

Cisco HyperFlex Software Advisory for HX Release 3.5(2d)

Cisco engineering has identified an issue with Cisco HyperFlex Release 3.5(2a) release that may affect your use of this software. Please review the Software Advisory notice here to determine if the issues apply to your environment and the steps required to address the issue.

Contents:

CDETS	Area of Impact	Date
CSCvw01432	Moving Existing HyperFlex ESXi Hosts to vCenter 7.0 Update 1 Can Unexpectedly Power Off and Delete the HyperFlex Controller VMs	Updated: January 8, 2021
CSCvq53058	Stretched Cluster Witness VM.	July 13, 2020
CSCvq18919	Better Split-Brain Handling and Failover	July 13, 2020
CSCvq17778	Better Split-Brain Handling and Failover	July 13, 2020
CSCvq89852	HXConnect	July 13, 2020
CSCvq06952	Snapshot Creation on Change Block Tracking (CBT)	February 27, 2020
CSCvp88515	Snapshot Creation on Change Block Tracking (CBT)	February 27, 2020
CSCvk62990	ESXi 6.0 Hosts Susceptible to PSOD	Update: January 14, 2019

Software Advisory for CSCvw01432: Moving Existing HyperFlex ESXi Hosts to vCenter 7.0 Update 1 Can Unexpectedly Power Off and Delete the HyperFlex Controller VMs

Updated: January 8, 2021

Dear Cisco Customer,

Cisco engineering has identified the following software issues with the release that you have selected that may affect your use of this software. Please review the Software Advisory notice here to determine if the issues apply to your environment. You may proceed to download this software if you have no concerns with the issue described.

For more comprehensive information about what is included in this software, refer to the Cisco software Release Notes, available from the [Product Selector tool](#). From this page, select the product you are interested in. Release Notes are under "General Information" on the product page.

Reason for Advisory:

An interoperability issue between HyperFlex and vCenter 7.0 Update 1 (U1) has been identified that impacts HyperFlex controller VMs that are managed by vSphere ESX Agent Manager (EAM). At the time of initial publishing of this notice, vCenter and ESXi 7.0 and 7.0 U1 were not supported with HyperFlex, and this notice was issued out of an abundance of caution. HyperFlex Data Platform software versions 4.0(2) & 4.5(1) and later now support vCenter Server 7.0 U1c (build 17327517) and higher.

HyperFlex Data Platform software adheres to a list of compatible ESXi and vCenter versions in all published software release notes. vCenter Server 7.0 U1c may be used with HyperFlex 4.0(2) and 4.5(1) and later releases. No vCenter Server 7 version prior to U1c may be used at any time. Refer to the [HyperFlex 4.0 Release Notes, table 6](#) & [Cisco HX Release 4.5\(x\) – Software Requirements, table 7](#) for full details.

Note: ESXi & vCenter compatibility with the HyperFlex Data Platform are independently qualified and listed separately in the product documentation. This field notice strictly covers vCenter Server 7.0 interoperability and does not address ESXi 7.0 interoperability.

vCenter Server 7.0 behavior prior to U1c:

Some HyperFlex clusters use a service in vCenter Server that is called EAM. This service is responsible for the lifecycle of the HyperFlex controller VMs (named stCtlVM in vCenter).

When existing HyperFlex ESXi hosts are moved from an existing vCenter Server to a new vCenter Server 7.0 U1, the HyperFlex controller VMs that reside on these hosts will be powered off and deleted. This issue impacts HyperFlex clusters that were initially installed with a HyperFlex Data Platform (HXDP) version earlier than 4.0(1a). By default, all clusters that were deployed initially on a version pre-4.0(1a) will continue to use the EAM service, even after an upgrade to version 4.0(1a) of HXDP or later.

Example scenarios are listed here:

- HyperFlex cluster initially deployed on 2.6(1a) and then upgraded to 4.0(2c).
 - Although the cluster is upgraded to a version later than 4.0(1a), it was initially deployed on a version pre-4.0(1a) and is therefore susceptible to this issue.
- HyperFlex cluster initially deployed on 4.0(1b).
 - This cluster is not impacted because it was initially deployed on version 4.0(1a) or later.
- HyperFlex cluster initially deployed on 3.5(2a), then upgraded to 4.0(2c), and then manually modified in order to remove the EAM configuration in accordance with Cisco documentation.
 - This cluster is still susceptible. The manual removal procedure does not resolve the problem.

