



## **Cisco HyperFlex Systems Troubleshooting Reference Guide, 5.0**

**First Published:** 2021-11-10

**Last Modified:** 2023-01-31

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021 –2023 Cisco Systems, Inc. All rights reserved.



## CONTENTS

[Full Cisco Trademarks with Software License](#) ?

---

### PREFACE

[Communications, Services, Bias-free Language, and Additional Information](#) ix

---

### CHAPTER 1

[New and Changed Information](#) 1

[New and Changed Information](#) 1

---

### CHAPTER 2

[HX Data Platform Troubleshooting](#) 3

[Troubleshooting HyperFlex Systems Overview](#) 3

[Contacting Cisco TAC](#) 4

---

### CHAPTER 3

[HX Data Platform Support](#) 5

[Using TAC Support for Troubleshooting](#) 5

[Common TAC-Assisted Support Topics](#) 5

[Auto Support and Smart Call Home for HyperFlex](#) 6

[Configuring Auto Support Using HX Connect](#) 7

[Configuring Notification Settings Using CLI](#) 8

[Configuring Smart Call Home for Data Collection](#) 9

[Cisco HX Data Platform Support Bundles](#) 11

[Generating a Support Bundle Using HX Connect](#) 12

[Collecting Basic Support Bundles from Controller VMs](#) 14

[Collecting Logs from Controller VMs: HX Release 4.0\(2a\)](#) 15

[Collecting Logs from Controller VMs](#) 16

[Generating an Audit Log](#) 16

[Collecting Logs from ESXi Hosts](#) 18

[Out of Space Error When Generating a Support Bundle](#) 18

Collecting Deployment Logs from the Cisco HX Data Platform Installer VM	18
Collecting Logs from the vCenter Server	19
Collecting Logs from the vSphere Web Client	19
Collecting Logs from the Cisco HXDP Plug-in	20
Downloading an Existing Support Bundle	20
Uploading Support Bundles	21
Uploading Support Bundles using ftp or sftp	21
Uploading Support Bundles using https	21

**CHAPTER 4****HX Data Platform Events 23**

Understanding System Events	23
View HX Data Platform Plug-in Events	24
Automatically Acknowledged Events	25
Manually Acknowledging Alarms Using HX Connect	26
Cluster Access Policy Compliance Events	26
Cluster Events	27
Critical Infrastructure Events	31
Disk Events	31
Host Events	34
Node Events	35
Encryption Events	37
Replication and Recovery Events	45
Security Events	47
Resource Monitoring Events and Alarms	48
Smart Call Home Events	53
Smart License Events	53
Snapshot Events	55
Space Utilization Events	56
Storage Events	59
Upgrade Events	59

**CHAPTER 5****Troubleshooting Topics 61**

Install and Upgrade Issues	61
Deploy IP Addresses Page Lists Duplicate Servers	61

Installation Fails when Manually Reboot FIs	61
During UCS Manager Only Upgrade, Controller VM Might Not Power On	62
Deploy or Upgrade Fail with Error: "NoneType" object has no attribute 'scsiLun'	62
Upgrade Fails to Enter Maintenance Mode	62
Upgrade Fails at vMotion Compatability Validation	63
Upgrade VM Power On Error: No compatible host was found	63
During Upgrade when two Nodes Fail, Controller VMs Power On Fail	63
Upgrade with Pre-6.5 vCenter Groups Some Controller VMs	64
A Node Fails to Upgrade due to vCenter Issues	64
HX Data Platform Installer Shows Host Managed by Different vCenter	64
Configuration Settings Differ Between HX Data Platform and UCS Manager	64
Cluster Creation Fails with DNS Error due to FQDN	65
Offline Upgrade Cluster Start Command Error: Node Not Available	66
vSphere Replication Plug-in Fails after HX Plug-in Deployed	66
Upgrade Fails but Reports All Nodes Up-to-Date	68
Restarting Online Upgrade Fails	68
Controller VM Fails to Power On During Cisco UCS Upgrade	69
Firmware Upgrade Fails from Server Storage Controller with Unsupported Board	69
A Node Fails to Upgrade due to vCenter Issues	69
Upgrade Stalls Waiting for Node to Return to Healthy State	70
Cluster Expansion Error: No cluster found	70
Cluster Expansion Fails when DNS Server is Not Found	70
Expand Cluster Fails with Stale HX Installer	72
Installation Fails when Secure Boot is Enabled	72
Host Issues	73
Post Manual ESX Installation statsd restart	73
scvmlclient Management Services Restarted when services.sh restart Issued	74
ESX Server Reboot on Upgrade Exited Maintenance Mode	74
EAM Did Not Start on Compute Node	74
Remove Node Fails when a Node is Down.	75
Rebooting an HA Enabled ESX Host	75
Node Failure While Simultaneously Adding Another Node to Cluster	75
Configure PCI Passthrough After Changing vNIC or vHBAs	76
Secure Boot Cannot Be Enabled After Upgrade	77

