Release Notes for Cisco HX Data Platform, Release 3.0

First Published: 2018-04-13

Last Modified: 2020-08-03

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 3.0, and describe the features, limitations and caveats for the Cisco HX Data Platform.

Revision History

Release	Date	Description
3.0(1i)	March 10, 2020	Added upgrade reminder for HyperFlex versions 3.0(1x) as this version is unsupported and declared end-of-life. **Note** 3.0(1i) may be used only as an intermediate upgrade target when upgrading older HX clusters to the latest recommended release.
3.0(1i)	December 13, 2019	Added CSCvs28167 to the list of Open Caveats for HX 3.0(1i).
3.0(1i)	September 10, 2019	Updated HUU/CIMC info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments.

Release	Date	Description
3.0(1i)	August 28, 2019	Updated HUU/CIMC recommended firmware versions for HyperFlex Releases 3.5(2e) and 3.5(2d).
3.0(1i)	August 21, 2019	Added Cisco IMC version support info in the HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments.
3.0(1i)	August 8, 2019	Added bullet describing the "Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases" in the Upgrade Guidelines section.
3.0(1i)	August 5, 2019	Added important note indicating that HyperFlex does not support UCS server firmware 4.0(4a), 4.0(4b), and 4.0(4c).
3.0(1i)	November 26, 2018	Added CSCvm97558 to the list of resolved caveats for 3.0(1i)
3.0(1i)	November 15, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1i).
3.0(1h)	November 7, 2018	Updated release notes for Cisco UCS Manager, release 3.2(3h).
3.0(1h)	October 23, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1h).
3.0(1e)	October 10, 2018	Updated "Performance Impact for CPU Side-Channel Vulnerability Fixes" section.
3.0(1e)	August 31, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1e).
3.0(1d)	July 31, 2018	Updated release notes for NVIDIA V100 GPUs.
3.0(1d)	July 23, 2018	Updated release notes for Cisco UCS Manager, release 3.2(3g).
3.0(1d)	July 9, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1d).

Release	Date	Description
3.0(1c)	June 25, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1c).
3.0(1b)	June 18, 2018	Added CSCvj82452 to the list of resolved caveats and Upgrade Advisory for LSI DDA Drivers.
	June 13, 2018	Updated release notes for support for VMware vSphere version, 6.5 U2.
	June 8, 2018	Added CSCvh68059 to the list of resolved caveats in 3.0(1a).
	June 7, 2018	Updated cluster limits under the "Cisco HX Data Platform Storage Cluster Specifications" section.
	May 30, 2018	Updated release notes for deferred Cisco UCS Manager Releases, 3.2(3a) and 3.2(3b).
	May 14, 2018	Updated release notes for Cisco HX Data Platform Software, Release 3.0(1b).
3.0(1a)	May 4, 2018	Added "Performance Impact for CPU Side-Channel Vulnerability Fixes" section.
	May 1, 2018	Added 2.6(1e) version updates in the "Software Requirements" section.
	April 13, 2018	Created release notes for Cisco HX Data Platform Software, Release 3.0(1a).

Upgrade Advisory for LSI DDA Drivers

If you are experiencing performance degradation due to installed LSI DDA drivers (Version 2.6.40.0, 11/24/2017) on your Windows hosts, please contact Cisco TAC for resolution of this issue.

New Features in Release 3.0

This release delivers key advancements in mission critical and cloud-native workload support.

New Software Features

- *Multiple Hypervisors*—Support for VMware ESXi and Microsoft Hyper-V Server 2016. For detailed information on installation requirements and procedures, see Cisco HyperFlex Systems Installation Guide for VMware ESXi, and Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V.
- *Stretched clusters*—Ability to stretch a single HyperFlex Cluster across two datacenters enabling the highest level of availability. For more information, see Cisco HyperFlex Systems Stretched Cluster Guide.
- *Kubernetes FlexVolume driver*—Enables the automated provisioning of persistent volumes to Kubernetes pods.
- *Improved resiliency*— Enabled by Logical Availability Zones (LAZ) that, when enabled, automatically partition the cluster so that it is more resilient to node and disk failure scenarios. This feature can only be enabled on HyperFlex clusters with 8 or more converged nodes.
- *Disaster Recovery Workflow*—Enhanced Disaster Recovery workflow (Planned and unplanned VM migration) using Cisco HX Connect. For more information, see Cisco HyperFlex Systems Administration Guide, Release 3.0
- *REST APIs*—Additional developer guidance in the form of a quick start guide for HyperFlex REST APIs on Cisco DevNet. For more information see, Cisco HyperFlex Systems REST API Getting Started Guide.
- *Linked mode*—HyperFlex Plugin Support for environments utilizing vCenter's enhanced linked mode feature.

New Hardware Features

- *Enhanced HX Capacity Scaling options*—
  - Large Form Factor (LFF) HX M5 240 chassis with support up to 12 drives 6TB or 8 TB drives. **Note:** Currently supported in HX M5 240 nodes only. For more information. see Cisco HyperFlex HX-Series Spec Sheets
  - Support for 1.8 TB SFF HDD with maximum cluster capacity of 18.06TiB for Hybrid M5 Edge nodes only. For more information, see Cisco HyperFlex HX-Series Spec Sheets

  Enhanced Node Scaling options—Support for up to 32-nodes converged (per cluster) with 32 compute nodes. Cisco HyperFlex HX-Series Spec Sheets.
- *Intel Optane Support for higher drive level performance and higher endurance*—HyperFlex has qualified the latest flash memory innovation, 3D XPoint. Added Intel Optane NVMe SSD HX-NVMEXP-I375 as a new caching drive option. **Note:**

  Supported in M5 All Flash configurations only. For more information, see Cisco HyperFlex HX-Series Spec Sheets, *Cisco HX220c M5 HyperFlex Node Installation Guide*, and *Cisco HX240c M5 HyperFlex Node Installation Guide*.
- *Support for AMD Multiuser GPU (MxGPU) hardware-based graphics virtualization on HX240c M5 nodes*—AMD FirePro S7150 series GPUs are now available for HX240c M5 nodes. These graphic accelerators enable highly secure, high performance, and cost effective VDI deployments. For more information see, see Cisco HyperFlex HX-Series Spec Sheets, and the Cisco HX240c M5 HyperFlex Node Installation Guide. For instructions on deployment, see: Deploying AMD GPUs.

- *Expanded HyperFlex Edge configurations*—New HyperFlex Edge ordering PIDs provide more flexibility, simplify configuration and lower costs. For more information see, see Cisco HyperFlex HX-Series Spec Sheets
- *Ability to orchestrate AI Workloads with Cisco HyperFlex and Kubernetes on NVIDIA GPUs*—The combination of Cisco HyperFlex and Kubernetes on NVIDIA V100(16GB) GPUs enables automated deployments and maximum utilization to allow scale across nodes and clouds to get most from the AI infrastructure. This is currently supported on HX 240 All Flash M5 nodes only. For more information, see Cisco HyperFlex HX-Series Spec Sheets.

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.

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 an HX Storage Cluster are compatible.

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

Important Do not upgrade to these versions of firmware.

Do not upgrade to these versions of UCS Manager.

- Verify that the preconfigured HX servers have the same version of Cisco UCS server firmware installed. If the Cisco UCS Fabric Interconnects (FI) firmware versions are different, see the *Cisco HyperFlex Systems Upgrade Guide* for steps to align the firmware versions.
- **M4**: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M4 or HX220c M4) deployments, verify that Cisco UCS Manager 3.1(3j) or later is installed. Contact Cisco TAC for guidance.
- **M5**: For NEW hybrid or All Flash (Cisco HyperFlex HX240c M5 or HX220c M5) deployments, verify that Cisco UCS Manager 3.2(3i) or later is installed.
- For SED-based HyperFlex systems, ensure that the A (Infrastructure) and C (Rack server) bundles are at Cisco UCS Manager version 3.1(3h) or higher for M4 SED systems. Ensure that all bundles are at Cisco UCS Manager version 3.2(3i) or higher for M5 SED systems.

- 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.
- Cisco UCS Manager version 4.0(1a) is supported on HX Data Platform version 3.0(1e) and later.

Table 1: HyperFlex Software Versions for M4/M5 Servers

HyperFlex Release	M4 Recommended FI/Server Firmware *(be sure to review important notes above)	M5 Recommended FI/Server Firmware *(be sure to review important notes above)
3.0(1i)	3.2(3h), 3.1(3j)	3.2(3h)
3.0(1h)	3.2(3h), 3.1(3j)	3.2(3h)
3.0(1e)	3.2(3h), 3.1(3j)	3.2(3h)
3.0(1d)	3.2(3h), 3.1(3h)	3.2(3h)
3.0(1c)	3.2(3h), 3.1(3h)	3.2(3h)
3.0(1b)	3.2(3d), 3.1(3h)	3.2(3d)
3.0(1a)	3.2(3d), 3.1(3f)	3.2(3d)
2.6(1e)	3.2(3d), 3.1(3f)	3.2(3d)
2.6(1d)	3.2(3d), 3.1(3c)	3.2(3d)
2.6(1b)	3.2(2d), 3.1(3c)	3.2(2d)
2.6(1a)	3.2(2d), 3.1(3c)	3.2(2d)

HyperFlex Edge and Firmware Compatibility Matrix for 3.x Deployments

Cisco HX Data Platform, Release 3.x based Deployments

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

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

Table 2: HX220c M4 / HXAF220c M4 Cluster

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

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

Table 3: HX220c M5 / HXAF220c M5 Cluster

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

HyperFlex Licensing

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

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

VMware vSphere Licensing Requirements

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

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

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

Note

- HX Nodes have OEM licenses preinstalled. If you delete or overwrite the content of the boot drives after receiving the HX servers, you also delete the factory-installed licenses.
- OEM license keys is a new VMware vCenter 6.0 U1b feature. Earlier versions do not support OEM licenses.
- All factory-installed HX nodes share the same OEM license key. With vSphere OEM keys, the `Usage` count can exceed the `Capacity` value.
- When you add an HX host to vCenter through the **Add Host** wizard, in the **Assign license** section, select the **OEM license**.

