a)

Resolved Caveats in Release 4.0(1b)

Defect ID	Symptom	First Release Affected	Resolved in Release
ESXi, Installation, Upgrade, Expansion, Management			

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs28167		2.6(1e)	4.0(1b)

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; in order 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 4.0(1b) patched OVA file for Cisco HyperFlex Data Platform Installer for VMware ESXi:
	`Cisco-HX-Data-Platform-Installer-v4.0.1b-33133p1-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.		

CSCvp64140	While running the HyperFlex installer with Windows Server Hyper-V, cluster creation process fails with the error: "`Failure occurred during Cluster Creation process: Unable to post the content to the downstream`". This symptom is encountered while deploying HyperFlex cluster with HyperFlex nodes configured/ordered with Cisco VIC 1457 MLOM (PID : HX-MLOM-C25Q-04) and Windows Server Datacenter or Core with Hyper-V version 2016 or 2019.	4.0(1a)	4.0(1b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvo69067	Adding Micron 5200 drive to a cluster fails to increase cluster capacity.	3.5(2b)	4.0(1b)

Resolved Caveats in Release 4.0(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
ESXi, Installation, Upgrade, Expansion, Management			

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvs28167		2.6(1e)	4.0(1a)

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; in order 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 4.0(1a) patched OVA file for Cisco HyperFlex Data Platform Installer for VMware ESXi:
	`Cisco-HX-Data-Platform-Installer-v4.0.1a-33028p1-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.		

CSCvk17250	Cluster instability when disks of different sector size placed in HX node.	3.0(1d)	4.0(1a)

CSCvo36198	When logged in, using a local HX user account instead of a Virtual Center account, an error message appears intermittently indicating "`Virtual Center unreachable`" or "`Resource information cannot be updated`", when VC is reachable.	3.5(1a)	4.0(1a)
			3.5(2c)

CSCvk38003	HXDP does not work with EMC RecoverPoint - needs to support VMware API (FSS-Readdir).	3.0(1d)	4.0(1a)
			3.5(2b)

Defect ID	Symptom	First Release Affected	Resolved in Release

CSCvm90352	Zookeeper (Exhibitor) process on HyperFlex Storage Controller VM's may fail to respond to requests when `/var/zookeeper` has low or no free space.	2.5(1c)
		2.5(1d)
			4.0(1a)
			3.5(2a)
			3.0(1i)

CSCvn02151	Use Asynchronous consolidation for HX Snapshots.	2.6(1c)	4.0(1a)
CSCvo90713	HX Quiesced Snapshot for Backup Vendors.	3.5(2b)	4.0(1a)
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)	4.0(1a)
CSCvn51562 CSCvo48463	Cisco HX Data Platform plugin fails to load in Windows vCenter Web Client 6.7 U1. The issue is not seen with VMware VCSA.	3.5(2a)	4.0(1a)
CSCvn73127	Kernel migration fails when a local datastore is searched for in the ESXi host.	3.0(1d)	3.5(2b)
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(2a)
CSCvi59119	Duplicating an existing datastore name that differs only in letter case might result in unknown behavior.	3.0(1a)	4.0(1a)
CSCvh80044	Hyper-V: 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)	4.0(1a)
CSCvc62266 CSCvm16157	After an offline upgrade, due to a VMware EAM issue, sometimes all the controller VMs do not restart. The `stcli start cluster` command returns an error: "Node not available".	2.0(1a)	4.0(1a)
Hyper-V			
CSCvn28721	Cluster expansion may fail with an error code `500 - operation timed out`.	3.5(2a)	4.0(1a)

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)
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)

Open Caveats in Release 4.0(2f)

There are no open caveats in this release.

Open Caveats in Release 4.0(2e)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs41324	Enabling K8/iscsi stack in HX cluster.	**1.** After upgrade enable iscsi as per steps in CSCvs41324 **2.** Delete the existing and new pod deployments only and recreate them.	4.0(2e)
CSCvq38279	When replicated DC was used, then during install time Windows failover cluster was not created successfully.	Clean up fail over cluster and recreate the fail over cluster. No need to touch HX storage cluster.	4.0(2e)
CSCvs62854	Upgrade failed at `Enter Platform Maintenance Mode` step on a HX 4.0(2a) cluster if the cluster expanded with new converged nodes.	**1.** Update VMUUID in /etc/springpath/secure/hxinstall_inventory.json to upper case for new nodes. **2.** Restart hxSvcMgr Repeat above procedure on all nodes.	4.0(2e)

Open Caveats in Release 4.0(2d)

There are no open caveats in this release.