Disk Issues	77
Cannot Allocate Data when All or Most Disks Fail on a Node in a 3 Node Storage Cluster	77
Removing Disks Causes Rebooting Storage Controller VM to Fail	78
Cluster Management IP Fails after NVME Drive is Replaced	78
Recover Failed SSD Hosting the Storage Controller VM	78
How to synchronize SCVM clock after an installation	79
vNode Scrubber Detection	79
VM Issues	79
Controller VM Prevented from Powering On	79
Automatic vMotion Migrations Fail with Timeout Error	80
Storage Controller VM Power On Fails with Two Node Failures	80
Fail Adding VM to Host with HA and DRS Enabled	80
Degraded Performance on VM with Disk Limit Shares	81
DRS Migrates VMs when Storage Cluster in Read Only State	81
VM Power On Fails Due to Stale EAM Extension	81
Deleting VM Folder or File Taking Very Long Time	83
VM Disk Usage and vCenter VM Committed Size Mismatch	83
Migrating a VM Task Fails and Results in Replication Error 10001	83
VM Migration Results in an Error	84
VM Migration BadVersionException Error	84
Datastore Issues	84
Removing Compute Node Did Not Remove Datastores	84
Adding Multiple Datastores Error: mountpoint not found on host	85
NFS All Paths Down with Message File Locked by Consumer on Host	85
Maximum Queue Depth Limit of NFS is not Optimal	85
Mounting Datastore Fails After Changing VLAN ID	86
Datastore fails to mount if data vSwitch has an existing vmkernel port configured with an IP address	86
Datastore Mount Fails after Motherboard Updated with ESXi 7.0 U3 Node	86
Remounting Datastore after Cluster Restart	87
Datastore Does Not Unmount When Storage I/O RM Daemon Running	87
Datastore Delete Fails with error: com.vmware.vim25.PlatformConfigFault	88
Datastore Not Accessible Due to Clock Skew	88
Datastore not synced results in errors during Disaster Recovery	89

ReadyClone, Snapshot, and Replication Issues	89
Replication Fails when using VMware Tools to Quiesce VMs	89
Replication Errors with VMware Guest Provider (quiesce)	90
Reprotect Operation Fails on a Recovered VM when Another Replication is In-progress on the Same VM	90
Migrate Task Fails When Target Datastore Is Renamed	90
Backup Workflows Fail with an error message	91
SRM Recovery Fails with vSphere Cluster Service (vCLS)	91
Backup Software Fails when SSLv3 Disabled	91
Recovery Fails after Renaming Datastore	92
Recover, Migrate, or Test Recovery fail when Silent Mode notification is enabled	92
Rebooting the Node Stops Recovery, Restore, and Clone	92
Rerunning stcli vm recover Command Fails	93
REST API Protection Group Filter Returns All Groups	93
VM Stunned During Snapshot Consolidation	94
Native Snapshots with Quiesce Option	94
Quiesce Based Snapshots Fail without an Error Message	94
vMotion Fails to Move a Native Snapshot Datastore	94
Reprotect option is unavailable for VMs in Protecting State	95
Site Recovery Manager Failover and Reprotect Issues	95
Cluster Issues	96
After Cluster reregister Controller VMs not EAM Agents	96
Cluster Becomes Unhealthy after Multiple Reregisters	96
ClusterNotConfigured Error after Node Removed	96
Cluster Capacity Higher than Individual Disks	96
Re-registering a Cluster Does Not Re-register Compute Nodes with EAM	97
Latency Spikes Seen for Workloads with Large Working Sets	97
Cluster Health Status Remains Unhealthy after Rebalance	98
NTP Not Configured on ESXi Hosts	98
Cluster Capacity Different Than Provisioned	98
Connectivity to Storage Controller VM when using vShield	99
Storage Cluster Missing from vCenter Cluster after Cluster Node Powered Off	99
Interface Issues	99
Multiple VM Power Operations Causes Task Queue to Error Out	99