  We obfuscate the actual vSphere OEM license key; for example, 0N085-XXXXX-XXXXX-XXXXX-10LHH.
- Standard, Essentials Plus, and ROBO editions are not available preinstalled on HX servers.

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

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

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

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

Software Requirements for VMware ESXi

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

Software Requirements for Microsoft Hyper-V

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

HyperFlex Software versions

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

- **Cisco HyperFlex M5 Converged nodes**— For Hybrid (Cisco HyperFlex HX240c M5, HX220c M5) and All Flash (Cisco HyperFlex HXAF240c M5, HXAF220c M5) verify that Cisco UCS Manager 3.2(3d) or later is installed. For detailed information on installation requirements and steps, see the *Cisco HyperFlex Systems Installation Guide on Microsoft Hyper-V*.

Table 4: Supported HyperFlex Software versions for M5 Servers on Hyper-V

HyperFlex Release	M5 Recommended UCS FI Firmware
3.0(1i)	3.2(3h)

Table 5: Supported Microsoft Software versions

Microsoft Component	Version
Windows Operating System (Windows OS)	Windows Server 2016 Datacenter OEM activated ISO and Retail ISOs are currently not supported. Earlier versions of Windows Server such as Windows 2012r2 are not supported. Non-English versions of the ISO are currently not supported.
Active Directory	A Windows 2012 or later domain and forest functionality level.

Supported Microsoft License Editions

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

Browser Recommendations

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.

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, at https://www.vmware.com/support/pubs/.

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

- **Cisco UCS Manager**

  The browser must support the following:

  - Java Runtime Environment 1.6 or later.
  - Adobe Flash Player 10 or higher is required for some features.

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

Cisco HX Data Platform Storage Cluster Specifications

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.

Node Limits for All Flash:

- Minimum converged nodes (per cluster): 3
- Maximum converged nodes (per cluster): 32
- (*Hyper-V, LFF, and Stretched Cluster only*) Maximum converged nodes (per cluster): 8
- Maximum compute-only nodes (per cluster): 32

Note The number of compute-only nodes cannot exceed two times the number of converged nodes. Additionally, the maximum number of compute-only nodes within a single cluster cannot exceed 32 nodes. Consider the following examples:

Number of converged nodes	Supported number of compute-only nodes
4	0-8
8	0-16
16	0-32
17-32	0-32 (max limit of 32 compute-only nodes)

Node Limits for Hybrid:

- Minimum converged nodes (per cluster): 3
- Maximum converged nodes (per cluster): 32
- (*Hyper-V, LFF, and Stretched Cluster only*) Maximum converged nodes (per cluster): 8
- Maximum compute-only nodes (per cluster): 32

Note The number of compute-only nodes cannot exceed two times the number of converged nodes. Additionally, the maximum number of compute-only nodes within a single cluster cannot exceed 32 nodes. Consider the following examples:

Number of converged nodes	Supported number of compute-only nodes
4	0-8
8	0-16
16	0-32
17-32	0-32 (max limit of 32 compute-only nodes)

Node Limits for HX Edge:

- For HX Edge configuration and node limitations, see the HyperFlex Edge Deployment Guide.

Cisco HX Data Platform storage clusters supported nodes:

- Converged nodes—**All Flash**: Cisco HyperFlex HXAF240c M5, HXAF220c M5, HXAF240c M4, and HXAF220c M4.
- Converged nodes—**Hybrid**: Cisco HyperFlex HX240c M5, HX220c M5, HX240c M4, and HX220c M4.
- Compute-only—Cisco B200 M3/M4, B260 M4, B420 M4, B460 M4, B480 M5, C240 M3/M4, C220 M3/M4, C480 M5, C460 M4, B200 M5, C220 M5, and C240 M5.

Upgrade Guidelines

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

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

Note 3.0(1i) may be used only as an intermediate upgrade target when upgrading older HX clusters to the latest recommended release.

HX Data Platform 1.7.x, 1.8.x, 2.0 and 2.1x clusters—Users from any version prior to 2.6(1a) must step through an intermediate version before upgrading to 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.
- **Software Advisory for M5 ESXi 6.0 Clusters**—For clusters running on ESXi 6.0, carefully review the related Software Advisory.
- **Required vCenter upgrade**—For enhanced security, HXDP release 3.5(1a) and later requires the use of TLS 1.2. Therefore, vCenter must be upgraded to 6.0 U3f or later prior to upgrading to HX 3.5. In addition, ESXi should be upgraded to 6.0 U3 or 6.5 U1 to meet HXDP 3.5 compatibility requirements.
- **Minimum HXDP version for upgrade**—HXDP clusters running 2.1(1c) or later may upgrade directly to 3.5.
- **HX Data Platform 1.7.x and 1.8.x clusters**—Users upgrading from 1.7.x, 1.8.x, 2.0.x,2.1(1a), or 2.1(1b) must step through an intermediate version before upgrading to 3.5 or later releases. For more information, see the Cisco HyperFlex Systems Upgrade Guide.
- **Cluster Readiness**—Ensure that the cluster is properly bootstrapped and the updated plug-in loaded before proceeding. **Cluster bootstrap is required for every upgrade.**
- **Initiating Upgrade—**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. Use the HX Connect UI when upgrading from 2.5(1a) or later releases. Do not use the Tech Preview UI (pre-2.5) for upgrades.
- **vSphere 5.5 Upgrades**—Users on vSphere 5.5 must upgrade to 6.0 U3/6.5 U1 before starting HX Data Platform upgrade. vSphere 5.5 support was deprecated with HX Data Platform 2.5(1a) and upgrade fails if attempted.
  - For HX220 users running 5.5, contact TAC for upgrade assistance.
  - For HX240 users running 5.5, upgrade components in the following order.
    1. Upgrade vCenter to 6.0 U3c or 6.5 U1. If upgrading to 6.5, you must upgrade your vCenter in place. Using a new vCenter 6.5 is not supported for users migrating from 5.5.
    2. Upgrade ESXi to 6.0/6.5 using the offline zip bundle.

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

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

  - **If Upgrading to vSphere 6.5**:
    - Certain cluster functions such as native and scheduled snapshots, ReadyClones, and Enter/Exit HX Maintenance Mode will not operate from the time the upgrade is started until the HX Data Platform upgrade to 2.5 or later is complete.
    - After upgrading ESXi using the offline zip bundle, use the ESX Exit Maintenance Mode option. The HX Exit Maintenance Mode option does not operate in the vSphere Web Client until the HX Data Platform upgrade is complete.

- **vSphere 6.0 Upgrades**—Users on vSphere 6.0 migrating to 6.5, upgrade components in the following order:
  1. HX Data Platform upgrade (and optionally the UCS firmware).
  2. Upgrade vCenter Server following VMware documentation and best practices. Optionally, deploy a new vCenter server and perform `stcli cluster reregister`.
  3. Upgrade ESXi to 6.5 using the offline zip bundle.
- **M4 Server Firmware Upgrades**—Server firmware should be upgraded to ensure smooth operation and to correct known issues.

  

  Note
  - Users are encouraged to upgrade to 3.1(3c) C-bundle or later whenever possible.
  - Users running C-bundle versions prior to 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(1d) or later.
- **M4/M5 Mixed Domains**—A mixed domain occurs when a new, separate M5 cluster is installed under the same UCS domain that contains existing M4 cluster(s). Under these conditions, orchestrated UCS server firmware upgrade will not operate until HX Data Platform 2.6 or later is installed on the M4 clusters. Therefore, it is best practice to first upgrade UCS server firmware to the latest 4.0(1) patch release prior to adding a new M5 cluster to the existing UCS domain. Additionally, any 1.7 HX Data Platform clusters must first be upgraded prior to adding any new M5 clusters to the same domain.
- **Maintenance Window**—If upgrading both HX Data Platform and UCS firmware, either a combined or split upgrade can be selected through the vSphere HX Data Platform depending on the length of the maintenance window. Direct firmware upgrade using server firmware auto install through Cisco UCS Manager should not be attempted. Instead, use the UCS server upgrade orchestration framework provided by the HX Data Platform.
- **HX Data Platform 2.1(1b) with SED**—Upgrading SED-ready systems running 2.1 require UCS infrastructure and server firmware upgrades.

Mixed Cluster Expansion Guidelines

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

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

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

- Only expansion workflow is supported to create a mixed cluster. Initial cluster creation with mixed M4/M5 servers is not supported.
- All M5 servers must match the form factor (220/240), type (Hybrid/AF), security capability (Non-SED only) & disk configuration (QTY, capacity, and non-SED) of the existing M4 servers.
  - HX220-M5 will use a maximum of 6 capacity disks (2 disk slots to remain empty) when mixed with HX220-M4.
- HX Edge, SED, LFF, Hyper-V, and Stretched Clusters do not support mixed M4 and M5 clusters.

Performance Impact for CPU Side-Channel Vulnerability Fixes

Spectre and Meltdown are names given to critical vulnerabilities existing in several modern CPUs. These vulnerabilities can allow an unprivileged local attacker, in specific circumstances, to access information belonging to other processes or the operating system kernel. The vulnerabilities are all variants of the same attack and differ in the way that speculative execution is exploited. For more information, see CPU Side-Channel Information Disclosure Vulnerabilities

When deploying the Spectre and Meltdown related security fixes in the hypervisor along with the updated microcode provided with patched server firmware and guest OS updates, customers may experience performance impact. Performance testing indicates that in a ESXi hypervisor environment the impact to HXDP storage performance on HyperFlex nodes is less than 5%.

For more information on Side Channel Analysis vulnerabilities that affect HyperFlex Controller VM software, see CSCvh68612.