Open Caveats in Release 4.0(2c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvu52699	The following symptoms can be observed after replacing HyperFlex server system board: 1. Intersight UI - Node is not listed in Hyperflex cluster detailed inventory view page. 2. Change of License tier for HyperFlex cluster fails and reverts back to old value (example - if changing from Base to Essentials, it fails and remains at Base) . From API, the HyperFlex node (server) object has old server serial number and physical server object has null value. UCSM inventory in Intersight has updated new server details and issue only with HX inventory in Intersight. HXDP Zookeeper is not updated with correct (new) Serial Number.	NA	3.5(2h)
CSCvq94466	Node expansion fails due to timeout.	NA	3.5(2d)
CSCvt35006	HyperFlex datastores may report high IO latency during CRM Master failover. If current CRM Master node reboots, the new CRM Master initialization can take more time and results in IO latency.	None, need to upgrade to fixed version.	3.5(2g)
CSCvu85439	HyperFlex Cluster may remain online, but datastores are not available and VM's become inaccessible.	Contact Cisco TAC. If issue is confirmed, TAC can stop storfs service on affected node to restore service.	3.5(2d)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvt10849		Manually copied hxupgrade_bundle.tgz file to the affected node /tmp file from the other node and restart stMgr from the CMIP node.	4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
	Error initiating upgrade: hxmanager.log 2020-02-18-07:38:29.937 [opID=865cc7297abb7] Got response 200 in 84.18382ms for GET https://localhost/rest/job?jobtype=check_cluster_upgrade_validations 2020-02-18-07:38:29.938 Completed 200 OK in 84.845132ms 2020-02-18-07:38:39.962 Started GET /hx/api/clusters/1/upgrade/clusterValidations 2020-02-18-07:38:39.963 [opID=6a028ec40c797] Request URL GET https://localhost/rest/job?jobtype=check_cluster_upgrade_validations 2020-02-18-07:38:40.025 [opID=6a028ec40c797] Got response 200 in 62.57042ms for GET https://localhost/rest/job?jobtype=check_cluster_upgrade_validations 2020-02-18-07:38:40.026 Completed 200 OK in 63.237634ms 2020-02-18-07:38:42.061 Started POST /hx/api/clusters/1/upgrade 2020-02-18-07:38:42.061 [opID=7575f932578db] Request URL POST https://localhost/rest/upgrade/cluster 2020-02-18-07:38:42.107 [opID=7575f932578db] Got response 500 in 46.104743ms for POST https://localhost/rest/upgrade/cluster 2020-02-18-07:38:42.107 [opID=7575f932578db] Error code 500 | <nil> | /rest/upgrade/cluster 2020-02-18-07:38:42.107 [error] 500|{"message":"Upgrade in progress","messageId":806}|/rest/upgrade/cluster 2020-02-18-07:38:42.107 Completed 500 Internal Server Error in 46.356043ms Cleaned up previous upgrade process using "stcli cluster upgrade --components		

Defect ID	Symptom	Workaround	Defect Found in Release
	hxdp" and then reinitiated. 2. Upgrade was stuck for 24hours with no progress. stMgr.log:- 2020-02-18-06:31:37.969 [] [opId=000e5553e678cf9f, operationId=000e5553e678cf9f(000e5553e678cf9f<000e5553e678cf9f] [pool-2-thread-21] DEBUG c.s.s.s.StMgrImpl$StMgrAPIWrapper$ - checkforUpgrade failed) {} java.io.FileNotFoundException: /tmp/hxupgrade_bundle.tgz (No such file or directory)		
CSCvr95936	HyperFlex TLS/SSL Server supports the use of Static Key Ciphers.	NA	4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvu92384		NA	4.0(2b)

Defect ID	Symptom	Workaround	Defect Found in Release
	root@HyperFlex-Installer-4.0.2b:~# post_install Select post_install workflow: 1. New/Existing Cluster 2. Expanded Cluster (for non-edge clusters) 3. Generate Certificate **Note** Workflow No.3 is mandatory to have unique SSL certificate in the cluster. By Generating this certificate, it will replace your current certificate. If you're performing cluster expansion, then this option is not required. Selection: 2 Expanded Cluster workflow selected Logging in to controller xx.xx.xx.xx HX CVM admin password: Getting ESX hosts from HX cluster... WARNING:root:Unable to fetch the deploymentMode from stcli. Will retry with fallback mechanism. vCenter URL: xx.xx.xx.xx Enter vCenter username (user@domain): administrator@vsphre.local vCenter Password: Found datacenter xxx Found cluster xx-xx post_install to be run for the following hosts: hx4 Enter ESX root password: Enter vSphere license key? (y/n) n Enable HA/DRS on cluster? (y/n) n		