HX Connect Data Does Not Refresh 100

Performance Charts Show a Gap while Node Rebooting 100

Cannot See the HX Data Platform Plug-in Through vSphere Clients 101

Performance Charts Display Not Formatted at 100% 101

HX Data Platform Plug-In Feature Not Performing 102





## Communications, Services, Bias-free Language, and Additional Information

---

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco Marketplace](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

### Documentation Feedback

To provide feedback about Cisco technical documentation, use the feedback form available in the right pane of every online document.

### Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.

### Bias-Free Language

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.





# CHAPTER 1

## New and Changed Information

---

- [New and Changed Information](#), on page 1

### New and Changed Information

This table summarizes the new and changed features for the Cisco HyperFlex Troubleshooting Guide, Release 5.0(x) and tells you where they are documented.

Feature	Description	Release/Date Added	Where Documented
HX License Information in Basic Support Bundles	Added information describing HX License Information in Basic Support Bundles.	HX 5.0(2b)	<a href="#">Generating a Support Bundle Using HX Connect</a> , on page 12
Cisco HyperFlex Troubleshooting Reference Guide	First version for Cisco HyperFlex Release 5.0	HX 5.0(1a)	N/A





## CHAPTER 2

# HX Data Platform Troubleshooting

---

- [Troubleshooting HyperFlex Systems Overview, on page 3](#)
- [Contacting Cisco TAC, on page 4](#)

## Troubleshooting HyperFlex Systems Overview

The HyperFlex product has multiple integrated components. This troubleshooting guide includes topics specific to HyperFlex. Some of the integrated components function outside of HyperFlex. Refer to the documentation for those products for additional assistance.

- **Cisco HyperFlex (HX) Data Platform** – See this troubleshooting guide. This includes troubleshooting for HX Data Platform installation, configuration, and management, UCS Manager to HyperFlex configuration, and vSphere to HyperFlex configuration.
- **Cisco UCS and UCS Manager** – See the UCS Manager documentation for general USC Manager issues.
- **Cisco HyperFlex servers** – See the hardware installation and maintenance guides for additional information.
- **VMware vSphere, vCenter, or ESX** – See the VMware documentation for general VMware related issues.

This HyperFlex Troubleshooting guide contains topics for:

- Common HyperFlex TAC assisted topics
- HyperFlex Data Platform support bundles
- HX Data Platform Event Messages
- HyperFlex component and process issues

The content in this Cisco HyperFlex Systems Troubleshooting Guide is supplement to the information provided in the HyperFlex Data Platform documentation. Refer to the HyperFlex Data Platform guides for requirements, practices, and procedures.

# Contacting Cisco TAC

You can open a Cisco Technical Assistance Center (TAC) support case to reduce time addressing issues and get efficient support directly with Cisco Prime Collaboration application.

For all customers, partners, resellers, and distributors with valid Cisco service contracts, Cisco Technical Support provides around-the-clock, award-winning technical support services. The Cisco Technical Support website provides online documents and tools for troubleshooting and resolving technical issues with Cisco products and technologies:

<http://www.cisco.com/techsupport>

Using the TAC Support Case Manager online tool is the fastest way to open S3 and S4 support cases. (S3 and S4 support cases consist of minimal network impairment issues and product information requests.) After you describe your situation, the TAC Support Case Manager automatically provides recommended solutions. If your issue is not resolved by using the recommended resources, TAC Support Case Manager assigns your support case to a Cisco TAC engineer. You can access the TAC Support Case Manager from this location:

<https://mycase.cloudapps.cisco.com/case>

For S1 or S2 support cases or if you do not have Internet access, contact the Cisco TAC by telephone. (S1 or S2 support cases consist of production network issues, such as a severe degradation or outage.) S1 and S2 support cases have Cisco TAC engineers assigned immediately to ensure your business operations continue to run smoothly.

To open a support case by telephone, use one of the following numbers:

- Asia-Pacific: +61 2 8446 7411
- Australia: 1 800 805 227
- EMEA: +32 2 704 5555
- USA: 1 800 553 2447

For a complete list of Cisco TAC contacts for Enterprise and Service Provider products, see <http://www.cisco.com/c/en/us/support/web/tsd-cisco-worldwide-contacts.html>.