Security Fixes

The following security issues are resolved:

Release	Defect ID	CVE	Description
3.0(1d)	CSCvk05700	CVE-2018-12015	The vulnerabilities associated with the "file" software package version included in the Cisco HX Data Platform.
3.0(1d)	CSCvk05679	CVE-2014-9620 CVE-2014-9653 CVE-2015-8865 CVE-2018-10360 CVE-2014-9621	The vulnerabilities associated with the "file" software package version included in the Cisco HX Data Platform.

Release	Defect ID	CVE	Description
3.0(1d)	CSCvj73115	CVE-2018-6798 CVE-2018-6913 CVE-2016-6185 CVE-2015-8853 CVE-2017-6512 CVE-2018-6797	The vulnerabilities associated with the "Ubuntu Perl" software package version included in the Cisco HX Data Platform.
3.0(1d)	CSCvj73114	CVE-2018-1123 CVE-2018-1122 CVE-2018-1126 CVE-2018-1125 CVE-2018-1124	The vulnerabilities associated with the "Ubuntu procps" software package version included in the Cisco HX Data Platform.
3.0(1d)	CSCvj63266	CVE-2018-1000300 CVE-2018-1000301	The vulnerabilities associated with the cURL software package version included in the Cisco HX Data Platform.
3.0(1d)	CSCvj61269	CVE-2018-0494	The vulnerabilities associated with the GNU Wget software package version included in the Cisco HX Data Platform.
3.0(1b)	CSCvi88572	CVE-2016-3186 CVE-2016-5102 CVE-2016-5318 CVE-2017-11613 CVE-2017-12944 CVE-2017-17095 CVE-2017-18013 CVE-2017-5563 CVE-2017-9117 CVE-2017-9147 CVE-2017-9935 CVE-2018-5784	The vulnerabilities of **Ubuntu** with the **LibTIFF** sofware package included in Cisco HyperFlex HX Data Platform.
3.0(1b)	CSCvi68120	CVE-2018-5146	The vulnerabilities associated with aversion of **libvorbis** software included in Cisco HyperFlex HX Data Platform.

Release	Defect ID	CVE	Description
3.0(1a)	CSCvh68612	CVE-2017-5715 CVE-2017-5753 CVE-2017-5754	The vulnerabilities associated with Spectre and Meltdowns impacting M5 servers. A complete resolution requires a CPU microcode update, a hypervisor update, and guest OS update for VMs running on the system. Refer to the Cisco Security Advisory for more information.
3.0(1a)	CSCvh55595	CVE-2017-5715 CVE-2017-5753 CVE-2017-5754	The vulnerability associated Side Channel Analysis.

Resolved Caveats in Release 3.0(1i)

Defect ID	Symptom	First Release Affected	Resolved in Release

CSCvm97558	The controller VM restarts and a message similar to the following is seen in the `/var/log/kern.log` file:
	
	`kernel: python invoked oom-killer: gfp_mask=0x24280ca, order=0, oom_score_adj=0`
	`kernel: python cpuset=/ mems_allowed=0`
	`kernel: CPU: 6 PID: 26270 Comm: python Tainted: G OE 4.4.0-112-generic #135-Ubuntu`
	`kernel: Hardware name: VMware, Inc. VMware Virtual Platform/440BX Desktop Reference Platform, BIOS 6.00 09/21/2015`
	`kernel: 0000000000000286 4ca55bb5d320913c ffff88084b8c7af8 ffffffff813fc233`
	`kernel: ffff88084b8c7cb0 ffff880c060ab800 ffff88084b8c7b68 ffffffff8120dafe`
	`kernel: ffff88084b8c7b18 ffffffff8114121a ffff88084b8c7b98 ffffffff811a8bd6`
	`kernel: Call Trace:`
	`kernel: [<ffffffff813fc233>] dump_stack+0x63/0x90`
	`kernel: [<ffffffff8120dafe>] dump_header+0x5a/0x1c5`
	`kernel: [<ffffffff8114121a>] ? __delayacct_freepages_end+0x2a/0x30`
	`kernel: [<ffffffff811a8bd6>] ? do_try_to_free_pages+0x2a6/0x3b0`
	`kernel: [<ffffffff811946a2>] oom_kill_process+0x202/0x3c0`
	`kernel: [<ffffffff81194ac9>] out_of_memory+0x219/0x460`
	`kernel: [<ffffffff8119aad5>] __alloc_pages_slowpath.constprop.88+0x965/0xb00`
	`kernel: [<ffffffff8119aef6>] __alloc_pages_nodemask+0x286/0x2a0`
	`kernel: [<ffffffff811e603d>] alloc_pages_vma+0xad/0x250`
	`kernel: [<ffffffff811c3b2e>] handle_mm_fault+0x110e/0x1820`
	`kernel: [<ffffffff8106b687>] __do_page_fault+0x197/0x400`
	`kernel: [<ffffffff8106b912>] do_page_fault+0x22/0x30`
	`kernel: [<ffffffff81849ac8>] page_fault+0x28/0x30`
	`kernel: Out of memory: Kill process 3243 (hxmanager) score 56 or sacrifice child`
		3.0(1c)	3.0(1i)

CSCvm66552	Multiple simultaneous 3.8TB SED SSD drive failures due to a drive firmware bug may cause the HX cluster to go offline. For more details, see the related Software Advisory.	3.0(1c)	3.0(1i)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvm59485	After a node is upgraded from release 2.5(1b) to 3.0(1c), the SED disks enter an ignored state.	3.0(1c)	3.0(1i)
CSCvm00138	In vCenter, the following warning appears for the ESXi host: `"Connectivity to Callhome endpoint failed."`	3.0(1c)	3.0(1i)
CSCvm89133	During an upgrade from release 2.5(1c) to 3.0(1e), a single and subsequently all VMs go down. The cluster does not start although the node that failed to upgrade is manually upgraded. Storfs process on the node panics and generates a core file in `/var/core`.	2.5(1c)	3.0(1i)
CSCvb54848	The vSphere Replication plugin fails to load after the HX Data Platform plugin is installed.	1.7.1	3.0(1i)

Resolved Caveats in Release 3.0(1h)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvm46965	In some cases with sustained write heavy workloads, cluster rebalance/upgrade may be unresponsive.	3.0(1d)	3.0(1h)
CSCvm22903	For Stretched cluster deployments, HX Installer may hang after the installation workflow is initiated.	3.0(1c)	3.0(1e)

Resolved Caveats in Release 3.0(1e)

Defect ID	Symptom	First Release Affected	Resolved in Release
Hyper-V			
CSCvk42435	VMs do not power on after Live Migration under the following conditions: • The Storage Controller VM on one node is down for several minutes, or • The entire host is down, or • VMs are powered on manually on other nodes after performing live migration.	3.0(1d)	3.0(1e)
CSCvk18743	The Storage Controller VM is down for an extended period of time that may cause the VMs to power off.	3.0(1b)	3.0(1e)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvi14568	Import VM operation fails using Hyper-V Manager (Remote) when the exported VM is stored on HX-Datastore.	3.0(1a)	3.0(1e)
ESXi			
CSCvm20044	HX Node may be impacted with a controller VM or related reboots/failures due to SAMSUNG Disk MZ7LM480HMHQ failures.	3.0(1c)	3.0(1e)

Resolved Caveats in Release 3.0(1d)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvi73801	The following roles are enabled by default on HX Hyper-V node: • Web Server • Remote access	3.0(1a)	3.0(1d)
CSCvj74517	Storfs service crashes on nodes running HXDP release 2.5(1c) with segmentation panic.	2.5(1c)	3.0(1d)

Resolved Caveats in Release 3.0(1c)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvk02032	HX Connect login fails after an upgrade from release 3.0(1b) to 3.0(1c).	2.6(1a)	3.0(1c)
CSCvj90352	When a physical disk is replaced by a new disk, unused and removed disks remain visible to the user in HX Connect UI and stcli.	3.0(1b)	3.0(1c)
CSCvj66157	SED drive failure may cause the UCS/HX cluster to go down.	2.6(1d)	3.0(1c)
CSCvf98675	During the reboot phase of deployment server disks were not detected. HX deployment fails with error: "Controller VM Install Device (tmpfs) not persistent. Cannot Install Packages".	2.6(1a)	3.0(1c)

Resolved Caveats in Release 3.0(1b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvj82452	LSI DDA driver version 2.60.40.0 installed on the Windows Host shows significant performance degradation	3.0(1a)	3.0(1b)
CSCvj18863	During or after an upgrade to 3.0(1a), the cluster may experience one or more nodes restarting due to a SW panic.	3.0(1a)	3.0(1b)
CSCvi95613	HX Installer fails when HX LFF servers are on the same UCS domain.	3.0(1a)	3.0(1b)
CSCvi92768	Adding two nodes via stcli to LAZ cluster results in an error.	3.0(1a)	3.0(1b)
CSCvi92070	If you select multiple VMs in **Failover Cluster Manager** UI, only one VM becomes highly available.	3.0(1a)	3.0(1b)
CSCvi84992	During offline upgrade, HX Connect displays the following a "`Server Call Failure`" error.	3.0(1a)	3.0(1b)
CSCvi82755	The HyperFlex Clusters list view in Cisco Intersight might not display the Hypervisor version for the ESXi cluster.	3.0(1a)	3.0(1b)
CSCvi76137	When a node is in maintenance mode, disk API is stuck.	3.0(1a)	3.0(1b)
CSCvi73807	Cloned user VM will not become a part of a failover cluster resource.	3.0(1a)	3.0(1b)
CSCvi73417	In HX Connect, while a snapshot is in progress, VMs list page appears as blank.	3.0(1a)	3.0(1b)
CSCvi71157	Storage migration for a virtual machine fails in Hyper-V Manager with an error, "There was an error during move operation."	3.0(1a)	3.0(1b)
CSCvi63527	Datastore accessibility on the host is not validated before putting a host in maintenance mode.	3.0(1a)	3.0(1b)
CSCvh54563	Generating HyperFlex logs on HX240 platform causes an all paths down state for the cluster. This issue is caused by the **vmw_ahci** driver in ESXi 6.5 GA release.	2.6(1b)	3.0(1b)
CSCvh99126	Creating virtual machine from template works, but without template workflow from VMM.	3.0(1a)	3.0(1b)
CSCvg41818	HX snapshot error when taking snapshot with GPU shared PCI device attached for powered off VM.	2.5(1c)	3.0(1b)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvg05306	In rare cases, after kernel upgrade, the gateway IP address changes for eth0 and eth1.	3.0(1a)	3.0(1b)
CSCvb54848	vSphere Replication Plug-in fails after HX Plug-in is deployed.	1.7.1	3.0(1b)