Defect ID	Symptom	Workaround	Defect Found in Release
	Disable SSH warning? (y/n) y Add vmotion interfaces? (y/n) y Existing cluster configuration for reference --- Netmask Vlan-Id IP-Address --- 255.255.255.0 vmotion xx.xx.xx.231 255.255.255.0 vmotion xx.xx.xx.232 255.255.255.0 vmotion xx.xx.xx.233 --- Enter Expanded node configuration- Netmask for vMotion: 255.255.255.0 VLAN ID: (0-4096) 3093 Expanded node configuration should match with existing cluster configuration. Kindly check the above reference information and retry. Netmask for vMotion: 255.255.255.0 VLAN ID: (0-4096) 3093 Expanded node configuration should match with existing cluster configuration. Kindly check the above reference information and retry.		
CSCvt22567	HXDP 3.5(2b) - Default VMware tools location change caused storfs restart.	Disable polling for update of VMware tools in VMX.	3.5(2g)
CSCvu29049	8-node cluster with SED enabled - upgrading from HX 3.5(2b) to 3.5(2h) - Auto-bootstrap failed, so manual bootstrap tried but still it was giving error.	Make sure the syslog is configured properly on storage controller VMs. If there are syslog configuration customizations done please revert them and try the upgrade again.	3.5(2h)
CSCvu07899	During post upgrade task, vCenter reregistration fails with **unknown host** message if the specified host name format is https://<ip>.	Re-registration workflow can be retried again by specifying valid host name format using command **stcli cluster reregister**.	4.0(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvu93214	Recover page in HX Connect shows error, but the recovery of VM operation is successful in the backend and displayed in the Activity tab.	Please check the Activity page for the status of the submitted operation.	4.0(2c)
CSCvu83071	If this is happening, you might see the deploy phase of the expand operation fail out with network connectivity related error messages. You will also see the VLAN ID set to 0 on the storage data port groups of the new server. If this defect has been triggered and the SCVM deploy phase is successful, the cluster expand phase will fail to mount HX datastores to the new host.	Watch the host port group configurations in vCenter while the deploy is running. If you add the VLAN IDs to the port groups fast enough each time the installer changes them to null then the expand will work without failing. If you are using this custom workflow on an HX release earlier than 4.0(2a), you can prevent the bug by only using one workflow option at a time: **1.** "I know what I'm doing" > Configure Hypervisor. **2.** Start the installer over > "I know what I'm doing" > Deploy HX Software. **3.** Start the installer over > "I know what I'm doing" > Expand Cluster.	3.5(2h)

Open Caveats in Release 4.0(2b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvt55712	When EAM service is not running, vCenter registration is failing.	Start/restart EAM service in vCenter.	4.0(2a)
CSCvt45344	HyperFlex Stretch Cluster saw poor application performance due to write latency, cluster remained unhealthy and rebalance was stuck. This defect is currently being used for investigative purposes.	Contact Cisco TAC if you believe you are seeing similar conditions.	3.5(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvt36374	/var/stv folder in Controller VM may become full.	Remove log files from /var/log/springpath folder when the log files are not getting rotated. Stop the relevant service, remove large log files from /var/log/springpath folder and restart the service.	3.5(2g)
CSCvt89709	In a Cluster configured with DR network, StMgr may not get initialized and stuck in a deadlock while enabling the IPTable rules (can be verified from the stMgr.log).	Restarting the stMgr will initialize it again.	4.0(2b)
CSCvv21905	**UCSM Read-Only** user missing error in the encryption page on HX Connect UI. Later while authenticating UCS-M 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)

Open Caveats in Release 4.0(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvx49418	HX converged servers deployed using HX Installer prior to HX release 4.0(2a), will have their storage data traffic going to Fabric Interconnect A as active and Fabric Interconnect B as standby. Expanded nodes using HX Installer 4.0(2a) or later will have their storage data traffic going to FI B as active and FI A as standby. Since HX storage traffic is not going to the same FI, the traffic needs to go to the upstream switch. This may lead to storage performance impact.	In a rolling fashion, update the storage vSwitch NIC teaming so that storage B NIC as active and storage A as standby for nodes deployed prior to HX 4.0(2a). Make sure the cluster is healthy before you begin this process and that the cluster returns to healthy state before moving on to the next node to perform the NIC teaming change. **Powershell to update the active / standby of the team -** **Set-NetLbfoTeamMember -Name storage-data-b -AdministrativeMode Active** **Set-NetLbfoTeamMember -Name storage-data-a -AdministrativeMode Standby**	4.0(2a)
CSCvs75553	When the user recovers a VM protected in a group, it along with other VMs in the group, move to "recovering" state in standalone mode. The selected VM should move to "recovered" state.	Click "recover" for the selected VM again under the standalone mode. This will move the VM to recovered state.	4.0(2a)
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)
CSCvs74286	All the disks in the node got locked after rebooting the node. We successfully unlocked using 'sed-client.sh -U' command, but wanted to test with another reboot, and drives locked again.	NA	4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs93245	After enable encryption ,HX connect encryption status showing Caution, but all of the disks enable encryption successful. -"Self Encrypting Drives Service ",is running on all of the nodes. -There is no error message of encryption. - USB0 interface is up. -All SED disks showing "supported": 1, "enabled": 1, "locked": 0,	To recover the VM, copy over the data disks and attach them to the new VM.	3.5(2g)
CSCvs54285	A cluster node running HX release 4.0(1b), may hang in the Linux kernel. This is classified as an oops and a deviation from the expected behavior.	Enable kernel.panic_on_oops as a persistent configuration. This will cause the node to panic and reboot immediately.	4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs41324	Enabling K8/iscsi stack in Hx cluster.		4.0(2a)