If a susceptible cluster is removed from the current vCenter and the corresponding ESXi hosts are added to a new vCenter 7.0 U1 instance, the controller VMs will be permanently deleted. As a result, the cluster becomes unavailable and, in some cases, the cluster data becomes unrecoverable.

Affected Hardware Platforms:

All HyperFlex converged nodes except for new hardware that requires HXDP version 4.0(1a) or later.

Symptom:

HyperFlex Controller VMs can power off suddenly and be deleted from disk by the EAM service in vCenter when adding the ESXi hosts that host the controller VMs to a vCenter prior to 7.0 U1c. This results in a loss of cluster availability. In some cases, the HyperFlex storage cluster can become unrecoverable.

Workaround:

VMware has enhanced the default EAM behavior in vCenter Server 7.0 U1c and later to prevent orphaned VM cleanup automatically for non-vCLS VMs. Fresh and upgraded vCenter Server installations will no longer encounter an interoperability issue with HyperFlex Data Platform controller VMs when running vCenter Server 7.0 U1c and later.

When used with HyperFlex, vCenter Server should never be manually configured to auto-clean up all orphaned VMs automatically. For further detail, refer to VMware KB 81352 available at <https://kb.vmware.com/s/article/81352>.

vCenter Server 7.0 prior to the U1c release should never be used with Cisco HyperFlex.

vCenter Server	EAM Behavior	HyperFlex Support Stance
All 6.x releases	No EAM Interop Issue.	Supported per HyperFlex release notes
7.0 (15952498)	No EAM Interop Issue. Unqualified by Cisco.	Unsupported
7.0a (16189094)	No EAM Interop Issue. Unqualified by Cisco.	Unsupported
7.0b (16386292)	No EAM Interop Issue. Unqualified by Cisco.	Unsupported
7.0c (16620007)	No EAM Interop Issue. Unqualified by Cisco.	Unsupported
7.0d (16749653)	No EAM Interop Issue. Unqualified by Cisco.	Unsupported
7.0 U1 (16860138)	EAM cleanup results in interoperability issue with HyperFlex	Unsupported
7.0 U1a (17004997)	EAM cleanup results in interoperability issue with HyperFlex	Unsupported
7.0 U1c (17327517) and later	Issue Resolved	Supported per HyperFlex release notes

Note: Manual removal of the EAM configuration in accordance with the documented procedure on Cisco.com will not prevent this issue.

More Information:

See Cisco Field Notice [FN70620](#).

If you require further assistance, or if you have any further questions regarding this field notice, please contact the Cisco Systems [Technical Assistance Center \(TAC\)](#) by one of the following methods:

- [Open a service request on Cisco.com](#)
- [By email or telephone](#)



Software Advisory for Cisco HyperFlex Stretched Cluster Operations

July 13, 2020

Dear Cisco Customer,

Cisco engineering has identified issues with Stretched Cluster configurations with HX 3.5(2d) release that may affect your use of this software. Please review the Software Advisory notice here to determine if the issues apply to your environment and the steps required to address the issue.

For more comprehensive information about what is included in this software, refer to the Cisco software Release Notes, available from the [Product Selector tool](#). From this page, select the product you are interested in. Release Notes are under "General Information" on the product page.

Affected Software and Replacement Solution (for issues listed above)		
Software Type	Software Affected	Software Solution
HX Data Platform	Version: HX Data Platform 3.5(2d) in a Stretched Cluster configuration	Version: Recommended release for 3.5(2x). ¹

Reason for Advisory:

This software advisory is a notification for issues found in HXDP 3.5(2d) that impacts all stretched cluster configurations. For this reason, HXDP 3.5(2d) should not be used for stretched cluster configurations.

Stretched Cluster Configuration Users:

If you have installed or already upgraded a Stretched Cluster configuration to HXDP 3.5(2d) please reinstall or upgrade to HXDP Recommended release for 3.5(2x). See the [Recommended Cisco HyperFlex HX Data Platform Software Releases](#) for current recommendations.

Non-Stretched Cluster Configuration Users:

HyperFlex Standard HX clusters are not affected by these defects.

¹ Cisco HX Software Recommendations are found in the [Recommended Cisco HyperFlex HX Data Platform Software Releases](#) Guide.

Affected Hardware Platforms:

HX220C-M5SX
HX240C-M5SX
HXAF220C-M5SX
HXAF240C-M5SX
HX240C-M5L

Symptoms & Conditions:

To be susceptible to the issues listed, the HyperFlex cluster must be running HXDP 3.5(2d) and configured in a Stretched Cluster configuration. HyperFlex Standard HX clusters are not affected by this defect.