Resolved Caveats in Release 3.0(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvh68059	Snapshot failure occurs and shows RAM disk (/var/) as full in the `vmkernel.log`.	2.6(1e)	3.0(1a)
CSCvg52652	HX Connect UI incorrectly displays the date for the ready clone activity.	2.5(1d)	3.0(1a)
CSCvg09618	Test-KMIP does not fail although primary (and only) configured server is down, and the key operation fails. **Note** Although there is no secondary KMIP Server IP address or name specified, ”testing KMIP server connectivity” task is being run for the previously configured secondary KMIP server.	2.5(1c)	3.0(1a)
CSCvg26090	In cases, where automatic synchronization between primary and secondary KMIP servers does not happen, any key operation(rekey/disable) would fail with the following error: “failure to locate key”.	2.5(1c)	3.0(1a)
CSCvf84968	Combined upgrade using stcli command fails without Cisco UCS Manager and vCenter credentials.	2.5(1c)	3.0(1a)
CSCvf12501	Sometimes, after fresh cluster creation, controller VM memory usage is high and in a critical state.	This is a known VMware issue. See the article, https://kb.vmware.com/s/article/2149787.	3.0(1a)
CSCve17284	During the HX cluster upgrade, performance charts shows IOPS drop temporarily.	2.0(1a)	3.0(1a)
CSCvd88557	When creating many datastores through the`stcli` command-line, a temporary error displays indicating that some datastores fail to mount.	2.0(1a)	3.0(1a)

Defect ID	Symptom	First Release Affected	Resolved in Release
CSCvc32497	Cluster creation or expansion fails when UCSM configuration is not chosen as an advanced option in the Cisco HX Data Platform Installer. This happens due to non-reachability of ESX. In the Cisco HX Data Platform Installer's configuration page, you will see that the default VLAN for hx-inband-mgmt 3091 is tagged to the ESX and not the user-specified VLAN.	2.0(1a)	3.0(1a)

Open Caveats in Release 3.0(1i)

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

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvs28167	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”.		2.6(1e)

Defect ID	Symptom	Workaround	Defect Found in Release
		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' 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. 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-v3.5.2-31725-esx.ova~~	

Defect ID	Symptom	Workaround	Defect Found in Release

		`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.	

CSCvn12846	When one or more nodes are down in the cluster, **Generate support bundle** button in the UI will work for the first time but will be grayed out after that.	The support bundles from the first attempt will be available for download. Until the nodes are back up, you can login to the controller VM to collect the support bundle using the **storfs-support** command.	3.0(1i)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: "There are locked drives on the system, unlock them and retry deployment." In addition, after upgrade from release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more information to recover from this issue.	3.0(1c)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster.	Ensure there are sufficient replication IP addresses available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
Management			
CSCvk39622	HX Connect displays alarms with a “Lockdown mode enabled on one or more nodes in the cluster” message. In addition, the alarms are manually reset to green.	To resolve this issue, do the following: **1.** Click **Manage** in the VMware Host Client inventory and click **Security & Users**. **2.** Click **Lockdown mode**. **3.** Click **Add user exception**, enter the name of the user, and click **Add exception**.	3.0(1d)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the 3.0(1b) following: **1.** Delete the dummy stCtlVMsfrom the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1e)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Hyper-V			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvm42278	Datastore access leads to frequent alerts. In addition, the SMB SCVM client log file (/var/log/springpath/debug-smbscvmclient.log) shows a message similar to the following for host data IP address of the host on which that controller is hosted.	Contact Cisco TAC for further assistance with this issue.	3.0(1e)
CSCvm05523	In rare cases, live migration of VMs fails with error code 0x138D.	Refer to the following workarounds from Microsoft and retry the operation: • Server 2016 S2D Cluster unable to Drain Role • Live migrations fail during drain from Cluster-Aware Updating • Draining Nodes for Planned Maintenance with Windows Server 2012	3.0(1e)
CSCvk37044	Deployment fails when the username in the user account contains a period (.").	To avoid this issue, do NOT use a period (".") in the username.	3.0(1e)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi59119	Duplicating an existing datastore name that differs only in letter case might result in unknown behavior.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
CSCvh25238	During HX Data Platform deployment, the controller VM may be assigned with only one DNS address if you add more than one IP address for DNS.	Usually the primary DNS is sufficient for the HX Controller VM to work. If you need additional DNS, edit the eth0 interface file in the controller VM to add the additional DNS.	3.0(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: **1.** Move the disk to the same folder on the datastore. **2.** Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
Encryption			
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCS Manager might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)

Open Caveats in Release 3.0(1h)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			
CSCvm42278	Datastore access leads to frequent alerts. In addition, the SMB SCVM client log file (/var/log/springpath/debug-smbscvmclient.log) shows a message similar to the following for host data IP address of the host on which that controller is hosted.	Contact Cisco TAC for further assistance with this issue.	3.0(1e)
CSCvm05523	In rare cases, live migration of VMs fails with error code 0x138D.	Refer to the following workarounds from Microsoft and retry the operation: • Server 2016 S2D Cluster unable to Drain Role • Live migrations fail during drain from Cluster-Aware Updating • Draining Nodes for Planned Maintenance with Windows Server 2012	3.0(1e)
CSCvk37044	Deployment fails when the username in the user account contains a period (.").	To avoid this issue, do NOT use a period (".") in the username.	3.0(1e)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvi59119	Duplicating an existing datastore name that differs only in letter case might result in unknown behavior.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
CSCvh25238	During HX Data Platform deployment, the controller VM may be assigned with only one DNS address if you add more than one IP address for DNS.	Usually the primary DNS is sufficient for the HX Controller VM to work. If you need additional DNS, edit the eth0 interface file in the controller VM to add the additional DNS.	3.0(1a)
Install, Upgrade, Expand			
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: "There are locked drives on the system, unlock them and retry deployment." In addition, after upgrade from release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more information to recover from this issue.	3.0(1c)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster.	Ensure there are sufficient replication IP addresses available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
Management			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvk39622	HX Connect displays alarms with a “Lockdown mode enabled on one or more nodes in the cluster” message. In addition, the alarms are manually reset to green.	To resolve this issue, do the following: **1.** Click **Manage** in the VMware Host Client inventory and click **Security & Users**. **2.** Click **Lockdown mode**. **3.** Click **Add user exception**, enter the name of the user, and click **Add exception**.	3.0(1d)
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the 3.0(1b) following: **1.** Delete the dummy stCtlVMsfrom the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1e)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
Encryption			

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

Open Caveats in Release 3.0(1e)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			
CSCvm42278	Datastore access leads to frequent alerts. In addition, the SMB SCVM client log file (/var/log/springpath/debug-smbscvmclient.log) shows a message similar to the following for host data IP address of the host on which that controller is hosted.	Contact Cisco TAC for further assistance with this issue.	3.0(1e)
CSCvm05523	In rare cases, live migration of VMs fails with error code 0x138D.	Refer to the following workarounds from Microsoft and retry the operation: • Server 2016 S2D Cluster unable to Drain Role • Live migrations fail during drain from Cluster-Aware Updating • Draining Nodes for Planned Maintenance with Windows Server 2012	3.0(1e)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvk37044	Deployment fails when the username in the user account contains a period (".").	To avoid this issue, do NOT use a period (".") in the username.	3.0(1e)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvi59119	Duplicating an existing datastore name that differs only in letter case might result in unknown behavior.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
CSCvh25238	During HX Data Platform deployment, the controller VM may be assigned with only one DNS address if you add more than one IP address for DNS.	Usually the primary DNS is sufficient for the HX Controller VM to work. If you need additional DNS, edit the eth0 interface file in the controller VM to add the additional DNS.	3.0(1a)
Install, Upgrade, Expand			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: "There are locked drives on the system, unlock them and retry deployment." In addition, after upgrade from release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more information to recover from this issue.	3.0(1c)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster.	Ensure there are sufficient replication IP addresses available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
Management			
CSCvk39622	HX Connect displays alarms with a “Lockdown mode enabled on one or more nodes in the cluster” message. In addition, the alarms are manually reset to green.	To resolve this issue, do the following: **1.** Click **Manage** in the VMware Host Client inventory and click **Security & Users**. **2.** Click **Lockdown mode**. **3.** Click **Add user exception**, enter the name of the user, and click **Add exception**.	3.0(1d)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the 3.0(1b) following: **1.** Delete the dummy stCtlVMsfrom the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1e)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: **1.** Move the disk to the same folder on the datastore. **2.** Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
Encryption			
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCS Manager might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)