Defect ID	Symptom	Workaround	Defect Found in Release
		Steps to enable K8 on HX cluster. Perform following steps on ALL the controller VMs: **1.** Update "/etc/init/scvmclient.conf" to enable tune "iscsiEnable" to "true". **# sed -ie "s/iscsiEnable=false/iscsiEnable=true/" /etc/init/scvmclient.conf** **2.** Run following initctl command to reload configuration: **# initctl reload-configuration** **3.** Restart scvmclient process. **# stop scvmclient; start scvmclient** **4.** Run following initctl command to mount SYSTEM DS. **# initctl emit --no-wait system-datastore-created** **5.** Verify that "iscsiEnable" tune is set to "true" using the following command. **# ps -eaf | grep scvmclient | grep -v grep** **root 6241 1 1 Dec06 ? 01:40:07 /opt/springpath/storfs-core/scvmclient -T logEnabled=false -T logSyslogEnabled=true -T logEchoToScreen=false -T statLoggingToFile=false -T statLoggingToSyslog=true -T logDir=/var/log/springpath -T nfsBackendServerList=10.107.48.100 -T iscsiEnable=true -T iscsiUseAsync=true -T**	

Defect ID	Symptom	Workaround	Defect Found in Release
		iscsiConfigFilePath=/nfs/SYSTEM/istgt.conf# ps -eaf | grep scvmclient | grep -v grep	
CSCvr83056	HyperFlex Datastore NFS Queue Depth shows as 256, which can lead to performance (including latency) issues.	Following procedure can be used to check NFS Queue Depth and increase if needed: **root@HXESXI1] vsish -e get /vmkModules/nfsclient/mnt/[DATASTORE]/properties | grep -i maxqdepth maxQDepth:256 <- Low value [root@HXESXI1] vsish -e set /vmkModules/nfsclient/mnt/[DATASTORE]/properties maxQDepth 1024 [root@HXESXI1] vsish -e get /vmkModules/nfsclient/mnt/[DATASTORE]/properties | grep -i maxqdepth maxQDepth:1024 <- Optimal value** This requires an ESXi host reboot to take effect. Please place the node in Hyperflex Maintenance Mode and gracefully reboot the node for the changes to be applied.	3.5(2e)
CSCvq38092	When a single node is offline in cluster, 'stcli cluster storage-summary' shows two nodes unavailable root@SCVM:~# stcli cluster storage-summary ... messages: -- Storage cluster is unhealthy. -- Storage nodes 192.168.1.4, 192.168.1.1 are unavailable.	This is cosmetic. The error will go away once the cluster returns to healthy status.	4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs69317	Cluster expansion fails at Config Installer stage when the root and admin password for the storage controller (SCVM) are different.	Modify both the root and admin password for the SCVMs to be the same. **stcli security password set -u admin**	3.5(2g)
CSCvs69154	After successfully changing the DNS server on the HX controller, we still can see the original DNS entry that was added during the deployment.	NA	3.5(2d)
CSCvs53555	You may see this error message after a failed upgrade or other task such as attempting to enter a node into HX maintenance mode:	Customers are strongly encouraged to work with TAC in order to identify and remediate this issue: **1.** Change 'getcluster' from 2m to 8m in /opt/springpath/storfs-mgmt/stMgr-1.0/conf/application.conf on all nodes. **2.** Restart stMgr on all nodes.	3.5(2e)
CSCvt13947	Receive the following alert/event in HX Connect: HX Controller VM {HOSTNAME} one or more configured DNS servers not responding.	Run the following command on each storage controller VM as root: **grep -i "monitor_dns_servers" /opt/springpath/hx-diag-tools/watchdog_config.json && sed -ie 's/"monitor_dns_servers": true/"monitor_dns_servers": false/' /opt/springpath/hx-diag-tools/watchdog_config.json && grep -i "monitor_dns_servers" /opt/springpath/hx-diag-tools/watchdog_config.json && restart watchdog**	4.0(2a)
CSCvs21562	Zookeeper fails to start while Exhibitor is running, however echo srvr returns nothing.	Delete any empty (Size 0) log files under /var/log/zookeeper/version-2 OR /var/log/zookeeper/standalone/version-2.	3.5(2b)

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 authenticating UCS-M 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)