A. Stretch Cluster Witness Related

CSCvq53058	When the Witness VM has a high RTT times (>50ms) from any of the two Stretched Cluster sites, there is a possibility under heavy transaction load that the cluster failover or failback times may be very high.
------------	---

B. Better Split-Brain Handling and Failover

CSCvq18919	During Stretched Cluster failover, zookeeper goes offline due to file write error.
CSCvq17778	After back-to-back failovers on one site, the cluster encountered a failure that lead to zookeeper going offline.

C. HX Connect

CSCvq89852	HX Connect reports Stretched Cluster as standard cluster after transition from offline to online state.
------------	---

Workaround: If a Stretched Cluster configuration has been installed or upgraded to HXDP 3.5(2d) (The affected release) - please upgrade to the HXDP Recommended release for 3.5(2x). See the "[Recommended Cisco HyperFlex HX Data Platform Software Releases](#)" for the current recommendations.

Software Advisory for CSCvq06952 and CSCvp88515: Snapshot Creation on Change Block Tracking (CBT) Enabled Fails

February 27, 2020

Dear Cisco Customer,

Cisco engineering has identified an incompatibility issue with the Cisco HyperFlex Native Snapshot API and VMware ESXi 6.7 U2 and U3. This combination may affect your use of this software. Please review the Software Advisory notice here to determine if the issues apply to your environment, and the steps required to address the issue.

For more comprehensive information about what is included in this software, refer to the Cisco software Release Notes, available from the Product Selector tool. From this page, select the product you are interested in. Release Notes are under "General Information" on the product page.

Affected Software and Replacement Solution for CSCvq06952, CSCvp88515		
Software Type	Software Affected	Software Solution
Cisco HyperFlex Data Platform VMware ESXi 6.7 U2 and U3 Hypervisor	Version: HXDP 3.5(2c), HXDP 3.5(2d) HXDP 3.5(2e) HXDP 4.0(1a) version running the ESXi 6.7 U2 hypervisor Affected Images (iso and zip bundles): ESXi 6.7 U2 and U3	Version: HX 3.5 release: HXDP 3.5(2g) HX 4.0 release: HXDP 4.0(2a) Replacement Images: None

Reason for Advisory:

This software advisory is a notification about an incompatibility found when using Cisco HyperFlex Data Platform Snapshot API and VMware ESXi 6.7 U2 and U3 Hypervisor. This combination is causing the snapshot creation on Change Block Tracking (CBT) enabled VM to fails with the error "Failed in vmreparent vmkfstools clone1"

Solution:

Users running Cisco HyperFlex 3.5 release should upgrade to 3.5(2g) to resolve this issue.

Users running Cisco HyperFlex 4.0 release should upgrade to 4.0(2a) to resolve this issue.

Affected Hardware Platforms:

HX220C-M5

HX220C-M5SX

HX220C-M5 Edge

HX240C-M5

HX240C-M5SX

HX220C-AF M5

HX240C-AF M5

HX240C-M5L

HXAF220c-M5SN

HXAF220C-M5SX

HXAF240C-M5SX

HX220C All NVMe M5

Symptom:

Due to a technical incompatibility between Cisco HyperFlex Data Platform's Native Snapshot APIs and the VMWare ESXi 6.7 U2 and U3 hypervisor, when Cisco HyperFlex Native Snapshot API is invoked for VMs that have Change Block Tracking (CBT) enabled, snapshots fail. After the snapshot failure, the backup process that initially called the native snapshot API will also fail.

Conditions:

The issue is limited to users with the following combination of Cisco HyperFlex Data Platform versions and VMWare ESXi hypervisor versions

- Clusters running one of the Cisco HyperFlex Data Platform versions listed above when a HyperFlex Native Snapshot is attempted for VMs that have CBT enabled.
- Cisco HyperFlex cluster that have upgraded to or fresh installed with the hypervisor version VMWare ESXi 6.7 U2 and U3.

Workaround: None. See Solution.

More Info: None.

Software Advisory for CSCvk62990: ESXi 6.0 Hosts Susceptible to PSOD

Updated: January 14, 2019

Dear Cisco Customer,

Cisco engineering has identified the following software issues with the release that you have selected that may affect your use of this software. Please review the Software Advisory notice here to determine if the issues apply to your environment. You may proceed to download this software if you have no concerns with the issue described.

For more comprehensive information about what is included in this software, refer to the Cisco software Release Notes, available from the [Product Selector tool](#). From this page, select the product you are interested in. Release Notes are under "General Information" on the product page.