Open Caveats in Release 3.0(1d)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvk42435	VMs do not power on after Live Migration under the following conditions: • The Storage Controller VM on one node is down for several minutes, or • The entire host is down, or • VMs are powered on manually on other nodes after performing live migration.	Contact Cisco TAC for further assistance on this issue.	3.0(1d)
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)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi73801	The following roles are enabled by default on HX Hyper-V node: • Web Server • Remote access	If needed, manually disable roles.	3.0(1a)
CSCvi73796	Windows Event Log or Viewer displays errors due to invalid VMQ network configuration.	This is a known issue and does not impact functionality.	3.0(1a)
CSCvi70645	Directory listing on a host appears as hung for a non-existent datastore.	Use CTRL+C to cancel the operation and re-try.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi37407	Service profile association is stuck at "`Waiting for Storage Subsystem to initialize`" stage after re-acknowledging the server.	If you see this issue, do the following: • Disassociate Service profile. • Decommission the server. • Re-acknowledge the server and associate the service profile again.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvi14568	Import VM operation fails using Hyper-V Manager (Remote) when the exported VM is stored on HX-Datastore.	If you see this issue, do the following: • Run the following Powershell commands: `C:\Users\Administrator.HV-AD10> Invoke-Command -Comp stfs-028b -Command { Import-VM -Path "\\ctlvm28b.hv-ad1.local\ds1\export\vm7\Virtual Machines\[illegible].vmcx" -VhdDestinationPath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SnapshotFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SmartPagingFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -Copy -GenerateNewId | Rename-VM -NewName vm102 }` • This is a known Microsoft issue. See: Microsoft article: Remote import-vm in Hyper-V.	3.0(1a)
CSCvh80044, CSCvi59119	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.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
Install, Upgrade, Expand			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: "There are locked drives on the system, unlock them and retry deployment." In addition, after upgrade from release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more information to recover from this issue.	3.0(1c)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster.	Ensure there are sufficient replication IPs available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)
CSCvg43082	During upgrade from 2.5(1c) to 2.6(1b), device connect might get disabled.	Enable device connect through HX Connect.	2.6(1a)
CSCvg05306	In rare cases, after kernel upgrade, the gateway IP address changes for eth0 and eth1.	Contact Cisco TAC for further guidance after collecting logs.	2.5(1a)
CSCvf98675	During the reboot phase of deployment server disks were not detected. HX deployment fails with error: "*Controller VM Install Device (tmpfs) not persistent. Cannot Install Packages*".	Re-acknowledge the server in UCS Manager. Then retry the deployment.	2.6(1a)
CSCvf93815	FI may reboot during cluster deployment, and installation hangs.	Reboot Installer VM and restart the HX cluster installation workflow.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf93812	Upgrade fails when the ESXi host is rebooted manually during the upgrade process. The rebooted ESXi host shows up as *'not-connected'/'not-responding'* in the VMware VCenter cluster. Subsequent upgrade attempts will fail during ESXi validation with authentication failure errors.	Manually recreate the HX user through vCenter. • Log in to ESXi command line. • Manually start vSphere HA (FDM) service. Reconnect the host to vCenter. Re-register this VM using vim command on the ESXi host. Please see VMWare KB: Cannot install the vSphere HA (FDM) agent on ESXi host (2007739) for more details. • Recover HX user. Contact Cisco TAC for assistance. • Add entry to */etc/hosts*. • Remount datastores. • Retry upgrade. Do not manually reboot any of the ESXi hosts in the cluster during the upgrade process.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)

CSCvb91838	Cluster expansion failed with no operational DNS server from the list.	If the DNS server becomes non-operational after deployment or cluster creation, add a new operational DNS to the controller. Use the following commands:
		`# stcli services dns remove --dns <non_operational_dns_ip>`
		`# stcli services dns add --dns <operational_dns_ip>`
			1.8(1c)

CSCvb29790	Cluster creation fails due to failure to locate vCenter server.	In the vSphere Web Client, change the vCenter host name to an IP address in the `config.vpxd.sso.sts.uri` variable.	1.8(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
Management			
CSCvk39622	HX Connect displays alarms with a “Lockdown mode enabled on one or more nodes in the cluster” message. In addition, the alarms are manually reset to green.	To resolve this issue, do the following: **1.** Click **Manage** in the VMware Host Client inventory and click **Security & Users**. **2.** Click **Lockdown mode**. **3.** Click **Add user exception**, enter the name of the user, and click **Add exception**.	3.0(1d)
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the following: **1.** Delete the dummy stCtlVMs from the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1b)
CSCvi34303	HX Connect UI displays an error when any table is exported in .CSV format, and opened in excel.	None.	3.0(1a)
CSCvg69228	Previously deleted/removed disks may be erroneously listed in the inventory.	This is a benign issue that can be safely ignored.	2.6(1b)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvg26340	For releases higher than 2.5(1c), `node failures tolerable to be` shows zero while the cluster is healthy.	If you encounter this issue, contact TAC for a workaround.	2.5(1c)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
CSCvf27609	Query for recovery job returns `summary_step_state` and `state` fields.	Refer to the `state` field only. Ignore the information in the `summary_step_state` field.	2.5(1a)
Encryption			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh79736	Disable remote security times out and retry disable operation fails in HX Connect.	After the disable operation times out, issue rescan-inventory for the servers whose SAS controllers are still secure. Then retry disable encryption from the HX Connect. Run the following commands from any controller VM: **`hx.py --get-org getOrg.json --server-serial <serial number of server>`** **`hx.py --rescan-drive-inventory -f getOrg.json --server-serial <serial number of server>`**	3.0(1a)
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCSM might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)
CSCvf04240	Encryption may not be enabled on the new node, after it is added to the cluster.	One of the potential causes is that the serial number was not reported correctly from the ESX host. Restart `hostd` service on the ESX host, and enable encryption on the cluster from the HX Connect UI. All the nodes that already have encryption enabled are not impacted.	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCve91866	Cannot modify encryption KMIP policy on UCSM to clear an IP address.	UCSM does not allow this behavior. On UCSM, delete the KMIP policy, adjusting for the IP addresses as needed, and retry the task.	2.5(1a)

Open Caveats in Release 3.0(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			
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)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi73801	The following roles are enabled by default on HX Hyper-V node: • Web Server • Remote access	If needed, manually disable roles.	3.0(1a)
CSCvi73796	Windows Event Log or Viewer displays errors due to invalid VMQ network configuration.	This is a known issue and does not impact functionality.	3.0(1a)
CSCvi70645	Directory listing on a host appears as hung for a non-existent datastore.	Use CTRL+C to cancel the operation and re-try.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, "`The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi37407	Service profile association is stuck at "`Waiting for Storage Subsystem to initialize`" stage after re-acknowledging the server.	If you see this issue, do the following: • Disassociate Service profile. • Decommission the server. • Re-acknowledge the server and associate the service profile again.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvi14568	Import VM operation fails using Hyper-V Manager (Remote) when the exported VM is stored on HX-Datastore.	If you see this issue, do the following: • Run the following Powershell commands: `C:\Users\Administrator.HV-AD10> Invoke-Command -Comp stfs-028b -Command { Import-VM -Path "\\ctlvm28b.hv-ad1.local\ds1\export\vm\Virtual Machines\[illegible].vmcx" -VhdDestinationPath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SnapshotFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SmartPagingFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -Copy -GenerateNewId | Rename-VM -NewName vm102 }` • This is a known Microsoft issue. See: Microsoft article: Remote import-vm in Hyper-V.	3.0(1a)
CSCvh80044, CSCvi59119	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.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
Install, Upgrade, Expand			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvm97558		**(Recommended)** If you see this issue, upgrade your cluster to release 3.0(1i). **(Optional)** In case an upgrade cannot be done immediately, disable HTTPS based nightly ASUP collection using the following command: (From a SSH session to the controller) `# stcli services sch disable` After subsequent upgrade to release 3.0(1i), enable ASUP using the following command: # stcli services sch enable	3.0(1c)

Defect ID	Symptom	Workaround	Defect Found in Release

	The controller VM restarts and a message similar to the following is seen in the `/var/log/kern.log` file:
	
	`kernel: python invoked oom-killer: gfp_mask=0x24280ca, order=0, oom_score_adj=0`
	`kernel: python cpuset=/ mems_allowed=0`
	`kernel: CPU: 6 PID: 26270 Comm: python Tainted: G OE 4.4.0-112-generic #135-Ubuntu`
	`kernel: Hardware name: VMware, Inc. VMware Virtual Platform/440BX Desktop Reference Platform, BIOS 6.00 09/21/2015`
	`kernel: 0000000000000286 4ca55bb5d320913c ffff88084b8c7af8 ffffffff813fc233`
	`kernel: ffff88084b8c7cb0 ffff880c060ab800 ffff88084b8c7b68 ffffffff8120dafe`
	`kernel: ffff88084b8c7b18 ffffffff8114121a ffff88084b8c7b98 ffffffff811a8bd6`
	`kernel: Call Trace:`
	`kernel: [<ffffffff813fc233>] dump_stack+0x63/0x90`
	`kernel: [<ffffffff8120dafe>] dump_header+0x5a/0x1c5`
	`kernel: [<ffffffff8114121a>] ?`		