Open Caveats in Release 4.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
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. **Note** This issue arises when M.2 drive running firmware is D0MU049 and it is upgraded to firmware D0MH072.	**Workaround: Pre upgrade (Can be performed remotely)** • Upgrade the server into fixed version with HUU. • No need to change server to standalone mode. • For more information see, https://tinyurl.com/vqnytww. **Workaround: Post failure (Requires onsite support)** **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 fail over cluster, and then recreate the fail over cluster. There is no need to touch the HX storage cluster.	3.5(2e)
CSCvq54992	Upgrade from release HX 3.5(2a) to 3.5(2x), 4.X does not upgrade scvmclient. Recovery point may not work.	Manually install / upgrade scvmclient VIB and make sure that it matches with HXDP version.	3.5(2b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvj22992	VM shows up on multiple nodes.	To recover the VM, copy over the data disks and attach them to the new VM.	3.0(1b)
CSCvm96629	Cluster experienced an APD due to incorrect network configuration.	Ensure all network configurations, including jumbo packet based end-to-end connectivity through top-of-rack switch is validated. May need to restart the controller VM for any changes to take effect, if the changes were made after the system is up and running.	3.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release

CSCvp23718	A cluster with 8TB or 12TB disk drives, may experience I/O stalls for several minutes after another node fails in the cluster.	After the installation (for new deployments) or after an upgrade to 4.0.1a (for existing deployments), perform the following steps on all Controller VMs:
		1. Edit the following tune files (on all Controller VMs):
		`/opt/springpath/config/lff.tunes`
		`/opt/springpath/config/vsi_1.6tb.tunes`
		2. Set
		`cleanerEnableSegSummaryCleaning`
		to "false"
		3. After editing the above tune files:
		a. SSH in to all the Controller VMs;
		b. Run " storfstool -- -Z"
		c. Run the below command and check the values. Tune value should be true
		# cat /tmp/stprocfs/system/tunables/cleanerEnableSegSummaryCleaning
		cleanerEnableSegSummaryCleaning=true
		d. Type below commands and applied the tunes changes dynamically
		# echo false > /tmp/stprocfs/system/tunables/cleanerEnableSegSummaryCleaning
		e. Verify that tune values are changed. Run below command. Value should be false.
		# cat /tmp/stprocfs/system/tunables/cleanerEnableSegSummaryCleaning
		cleanerEnableSegSummaryCleaning=false
		f. umount /tmp/stprocfs
			4.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvq53058	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.	NA	3.5(2a)
CSCvp36364	This product includes Third-party Software that is affected by the vulnerabilities identified by the following Common Vulnerability and Exposures (CVE) IDs: CVE-2016-0762, CVE-2016-6797, CVE-2016-6816, CVE-2016-8735, CVE-2017-5647, CVE-2017-12615, CVE-2017-12616, CVE-2017-12617, CVE-2017-7674, CVE-2018-1304, CVE-2018-8014, CVE-2018-1336, CVE-2018-8034, CVE-2018-11784, CVE-2019-0232	NA	4.0(1a)
CSCvq11456	stcli cluster info command needs to provide UCSM VIP address. Currently it shows ucsm-host.com	NA	3.5(2d)
CSCvq32530	HyperFlex upgrade validation failed because cluster has extraneous stNodes in internal database.	No workarounds currently. The stale entries do not affect the day to day operations of HyperFlex.	3.5(2a)
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)
CSCvq39471	When using motherBoardReplace-1.2 to clean ZK from old stNode/pNodes resulted in unmounted datastore and making the size of the Hyperflex Datastores 0 resulting in all VMs in the cluster going offline.	Run the stcli datastore mount command, to remount the datastore.	3.5(2b)
CSCvq66245	Currently HyperFlex installs on Hyper-V do not contain the "stcli security whitelist" commandset.	NA	4.0(1a)

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 authenticating UCS-M 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)

Open Caveats in Release 4.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
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`		4.0(1a)

<table>
<tr><th>Defect ID</th><th>Symptom</th><th>Workaround</th><th>Defect Found in Release</th></tr>
<tr><td></td><td></td><td>1. Set Jumbo frame in Windows/Hypervisor for incoming node:
Get-NetAdapter storage-data-a | Get-NetAdapterAdvancedProperty -RegistryKeyword *JumboPacket | Set-NetAdapterAdvancedProperty -RegistryValue 9014
Get-NetAdapter storage-data-b | Get-NetAdapterAdvancedProperty -RegistryKeyword *JumboPacket | Set-NetAdapterAdvancedProperty -RegistryValue 9014
2. Verify all interfaces are set correctly to support Jumbo frames:
Get-NetAdapter *storage* | Get-NetAdapterAdvancedProperty | ? RegistryKeyword -match "jumbo" | ft -auto
You should receive the following message:
`Name DisplayName DisplayValue RegistryKeyword RegistryValue`
`---- ----------- ------------ --------------- -------------`
`vswitch-hx-storage-data Jumbo Packet 9014 Bytes *JumboPacket {9014}`
`storage-data-b Jumbo Packet Bytes 9014 *JumboPacket {9014}`
`storage-data-a Jumbo Packet Bytes 9014 *JumboPacket {9014}`
3. Reboot only the incoming node or the node being expanded.
4. Retry Cluster Expansion in</td><td></td></tr>
</table>