Affected Software and Replacement Solution for CSCvk62990		
Software Type	Software Affected	Software Solution
VMware ESXi 6.0 Hypervisor	<p>Version: HX Data Platform 2.6, 3.0, 3.5 running an ESXi 6.0 hypervisor version listed below or any other VMware 6.0 build:</p> <p>Affected Images (iso and zip bundles): <i>ESXi 6.0.U3patch1 (build 5572656)</i> <i>ESXi 6.0.U3patch6 (build 6921384)</i> <i>ESXi 6.0.U3e (build 7967664)</i></p>	<p>Version: Patched VMware ESXi 6.0 hypervisor listed below:</p> <p>Replacement Images:</p> <p><i>Cisco HX Custom Image for ESXi 6.0 EP19 Offline Bundle for Upgrading from prior ESXi versions (build 10719132)</i></p> <p><i>Cisco HX Custom Image for ESXi 6.0 U3 EP18 Offline Bundle for Upgrading from prior ESXi versions (build 10474991)</i></p> <p><i>Cisco HX Custom Image for ESXi 6.0 U3-9919195 Offline Bundle for Upgrading from prior ESXi versions (build 9919195)</i></p>

Reason for Advisory:

This software advisory addresses a software driver issue found in the VMware ESXi 6.0 hypervisor.

CSCvk62990 - VMware ESXi 6.0 Hosts Susceptible to PSOD Triggered by xhci Driver Race Condition**Affected Hardware Platforms:**

HX220C-M5SX
HX240C-M5SX
HXAF220C-M5SX
HXAF240C-M5SX
HX240C-M5L

Note: This software issue does not affect HyperFlex M4 generation hardware.

Symptom:

A software bug exists in the VMware xhci-xhci vmklinux driver that may cause an ESXi host to crash (PSOD) due to a null pointer exception in the driver triggered by a race condition. This VMware bug has been observed sparingly during installation and upgrade activity on HyperFlex clusters running M5 hardware with the ESXi 6.0 hypervisor. All previously posted Cisco customized ESXi 6.0 builds are susceptible, including the latest patch builds available directly from VMware.

Conditions:

To be susceptible, the HyperFlex cluster must be running an M5 hardware platform as listed above and must also be running ESXi 6.0. All HyperFlex M4 ESXi 6.0 & 6.5 clusters, M4 & M5 ESXi 6.5 clusters, and M5 Hyper-V 2016 clusters are unaffected.

Workaround:

The issue is acknowledged by VMware and a fix has been tested, released, and posted on Cisco.com for ESXi 6.0. The fixed build can be found in the HyperFlex Downloads section on Cisco.com as ESXi 6.0 U3-9919195. Customers may opt to upgrade to this patch release or decide to upgrade to ESXi 6.5, that uses a different native driver framework that is not susceptible to this defect. Customers are encouraged to upgrade ESXi first before upgrading to a newer HyperFlex release. Upgrades should be performed using the esxcli software profile commands as documented in the [upgrade guide](#). Other ESXi upgrade mechanisms, including using vSphere update manager (VUM) or ISO driven upgrades, should not be used on HyperFlex servers.

More Information:

The vulnerable xhci-xhci driver version can be identified by running “esxcli software vib list” on any ESXi shell. Any version of xhci-xhci older than 1.0-3vmw.600.3.104.10280114 does not contain the fix. Please note the version numbers are similar and the underlined portion should be compared.

If the issue is encountered, it may be identified by a PF Exception 14 PSOD, followed by mention of xhci_endpoint in the stack trace. If this condition is met, a reboot of the host will bring it back into service without any additional user intervention required. VM workload disruption is possible, although Cisco has seen the issue only occur on fresh install prior to workload commission and during upgrade only when all VMs have been live migrated to other hosts as part of the normal upgrade process. Out of an abundance of caution, customers that must remain on ESXi 6.0 are encouraged to upgrade to the ESXi 6.0 U3-9919195 release prior to upgrading HyperFlex clusters.

This fix will become generally available in the next bugfix release from VMware. However, it is always recommended to download the latest Cisco customized ESXi zip bundles for upgrading HyperFlex clusters directly from Cisco.com.

Note:

VMware build 9919195 is a tracking build number provided by VMware that contains additional fixes to misc-drivers and sata-ahci VIBs. The esx-base version in this build remains at 9313334. In all VMware management interfaces, UI, CLI, HX Connect, and Cisco Intersight, the build number will be displayed as 9313334 and not 9919195.