Defect ID	Symptom	Workaround	Defect Found in Release
	__delayacct_freepages_end+0x2a/0x30 kernel: [<ffffffff811a8bd6>] ? do_try_to_free_pages+0x2a6/0x3b0 kernel: [<ffffffff811946a2>] oom_kill_process+0x202/0x3c0 kernel: [<ffffffff81194ac9>] out_of_memory+0x219/0x460 kernel: [<ffffffff8119aad5>] __alloc_pages_slowpath.constprop.88+0x965/0xb00 kernel: [<ffffffff8119aef6>] __alloc_pages_nodemask+0x286/0x2a0 kernel: [<ffffffff811e603d>] alloc_pages_vma+0xad/0x250 kernel: [<ffffffff811c3b2e>] handle_mm_fault+0x110e/0x1820 kernel: [<ffffffff8106b687>] __do_page_fault+0x197/0x400 kernel: [<ffffffff8106b912>] do_page_fault+0x22/0x30 kernel: [<ffffffff81849ac8>] page_fault+0x28/0x30 kernel: Out of memory: Kill process 3243 (hxmanager) score 56 or sacrifice child		
CSCvm59485	After a node is upgraded from release 2.5(1b) to 3.0(1c), the SED disks enter an ignored state.	When you see this issue, create the following file manually and reboot the system to claim the drives. **cat /etc/springpath/sed_capability.conf** **sed_capable_cluster=True**	3.0(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvm00138	In vCenter, the following warning appears for the ESXi host: `"Connectivity to Callhome endpoint failed."`	When you see this issue, manually clear the fault from vCenter. On each HyperFlex controller VM, complete these steps: open the following 1. In a text editor, open `/etc/cron.d/callhome`. 2. Comment the following two lines: 32 * * * * root sleep $(expr $RANDOM \% 90); /bin/asupcli ping --notify true 5 * * * * root sleep $(expr $RANDOM \% 300); /bin/asupcli post --type alert --event-name heartbeat	3.0(1c)
CSCvh04307	Installing software packages on the storage controller VM fails with the following error: “There are locked drives on the system, unlock them and retry deployment.” In addition, after upgrade from release 2.6(1e) to 3.0(1c), the following conditions are seen: • The upgrade is stuck on "checking cluster readiness" state for a long time. • The stcli cluster information shows the SED disks as unavailable, and hence the cluster cannot recover to a healthy state.	Contact Cisco TAC for more information to recover from this issue.	3.0(1c)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster.	Ensure there are sufficient replication IPs available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvg43082	During upgrade from 2.5(1c) to 2.6(1b), device connect might get disabled.	Enable device connect through HX Connect.	2.6(1a)
CSCvg05306	In rare cases, after kernel upgrade, the gateway IP address changes for eth0 and eth1.	Contact Cisco TAC for further guidance after collecting logs.	2.5(1a)
CSCvf98675	During the reboot phase of deployment server disks were not detected. HX deployment fails with error: "*Controller VM Install Device (tmpfs) not persistent. Cannot Install Packages*".	Re-acknowledge the server in UCS Manager. Then retry the deployment.	2.6(1a)
CSCvf93815	FI may reboot during cluster deployment, and installation hangs.	Reboot Installer VM and restart the HX cluster installation workflow.	2.6(1a)
CSCvf93812	Upgrade fails when the ESXi host is rebooted manually during the upgrade process. The rebooted ESXi host shows up as *'not-connected'/'not-responding'* in the VMware VCenter cluster. Subsequent upgrade attempts will fail during ESXi validation with authentication failure errors.	Manually recreate the HX user through vCenter. • Log in to ESXi command line. • Manually start vSphere HA (FDM) service. Reconnect the host to vCenter. Re-register this VM using vim command on the ESXi host. Please see VMWare KB: Cannot install the vSphere HA (FDM) agent on ESXi host (2007739) for more details. • Recover HX user. Contact Cisco TAC for assistance. • Add entry to */etc/hosts*. • Remount datastores. • Retry upgrade. Do not manually reboot any of the ESXi hosts in the cluster during the upgrade process.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)
CSCvb91838	Cluster expansion failed with no operational DNS server from the list.	If the DNS server becomes non-operational after deployment or cluster creation, add a new operational DNS to the controller. Use the following commands: **`# stcli services dns remove --dns <non_operational_dns_ip>`** **`# stcli services dns add --dns <operational_dns_ip>`**	1.8(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvb29790	Cluster creation fails due to failure to locate vCenter server.	In the vSphere Web Client, change the vCenter host name to an IP address in the `config.vpxd.sso.sts.uri` variable.	1.8(1a)
Management			
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the following: **1.** Delete the dummy stCtlVMs from the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1b)
CSCvi34303	HX Connect UI displays an error when any table is exported in .CSV format, and opened in excel.	None.	3.0(1a)
CSCvg69228	Previously deleted/removed disks may be erroneously listed in the inventory.	This is a benign issue that can be safely ignored.	2.6(1b)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvg26340	For releases higher than 2.5(1c), `node failures tolerable to be` shows zero while the cluster is healthy.	If you encounter this issue, contact TAC for a workaround.	2.5(1c)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
CSCvf27609	Query for recovery job returns `summary_step_state` and `state` fields.	Refer to the `state` field only. Ignore the information in the `summary_step_state` field.	2.5(1a)
Encryption			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh79736	Disable remote security times out and retry disable operation fails in HX Connect.	After the disable operation times out, issue rescan-inventory for the servers whose SAS controllers are still secure. Then retry disable encryption from the HX Connect. Run the following commands from any controller VM: **`hx.py --get-org getOrg.json --server-serial <serial number of server>`** **`hx.py --rescan-drive-inventory -f getOrg.json --server-serial <serial number of server>`**	3.0(1a)
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCSM might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)
CSCvf04240	Encryption may not be enabled on the new node, after it is added to the cluster.	One of the potential causes is that the serial number was not reported correctly from the ESX host. Restart `hostd` service on the ESX host, and enable encryption on the cluster from the HX Connect UI. All the nodes that already have encryption enabled are not impacted.	2.5(1a)
CSCve91866	Cannot modify encryption KMIP policy on UCSM to clear an IP address.	UCSM does not allow this behavior. On UCSM, delete the KMIP policy, adjusting for the IP addresses as needed, and retry the task.	2.5(1a)

Open Caveats in Release 3.0(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			
CSCvk18743	The Storage Controller VM is down for an extended period of time that may cause the VMs to power off.	In most cases, the Storage Controller VM comes back automatically within a couple of minutes. During the time the controller VM is down, the VM pauses temporarily, IO suspends and automatically resumes once the Controller VM is back up. Live migration is not required. If the controller VM does not come back in a couple of minutes, migrate the VMs from the affected HyperFlex Hyper-V host until the Storage controller VM can be powered up. Contact Cisco TAC for further assistance with resolution of this issue.	3.0(1b)
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)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi73801	The following roles are enabled by default on HX Hyper-V node: • Web Server • Remote access	If needed, manually disable roles.	3.0(1a)
CSCvi73796	Windows Event Log or Viewer displays errors due to invalid VMQ network configuration.	This is a known issue and does not impact functionality.	3.0(1a)
CSCvi70645	Directory listing on a host appears as hung for a non-existent datastore.	Use CTRL+C to cancel the operation and re-try.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)
CSCvi37407	Service profile association is stuck at "`Waiting for Storage Subsystem to initialize`" stage after re-acknowledging the server.	If you see this issue, do the following: • Disassociate Service profile. • Decommission the server. • Re-acknowledge the server and associate the service profile again.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)
CSCvi14568	Import VM operation fails using Hyper-V Manager (Remote) when the exported VM is stored on HX-Datastore.	If you see this issue, do the following: • Run the following powershell commands: `C:\Users\Administrator.HV-AD10> Invoke-Command -Comp stfs-028b -Command { Import-VM -Path "\\ctlvm28b.hv-ad1.local\ds1\export\vm7\Virtual Machines\[illegible].vmcx" -VhdDestinationPath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SnapshotFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SmartPagingFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -Copy -GenerateNewId | Rename-VM -NewName vm102 }` • This is a known Microsoft issue. See: Microsoft article: Remote import-vm in Hyper-V.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh80044, CSCvi59119	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.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
Install, Upgrade, Expand			
CSCvh51922	When using the LAN on motherboard (LOM) ports in shared LOM mode on HX M5 hardware, the link speed may not fully negotiate to 1000Mbps and may negotiate at the slower 100Mbps.	Manually set all switchport speeds to 1000 and do not rely on auto-negotiate. Reference the HX Edge install guide for sample configurations. Alternatively, set the CIMC to use dedicated mode.	2.6(1b)
CSCvh09129	Cluster Expansion: Validation (sufficient DR IP) should occur before adding the node to the cluster	Ensure there are sufficient replication IPs available for assignment to new nodes in the cluster. If necessary, modify the replication network configuration to include additional IP ranges.	2.6(1a)
CSCvg43082	During upgrade from 2.5.1c to 2.6.1b, device connect might get disabled.	Enable device connect through HX Connect.	2.6(1a)
CSCvg05306	In rare cases, after kernel upgrade, the gateway IP address changes for eth0 and eth1.	Contact Cisco TAC for further guidance after collecting logs.	2.5(1a)
CSCvf98675	During the reboot phase of deployment server disks were not detected. HX deployment fails with error: "*Controller VM Install Device (tmpfs) not persistent. Cannot Install Packages*".	Re-acknowledge the server in UCS Manager. Then retry the deployment.	2.6(1a)
CSCvf93815	FI may reboot during cluster deployment, and installation hangs.	Reboot Installer VM and restart the HX cluster installation workflow.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf93812	Upgrade fails when the ESXi host is rebooted manually during the upgrade process. The rebooted ESXi host shows up as *'not-connected'/'not-responding'* in the VMware VCenter cluster. Subsequent upgrade attempts will fail during ESXi validation with authentication failure errors.	Manually recreate the HX user through vCenter. • Log in to ESXi command line. • Manually start vSphere HA (FDM) service. Reconnect the host to vCenter. Re-register this VM using vim command on the ESXi host. Please see VMWare KB: Cannot install the vSphere HA (FDM) agent on ESXi host (2007739) for more details. • Recover HX user. Contact Cisco TAC for assistance. • Add entry to */etc/hosts*. • Remount datastores. • Retry upgrade. Do not manually reboot any of the ESXi hosts in the cluster during the upgrade process.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)

CSCvb91838	Cluster expansion failed with no operational DNS server from the list.	If the DNS server becomes non-operational after deployment or cluster creation, add a new operational DNS to the controller. Use the following commands:
		`# stcli services dns remove --dns <non_operational_dns_ip>`
		`# stcli services dns add --dns <operational_dns_ip>`
			1.8(1c)