Defect ID	Symptom	Workaround	Defect Found in Release
		installer (from the point where it errored out previously).	
CSCvp12241	Two node HyperFlex Edge cluster may not failback successfully and return to healthy. This can occur if connection to Intersight is highly impaired (e.g. transaction latencies exceeding 100ms).	Confirm that both nodes are up and running and you have given it some time to heal (a few hours) before attempting the workaround. If not healed and failed-back, run the following command on both controller VMs (preferably simultaneously on both nodes) to restart: **restart hxRoboController**.	4.0(1a)
CSCvp20102	Datastore create/delete fails when VC is not available for ESXi HX cluster (regular or stretch or 2N robo).	Ensure VC is available before performing DS operations.	4.0(1a)

CSCvk23212	Emergency cluster shuts down in HX 3.0(1b) after exiting host from HX Maintenance Mode and storfs panic on other node.	Avoid removal and re-insertion of drives that are regarded as 'good'. If you have performed testing that has removed and reinserted 'good' drives, contact Cisco TAC for further instructions.	3.0(1b)
			3.5(1a)

CSCvm55176	During Hyper-V installation, if you choose to perform constrained delegation later, sometimes it takes an excessive amount of time to reflect on HX Connect UI.	Wait at least 30 minutes for the AD policy to take effect. If the issue is not resolved, reboot 1 host at a time using maintenance mode.	3.5(1a)
CSCvq04252	HX 3.5(2b) installer fails on the hypervisor configuration step with no visible error. UCSM configuration completes, but Hypervisor Configuration seems to not start.	Use the **I know what I'm doing workflow** for hypervisor configuration, deploy HX software, and cluster creation.	3.5(2a)

CSCvp17427	Stcli cluster storage-summary takes a long time to return on 16+ node cluster when one node reboots.	Please wait for cluster to heal and re-run command.	3.5(2b)
			3.5(1a)

CSCvm53679	HX Hyper-V installation fails and `HXBootstrap.log` contains the following message:
	`Unable to find a default server with Active Directory Web Services running`
		This error indicates that Windows failed to find a domain controller. Please add a specific IP of Domain Controller in the Advanced input of HXInstaller.	3.5(1a)
			3.0(1e)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvp97422	After the network partition heals, the datastore on one of the nodes remain unavailable for some time.	NA	4.0(1a)
CSCvp21417	Deploy of EMC RecoverPoint fails with error: "`Failed finding repository device in the vRPA view`".	If upgrading to 3.5(2b) or 4.0I(1a), enabling the RecoverPoint feature may require you to perform a rolling, node by node HX Maintenance Mode in the cluster. Ensure that the cluster is healthy and can tolerate one node failure (for 3 or 4 node clusters) and 2 node failures (for 5 or greater than 5 node clusters).	3.5(2b)

CSCvo89507	In the event of adding unsupported Micron 5200 drives to an HX cluster, and then upgrading HX to a release that supports them, the drives could get locked if the cluster has remote security enabled (during certain cases like continuous reboot of controller VMs).	Remove the Micron 5200 drives from the system. Upgrade to Release 3.5(2b), then, follow disk-add expansion workflow.	4.0(1a)
			3.5(2b)

CSCvo83276	VM powers off during backup VM snapshot.	Retake the snapshot.	3.5(1a)

CSCvn11045	HX node keeps crashing after node is restarted.	1. Verify if the interface is up and if you can ping the loopback interface:
		ifconfig -aping 127.0.0.1
		2. Bring up the loopback interface: **ip link set lo up**
		3. Check that the service is running:**status scvmclientstatus storfs**
		4. Start the following services:**start scvmclientstart storfs**
			3.5(1a)
			3.0(1e)

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)

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 authenticating UCS-M 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)

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
4.0(2f)	August 25, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.1(3e) is qualified for HX 4.0(2x) releases.
4.0(2f)	August 9, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.0(4m), and 4.1(3d) are qualified for HX 4.0(2x) releases.
4.0(2f)	June 21, 2021	Added Single Socket support in New Features, on page 2.
4.0(2f)	June 3, 2021	Created release notes for Cisco HX Data Platform Software, Release 4.0(2f).
4.0(2e)	May 7, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.0(4l) is qualified for HX 4.0(2x) releases.

Release	Date	Description
4.0(2e)	April 29, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.1(3c) is qualified for HX release 4.0(2d) and 4.0(2e).
4.0(2e)	April 28, 2021	Added support for Cisco HyperFlex HTML5 Plugin for VMware vCenter version 2.1.0.
4.0(2e), 4.0(2d), 4.0(2c	March 30, 2021	Updated Software Requirements for VMware ESXi to indicate support for VMware vCenter Versions 7.0 U1c through 7.0 U1d builds for HX 4.0(2e), 4.0(2d) and 4.0(2c).
4.0(x)	March 19, 2021	Updated link to indicate UCSM 4.1(2f) is the recommended Host Upgrade Utility (HUU) for M5 for HX 4.0(x).
4.0(2e)	March 17, 2021	Created release notes for Cisco HX Data Platform Software, Release 4.0(2e).
4.0(2x)	March 11, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.0(4k) is the recommended release.
4.0(2x)	February 18, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.1(2c) qualified for HX 4.0(2x).
4.0(2x)	February 10, 2021	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 to indicate UCSM 4.1(3b) qualified for HX 4.0(2x).
4.0(2c), 4.0(2d)	December 18, 2020	Updated Software Requirements for VMware ESXi - 4.0(x) Releases, on page 12 to indicate limitations for using vCenter 7.0 U1 with a 4.0(2x) HXDP cluster. Cisco HyperFlex CSI Interoperability Metrics added support for CCP, and Anthos Versions.

Release	Date	Description
4.0(2c)	December 7, 2020	Added support for Kubernetes Version 1.17
4.0(2d)	November 18, 2020	Created release notes for Cisco HX Data Platform Software, Release 4.0(2d).
4.0(2c)	October 29, 2020	Updated New Features, on page 2 with Cisco HyperFlex HTML5 Plugin 2.0.0.
4.0(2c)	October 22, 2020	Updated support for scale limits increase, new drives, and single socket configuration New Features, on page 2, Cisco HX Data Platform Compatibility and Scalability Details - 4.0(x) Releases, on page 16.
4.0(1x)	September 30, 2020	HX 4.0(1x) - End-of-Life Cisco HX Data Platform Software Version 4.0(1x) Product Bulletin.
4.0(2c)	September 24, 2020	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 HyperFlex Software Versions to indicate UCSM 4.1(1e) qualified for HX 4.0(2a), HX 4.0(2b) and HX 4.0(2c) releases.
4.0(2c)	September 14, 2020	Updated CIMC, and Host Upgrade Utility (HUU) for M5 to UCS 4.1(1h) for HX 4.0(2c).
4.0(1b), 4.0(2a), 4.0(2b)	September 4, 2020	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 HyperFlex Software Versions Recommended FI/Server Firmware and Software Requirements for Microsoft Hyper-V - 4.0(x) Releases, on page 14 to 4.0(4i) for 4.0(1a), 4.0(1b), 4.0 (2a), and 4.0(2b) releases.
4.0(2c)	August 21, 2020	Added support for HyperFlex Edge Short Depth Servers; All-flash (HXAF240c-M5SD) and Hybrid (HX240c-M5SD).

Release	Date	Description
4.0(2c)+	August 11, 2020	• Added column for M4/M5 Qualified FI/Server Firmware. Listed USC-M 4.1(2a) as qualified for HX 4.0(2c), 4.0(2b), and 4.0(1b). • Added CSCvv21905 to the list of Open Caveats for HX 4.0(2c), 4.0(2b), 4.0(2a), 4.0(1b), 4.0(1a).
4.0(1b)	July 23, 2020	Updated Recommended FI/Server Firmware - 4.0(x) Releases, on page 8 HyperFlex Software Versions starting with Release 4.0(1b): Added qualification for Cisco UCS Manager 4.0(4i), and 4.1(1d).
4.0(2c)	July 21, 2020	Added CSCvv05705 to list of Resolved Caveats for HX 4.0(2c).
4.0(2c)	July 14, 2020	Created release notes for Cisco HX Data Platform Software, Release 4.0(2c).
4.0(2b)	July 1, 2020	Updated Release 4.0(2b) support for ESXi 6.7 3 EP19 and ESX 6.7 U3 EP15. For more information, see Recommended Cisco HyperFlex HX Data Platform Software Releases - for Cisco HyperFlex HX-Series Systems
4.0(2b)	May 11, 2020	Added OTV for Stretched Cluster update in New Features for the 4.0(2b) release.
4.0(2a)	May 5, 2020	Updated Host Upgrade Utility (HUU) for M5 to UCS 4.0(4k) for HX 4.0(2a).
4.0(2b)	April 22, 2020	Created release notes for Cisco HX Data Platform Software, Release 4.0(2b).
4.0(1b)	March 30, 2020	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4h) for HX 4.0(1b).

Release	Date	Description
4.0(2a)	March 24, 2020	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4h) for HX 4.0(2a).
4.0(2a)	February 11, 2020	Created release notes for Cisco HX Data Platform Software, Release 4.0(2a).
3.5(2c)	January 15, 2020	Updated release notes for deferred Cisco HyperFlex Release HX 3.5(2c).
4.0(1b)	December 23, 2019	Updated M4 and M5 Recommended FI/Server Firmware to UCS 4.0(4e) for HX 4.0(1b), 4.0(1a), 3.5(2f), 3.5(2e), and 3.5(2d).
4.0(1b)	December 13, 2019	Added CSCvs28167 to the list of Resolved Caveats for HX 4.0(1b) and HX 4.0(1a).
4.0(1b)	November 25, 2019	Added CSCvs02466 to the list of Open Caveats.
4.0(1b)	November 7, 2019	Updated info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments. Updated info in the Storage Cluster Specifications.
4.0(1b)	October 25, 2019	Added CSCvj95606 and CSCvq24176 to the list of Security Fixes.
4.0(1b)	October 8, 2019	Updated Recommended FI/Server Firmware versions.
4.0(1b)	September 30, 2019	Updated HUU/CIMC info in the HyperFlex Edge and Firmware Compatibility Matrix for 4.x Deployments.
4.0(1b)	September 17, 2019	Added CSCvq41985 to new section for "Related Caveats".
4.0(1b)	September 16, 2019	Updated HUU/CIMC info in the HyperFlex Edge and Firmware Compatibility Matrix for 4.x Deployments.