CSCvb29790	Cluster creation fails due to failure to locate vCenter server.	In the vSphere Web Client, change the vCenter host name to an IP address in the `config.vpxd.sso.sts.uri` variable.	1.8(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
Management			
CSCvj31645	In rare cases, duplicate or dummy storage controller VMs (stCtlVMs) running windows appear in ESXi clusters.	If you see this issue, perform the following: **1.** Delete the dummy stCtlVMs from the vCenter. **2.** Cleanup the old extensions. **3.** Re-register to the original vCenter.	3.0(1b)
CSCvi34303	HX Connect UI displays an error when any table is exported in .CSV format, and opened in excel.	None.	3.0(1a)
CSCvg69228	Previously deleted/removed disks may be erroneously listed in the inventory.	This is a benign issue that can be safely ignored.	2.6(1b)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvg26340	For releases higher than 2.5(1c), `node failures tolerable to be` shows zero while the cluster is healthy.	If you encounter this issue, contact TAC for a workaround.	2.5(1c)
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
CSCvf27609	Query for recovery job returns `summary_step_state` and `state` fields.	Refer to the `state` field only. Ignore the information in the `summary_step_state` field.	2.5(1a)
Encryption			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh79736	Disable remote security times out and retry disable operation fails in HX Connect.	After the disable operation times out, issue rescan-inventory for the servers whose SAS controllers are still secure. Then retry disable encryption from the HX Connect. Run the following commands from any controller VM: **`hx.py --get-org getOrg.json --server-serial <serial number of server>`** **`hx.py --rescan-drive-inventory -f getOrg.json --server-serial <serial number of server>`**	3.0(1a)
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCSM might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)
CSCvf04240	Encryption may not be enabled on the new node, after it is added to the cluster.	One of the potential causes is that the serial number was not reported correctly from the ESX host. Restart `hostd` service on the ESX host, and enable encryption on the cluster from the HX Connect UI. All the nodes that already have encryption enabled are not impacted.	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCve91866	Cannot modify encryption KMIP policy on UCSM to clear an IP address.	UCSM does not allow this behavior. On UCSM, delete the KMIP policy, adjusting for the IP addresses as needed, and retry the task.	2.5(1a)

Open Caveats in Release 3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
Hyper-V			
CSCvi92070,	If you select multiple VMs in **Failover Cluster Manager** UI, only one VM becomes highly available.	Use the following PowerShell command instead of **Failover Cluster Manager**: `Get-VM -Name *print* | Add-ClusterVirtualMachineRole`	3.0(1a)
CSCvi73818	Windows Event Viewer shows an error, "RSS and VMQ - base processor and max processors overlap."	This issue can be safely ignored and has no impact on functionality. Use the `set -netadaptervmq` powershell command to update the processors used for VMQ.	3.0(1a)
CSCvi73807	Cloned user VM will not become a part of a failover cluster resource.	Use Failover Cluster Manager (FCM) to make it a clustered resource.	3.0(1a)
CSCvi73801	The following roles are enabled by default on HX Hyper-V node: • Web Server • Remote access	If needed, manually disable roles.	3.0(1a)
CSCvi73796	Windows Event Log or Viewer displays errors due to invalid VMQ network configuration.	This is a known issue and does not impact functionality.	3.0(1a)
CSCvi71157	Storage migration for a virtual machine fails in Hyper-V Manager with an error, "There was an error during move operation."	Storage migration within HX storage is not supported. However, Storage migration to or from a non-hx storage is supported.	3.0(1a)
CSCvi70645	Directory listing on a host appears as hung for a non-existent datastore.	Use CTRL+C to cancel the operation and re-try.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi63527	Datastore accessibility on the host is not validated before putting a host in maintenance mode.	Before entering the host in maintenance mode, ensure that all physical hosts have access to the datastore. If not, do not attempt maintenance mode.	3.0(1a)
CSCvi56910	Shortly after datastore creation, directory listing may fail with an error, " `The process cannot access the file \\xyz.cloud.local\\ds1 as it is being used in another process.`"	Wait for a minutes after datastore creation and retry the command.	3.0(1a)
CSCvi37407	Service profile association is stuck at "`Waiting for Storage Subsystem to initialize`" stage after re-acknowledging the server.	If you see this issue, do the following: • Disassociate Service profile. • Decommission the server. • Re-acknowledge the server and associate the service profile again.	3.0(1a)
CSCvi16323	You may not be able to navigate to a specific HX-Datastore in Hyper-V Manager (Remote) as the HX-Datastore is grayed-out, and the **Inspect Disk** option is unavailable.	This is a known issue.	3.0(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvi14568	Import VM operation fails using Hyper-V Manager (Remote) when the exported VM is stored on HX-Datastore.	If you see this issue, do the following: • Run the following powershell commands: `C:\Users\Administrator.HV-AD10> Invoke-Command -Comp stfs-028b -Command { Import-VM -Path "\\ctlvm28b.hv-ad1.local\ds1\export\vm\Virtual Machines\[illegible].vmcx" -VhdDestinationPath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SnapshotFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -SmartPagingFilePath "\\ctlvm28b.hv-ad1.local\ds1\i14" -Copy -GenerateNewId | Rename-VM -NewName vm102 }` • This is a known Microsoft issue. See: Microsoft article: Remote import-vm in Hyper-V.	3.0(1a)
CSCvh99126	Creating virtual machine from template works, but without template workflow from VMM.	If you see this issue, do the following: • Resolve SMB access point to the cluster management IP using etc/hosts file on the VMM console host. • Execute the following command on all controller VMs before adding the HX Share to VMM: `iptables -I INPUT 1 -i eth0 -p tcp --dport 445 -j ACCEPT` The command opens port 445 on management networ	3.0(1a)
CSCvh80044, CSCvi59119	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.	This is a known limitation. Case-sensitivity is currently not supported, and will be addressed in future releases.	3.0(1a)
Install, Upgrade, Expand			

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvj18863	During or after an upgrade to release 3.0(1a), OR a subsequent upgrade from 3.0(1a), the cluster may experience one or more nodes restarting due to a SW panic.	Contact Cisco TAC for further assistance with this issue. In addition, note the following: • If you have already upgraded to 3.0(1a) from a previous release, contact Cisco TAC to proactively address this exposure. • This is not an issue for a fresh install of 3.0(1a).	3.0(1a)
CSCvi84992	During offline upgrade, HX Connect displays the following a "`Server Call Failure`" error.	If you see this issue, do the following: 1. Log out from HX Connect, and log back in. Go to the **Upgrade** page. 2. Execute the following command: `stcli cluster start` 3. Refresh HX Connect to see the updated information.	3.0(1a)
CSCvh54563	Generating HyperFlex logs on HX240 platform causes an all paths down state for the cluster. This issue is caused by the vmw_ahci driver in ESXi 6.5 GA release.	If this issue occurs, do one of the following: • (Recommended) Upgrade to VMware vSphere 6.5 patch U1. • Disable the vmw_ahci driver on ESXi host one node at a time, making sure that the cluster is healthy before moving to the next node. Use the following steps: 1. Run the following command on the ESXi nodes: `# esxcli system module set --enabled=false --module=vmw_ahci` • Reboot the node.	2.6(1b)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvh51922	When using the LAN on motherboard (LOM) ports in shared LOM mode on HX M5 hardware, the link speed may not fully negotiate to 1000Mbps and may negotiate at the slower 100Mbps.	Manually set all switchport speeds to 1000 and do not rely on auto-negotiate. Reference the HX Edge install guide for sample configurations. Alternatively, set the CIMC to use dedicated mode.	2.6(1b)
CSCvg41818	HX snapshot error when taking snapshot with GPU shared PCI device attached for powered off VM.	You cannot create a snapshot for VMs with GPUs. These VMs must be powered off for taking a snapshot.	2.5(1c)
CSCvg43082	During upgrade from 2.5.1c to 2.6.1b, device connect might get disabled.	Enable device connect through HX Connect.	2.6(1a)
CSCvg05306	In rare cases, after kernel upgrade, the gateway IP address changes for eth0 and eth1.	Contact Cisco TAC for further guidance after collecting logs.	2.5(1a)
CSCvf98675	During the reboot phase of deployment server disks were not detected. HX deployment fails with error: "*Controller VM Install Device (tmpfs) not persistent. Cannot Install Packages*".	Re-acknowledge the server in UCS Manager. Then retry the deployment.	2.6(1a)
CSCvf93815	FI may reboot during cluster deployment, and installation hangs.	Reboot Installer VM and restart the HX cluster installation workflow.	2.6(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf93812	Upgrade fails when the ESXi host is rebooted manually during the upgrade process. The rebooted ESXi host shows up as *'not-connected'/'not-responding'* in the VMware VCenter cluster. Subsequent upgrade attempts will fail during ESXi validation with authentication failure errors.	Manually recreate the HX user through vCenter. • Log in to ESXi command line. • Manually start vSphere HA (FDM) service. Reconnect the host to vCenter. Re-register this VM using vim command on the ESXi host. Please see VMWare KB: Cannot install the vSphere HA (FDM) agent on ESXi host (2007739) for more details. • Recover HX user. Contact Cisco TAC for assistance. • Add entry to */etc/hosts*. • Remount datastores. • Retry upgrade. Do not manually reboot any of the ESXi hosts in the cluster during the upgrade process.	2.6(1a)
CSCve73004	UCS Manager does not update the disk firmware status, if a firmware upgrade from 2.1(1b) to 2.5 was initiated by the HX Data Platform.	Perform a soft reset: # `CIMC-soft-rest`	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvc62266	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".	Manually power on the controller VM and start the cluster. **1.** Manually power on the controller VMs. • Log in to the vSphere Web Client. • Locate the controller VMs that are not powered on. From the vCenter Navigator select, **Inventory Lists** > **Virtual Machines** > **vm**. Storage controller VMs have the prefix, `stCtlVM`. • From the Actions menu, select **Power** > **Power On**. **2.** Restart the storage cluster. • Log in to the command line of any controller VM. • Run the command: `# stcli cluster start`	2.0(1a)
CSCvb94112	HX Installer may be stuck at Cluster Expansion Validation screen during the cluster expansion process.	**1.** Check logs to verify that the expansion workflow is hung. **2.** In your browser, type *http://ip_of_installer/api/reset* to restart the workflow.	1.8(1c)

CSCvb91838	Cluster expansion failed with no operational DNS server from the list.	If the DNS server becomes non-operational after deployment or cluster creation, add a new operational DNS to the controller. Use the following commands:
		`# stcli services dns remove --dns <non_operational_dns_ip>`
		`# stcli services dns add --dns <operational_dns_ip>`
			1.8(1c)

CSCvb29790	Cluster creation fails due to failure to locate vCenter server.	In the vSphere Web Client, change the vCenter host name to an IP address in the `config.vpxd.sso.sts.uri` variable.	1.8(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
Management			
CSCvi95613	HX Installer fails when HX LFF servers are on the same UCS domain.	If deploying an LFF HX cluster in a shared UCS domain with other HX clusters, you must use upgrade to 3.0(1a) before attempting expansion. New non-LFF cluster deployments on the same domain as an LFF HX cluster must also use the 3.0(1a) or later HX Installer.	3.0(1a)
CSCvi92768	Adding two nodes via stcli to LAZ cluster results in an error.	Due to a known issue in the current release, multiple node can lead to cluster unavailability. This bug will be addressed in the next release.	3.0(1a)
CSCvi82755	The HyperFlex Clusters list view in Cisco Intersight might not display the Hypervisor version for the ESXi cluster.	From Cisco Intersight, cross launch into HX Connect to see the running ESXi version.	3.0(1a)
CSCvi76137	When a node is in maintenance mode, disk API is stuck.	Bring the node back up.	3.0(1a)
CSCvi73417	In HX Connect, while a snapshot is in progress, VMs list page appears as blank.	None.	3.0(1a)
CSCvi34303	HX Connect UI displays an error when any table is exported in .CSV format, and opened in excel.	None.	3.0(1a)
CSCvg69228	Previously deleted/removed disks may be erroneously listed in the inventory.	This is a benign issue that can be safely ignored.	2.6(1b)
CSCvg47332	Using the the quiesce option for a VM with a HX snapshot may cause the VM to be stunned.	If you plan to use the quiesce option, do not use it for a VM that has a HX snapshot. If you need to use the quiesce option, delete all HX snapshots and use VMware snapshots.	2.1(1b)
CSCvg26340	For releases higher than 2.5(1c), `node failures tolerable to be` shows zero while the cluster is healthy.	If you encounter this issue, contact TAC for a workaround.	2.5(1c)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf90091	When incorrect gateway is provided, errors are seen in a cluster, after cluster creation.	Log in to the controller VM and correct the gateway.	2.5(1c)
CSCvf25130	HX Connect times out after 30 minutes.	When left idle for more than 30 minutes, the HX Connect Virtual Machine page times out. When you return to a page and click anywhere, refreshed data might be incomplete or you might receive the following error: *VI SDK invoke exception: nested exception is:* `com.vmware.vim25. Not Authenticated`. Retry refresh HX Connect through the browser or HX Connect buttons. Alternatively, log out of HX Connect and log back in. This is a known VMware issue. See VMware KB, vCenter Server logs report the error: SOAP session count limit reached (2004663).	2.5(1a)
CSCve17284	Performance charts show a gap for several minutes during an All Flash cluster upgrade.	This is expected behavior because the reporting services are taken down during the upgrade. Only the Reporting Chart is affected, not the actual performance.	2.5(1a)
Replication			
CSCvf29202	Recovery might not include disks that are not in the same folder on a datastore as the virtual machine being protected.	If any virtual machine disk resides outside the same folder and datastore of a protected virtual machine: 1. Move the disk to the same folder on the datastore. 2. Then add (re-add) the disk to the virtual machine. This ensures protection and recovery work successfully.	2.5(1a)
CSCvf27609	Query for recovery job returns `summary_step_state` and `state` fields.	Refer to the `state` field only. Ignore the information in the `summary_step_state` field.	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvb54848	vSphere Replication Plug-in fails after HX Plug-in is deployed.	To prevent the issue, first install the vSphere Replication plug-in, and then install the HX Data Platform plug-in. For complete steps for uninstalling required elements and reinstalling them in the supported order, see the 2.5 Release Troubleshooting guide.	1.7.1
Encryption			
CSCvh79736	Disable remote security times out and retry disable operation fails in HX Connect.	After the disable operation times out, issue rescan-inventory for the servers whose SAS controllers are still secure. Then retry disable encryption from the HX Connect. Run the following commands from any controller VM: **`hx.py --get-org getOrg.json --server-serial <serial number of server>`** **`hx.py --rescan-drive-inventory -f getOrg.json --server-serial <serial number of server>`**	3.0(1a)
CSCvf17183	If a CIMC reboots while a `modify-security` command is in-progress, and the server is secured with local key management, a subsequent `disable-security` command may fail, because the server doesn't know the correct key to use.	CIMC was rebooted while a `modify-security` command was in-progress. Log in to the controller VM and use the sed-client to update the physical drive keys to match the server's key.	2.5(1a)
CSCvf06510	UCSM might indicate partially disabled encryption security.	No action required. This is a sync issue between reporting interfaces. To verify from HX Connect, select **System Information** > **Disks** > **Security**. All disks and the controller VM should indicate *Security Disabled*.	2.5(1a)

Defect ID	Symptom	Workaround	Defect Found in Release
CSCvf04240	Encryption may not be enabled on the new node, after it is added to the cluster.	One of the potential causes is that the serial number was not reported correctly from the ESX host. Restart `hostd` service on the ESX host, and enable encryption on the cluster from the HX Connect UI. All the nodes that already have encryption enabled are not impacted.	2.5(1a)
CSCve91866	Cannot modify encryption KMIP policy on UCSM to clear an IP address.	UCSM does not allow this behavior. On UCSM, delete the KMIP policy, adjusting for the IP addresses as needed, and retry the task.	2.5(1a)

HX Data Platform Preinstallation Checklist, Installation and Configuration, Administration, Upgrade, CLI Reference, and Troubleshooting Guides

Document/Description	Quick Links
Preinstallation Checklist for VMware ESXi	
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.	Preinstallation Checklist for VMware ESXi
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.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Upgrade Guides for VMware ESXi	
Provides information on how to upgrade an existing installation of Cisco HX Data Platform, upgrade guidelines, and information about various upgrade tasks.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases	

Document/Description	Quick Links
Guides Cisco HyperFlex users who need to upgrade their 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.	Cisco HyperFlex Systems Upgrade Guide for Unsupported Cisco HX Releases
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.	Cisco HyperFlex Release 5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Preinstallation Checklist for Cisco HyperFlex Edge	
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.	Preinstallation Checklist for Cisco HyperFlex Edge
Edge Deployment Guide	
Provides deployment procedures for HyperFlex Edge, designed to bring hyperconvergence to remote and branch office (ROBO) and edge environments.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Network and External Storage Management Guide	
Provides information about HyperFlex Systems specific network and external storage management tasks.	Network and External Storage Management Guide
Installation Guide on Microsoft Hyper-V	
Provides installation and configuration procedures on how to install and configure Cisco HyperFlex Systems on Microsoft Hyper-V.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Cisco HyperFlex Upgrade Guide for Microsoft Hyper-V	

Document/Description	Quick Links
Provides information on how to upgrade an existing installation of Cisco HX Data Platform, upgrade guidelines, and information about various upgrade tasks.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
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 HX Connect, HX Data Platform Plug-in, and the `hxcli` commands.	Cisco HyperFlex Release 5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
Cisco HyperFlex Systems Network and External Storage Management Guide for Microsoft Hyper-V	
Overview of the network and external storage architecture for Cisco HyperFlex Systems.	Cisco HyperFlex Systems Network and External Storage Management Guide for Microsoft Hyper-V
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.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
Kubernetes Integration	
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.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
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.	5.0
	Cisco HyperFlex Release 4.5x
	Cisco HyperFlex Release 4.0x
	Cisco HyperFlex Release 3.5x
HyperFlex Intersight Installation Guide	

Document/Description	Quick Links
Provides installation, configuration, and deployment procedures for HyperFlex Intersight, designed to deliver secure infrastructure management anywhere from the cloud.	HyperFlex Intersight Installation Guide
Cisco HyperFlex SD-WAN Deployment Guide	
Feature preview for deploying the SD-WAN solution on a HyperFlex cluster. Cisco recommends that you test this feature on a test network/system (Not for use in your production environment).	Cisco HyperFlex SD-WAN Deployment Guide
Cisco HX Data Platform Security Hardening Guide	
Provides recommended configuration settings and deployment architectures for HXDP-based solutions.	Cisco HX Data Platform Security Hardening Guide
Provides additional vCenter Security Hardening settings.	How to Configure vCenter Security Hardening Settings
Tech Notes	
Provides information on recommended FI/Server firmware.	TechNotes
Troubleshooting Guide	
Provides troubleshooting for installation, configuration, Cisco UCS Manager to Cisco HyperFlex configuration, and VMware vSphere to HyperFlex configuration. In addition, this guide provides information about understanding system events, errors, Smart Call Home, and Cisco support.	Troubleshooting Guide
Command Line Interface (CLI) Guide	
Provides CLI reference information for HX Data Platform `stcli` commands.	Command Line Interface (CLI) Guide
Rest API Guides	
Provides information related to REST APIs that enable external applications to interface directly with the Cisco HyperFlex management plane.	REST API Getting Started Guide REST API Reference
Cisco HyperFlex PowerShell Cmdlets for Disaster Recovery	
Provides information on how to use the Cisco PowerShell Cisco HXPowerCLI cmdlets for Data Protection.	Cisco HyperFlex PowerShell Cmdlets for Disaster Recovery
Cisco HxBench Getting Started Guide	
This document describes how to use the Cisco HxBench storage performance testing tool to measure the storage infrastructure.	Cisco HxBench Getting Started Guide