Release	Date	Description
4.0(1b)	September 10, 2019	Updated HUU/CIMC info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x and 4.x Deployments.
4.0(1b)	August 28, 2019	Updated HUU/CIMC recommended firmware versions for HyperFlex Releases 4.0(1b), 3.5(2e) and 3.5(2d).
4.0(1b)	August 23, 2019	Updated Recommended FI/Server Firmware versions for HyperFlex Releases 3.5(2e) and 3.5(2d).
4.0(1b)	August 21, 2019	Added Cisco IMC version support info in the HyperFlex Edge and Firmware Compatibility Matrix for 4.x Deployments.
4.0(1b)	August 19, 2019	Created release notes for Cisco HX Data Platform Software, Release 4.0(1b).
4.0(1a)	August 8, 2019	Added bullet describing the "Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases" in the Upgrade Guidelines section.
4.0(1a)	August 5, 2019	Added important note indicating HyperFlex does not support UCS server firmware 4.0(4a), 4.0(4b), and 4.0(4c).
4.0(1a)	July 25, 2019	Updated component info for HX220c M5/HXAF220c M5 Cluster to VIC 1457 in "HyperFlex Edge and Firmware Compatibility Matrix for 4.x Deployments section".
4.0(1a)	July 22, 2019	• Added Release 3.5(2e) support for ESXi 6.7 U2, and updated notes for 3.5(2d), 3.5(2c) and 3.5(2b) support for ESXi 6.7 U2 in "Supported VMware vSphere Versions and Editions". • Added Witness Node Version for Release 3.5(2e).

Release	Date	Description
4.0(1a)	July 5, 2019	Added CSCvq39523 to the list of Open Caveats for Release 4.0(1a).
4.0(1a)	July 1, 2019	• Updated Important Note 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".
4.0(1a)	June 20, 2019	• Added CSCvp64140 and CSCvp98910 to the list of Open Caveats for Release 4.0(1a). • Updated HyperFlex Edge and Firmware Compatibility Matrix tables.
4.0(1a)	June 17, 2019	Updated Browser Recommendations.
4.0(1a)	May 31, 2019	Added information indicating 4.0(1a) features are supported on the Intersight Virtual Appliance and on Intersight.com.
4.0(1a)	May 21, 2019	• Updated M4/M5 Recommended FI/Server Firmware for 3.5(2b) and 4.0(1a). • Added bullet describing recommended use of the Hypercheck Health Check Utility in the "Upgrade Guidelines" section.
4.0(1a)	May 14, 2019	• Added New Feature description for DISA STIG Compliance. • Updated Storage Cluster Specifications for Hyper-V. • Added CSCvp66277 to list of Open Caveats.

Release	Date	Description
4.0(1a)	May 8, 2019	• Updated VMware vCenter Versions for 4.0(1a). • Updated Supported Microsoft Software versions.
4.0(1a)	May 3, 2019	• Added CSCvo36198 and CSCvk38003 to the list of Resolved Caveats. • Added CSCvp21417 to the list of Open Caveats.
4.0(1a)	April 29, 2019	Created release notes for Cisco HX Data Platform Software, Release 4.0(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.
Administration Guide for Hyper-V	Provides information about how to manage and monitor the Hyper-V cluster, encryption, data protection (replication and recovery), ReadyClones, Hyper-V Checkpoints, and user management. Interfaces include Cisco HyperFlex Systems, and the `hxcli` commands.

Document	Description
Administration Guide for Kubernetes	Provides information about HyperFlex storage integration for Kubernetes, information on Kubernetes support in HyperFlex Connect, and instructions on how to configure Cisco HyperFlex Container Storage Interface (CSI) storage integration for both the Cisco container platform and the RedHat OpenShift container platform.
Administration Guide for Citrix Workspace Appliance	Provides installation, configuration, and deployment procedures for a HyperFlex system to connect to Citrix Workspaces and associated Citrix Cloud subscription services such as Citrix Virtual Apps and Desktops Services. The Citrix Ready HCI Workspace Appliance program enables a Cisco HyperFlex System deployed on Microsoft Hyper-V to connect to Citrix Cloud.
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.
TechNotes	Provides independent knowledge base articles.
Release Notes for UCS Manager, Release 4.0	Provides information on recommended FI/Server firmware.