For a complete list of Cisco Small Business Support Center (SBSC) contacts, see <http://www.cisco.com/c/en/us/support/web/tsd-cisco-small-business-support-center-contacts.html>.



## CHAPTER 3

# HX Data Platform Support

---

- [Using TAC Support for Troubleshooting, on page 5](#)
- [Auto Support and Smart Call Home for HyperFlex, on page 6](#)
- [Cisco HX Data Platform Support Bundles, on page 11](#)

## Using TAC Support for Troubleshooting

This section lists common TAC support-assisted tasks, and describes how to configure automatic support options and create HyperFlex support bundles.

Support bundles are a collection of logs from several components within and used by the HX Data Platform. This includes:

- HX Data Platform Installer VM—Logs provide information about installation.
- Controller VM—Logs provide information about the HX Data Platform file system, cluster creation, and cluster expansion.
- VMware ESXi host—Logs provide information about the nodes in the HX storage cluster.
- VMware vCenter—Logs provide information about the HX Data Platform Plug-in and the vCenter server.

TAC uses support bundles to assist in troubleshooting your issues.

## Common TAC-Assisted Support Topics

The following is a brief list of support topics that are best handled with Technical Assistance Center (TAC) assistance:

- Adding disks to a node in the HX storage cluster fails to be recognized.
- Adding a node to the HX storage cluster fails.
- Changing the IP address for the HX storage cluster.
- Destroying a cluster, including destroying an encrypted cluster.
- Downgrading the HX Data Platform version.
- Failing HX storage cluster creation.

- Modifying a node rebalance timeout.
- Modifying VDI or VSI optimized deployment for the HX storage cluster.
- Removing a node in a 3-node cluster.
- Replacing a node in a 4-node cluster.
- Replacing a housekeeping SSD on HX240c servers.
- Reusing a removed node in another HX storage cluster.
- Reusing disks from a removed node.
- Setting the `cleaner` schedule for recovering storage.
- Setting the MTU value to something other than 9000.
- Sizing guidance for using non-default larger capacity SSD or HDD per server is supported.
- Uninstalling the HX Data Platform.
- Uninstalling the HX storage cluster.
- Upgrading the HX Data Platform from a version older than HX Data Platform version 1.7.1.
- Using the `stcli` commands `whitelist` or `recreate`.

## Auto Support and Smart Call Home for HyperFlex

You can configure the HX storage cluster to send automated email notifications regarding documented events. You can use the data collected in the notifications to help troubleshoot issues in your HX storage cluster.




---

**Note** Auto Support (ASUP) and Smart Call Home (SCH) support the use of a proxy server. You can enable the use of a proxy server and configure proxy settings for both using HX Connect.

---

### Auto Support (ASUP)

Auto Support is the alert notification service provided through HX Data Platform. If you enable Auto Support, notifications are sent from HX Data Platform to designated email addresses or email aliases that you want to receive the notifications. Typically, Auto Support is configured during HX storage cluster creation by configuring the SMTP mail server and adding email recipients.




---

**Note** Only unauthenticated SMTP is supported for ASUP.

---

If the **Enable Auto Support** check box was not selected during configuration, Auto Support can be enabled post-cluster creation using the following methods:



Post-Cluster ASUP Configuration Method	Associated Topic
HX Connect user interface	<a href="#">Configuring Auto Support Using HX Connect, on page 7</a>
Command Line Interface (CLI)	<a href="#">Configuring Notification Settings Using CLI, on page 8</a>
REST APIs	Cisco HyperFlex Support REST APIs on <a href="#">Cisco DevNet</a> .

Auto Support can also be used to connect your HX storage cluster to monitoring tools.

### Smart Call Home (SCH)

Smart Call Home is an automated support capability that monitors your HX storage clusters and then flags issues and initiates resolution before your business operations are affected. This results in higher network availability and increased operational efficiency.

Call Home is a product feature embedded in the operating system of Cisco devices that detects and notifies the user of a variety of fault conditions and critical system events. Smart Call Home adds automation and convenience features to enhance basic Call Home functionality. After Smart Call Home is enabled, Call Home messages/alerts are sent to Smart Call Home.

Smart Call Home is included with many Cisco service contracts and includes:

- Automated, around-the-clock device monitoring, proactive diagnostics, real-time email alerts, service ticket notifications, and remediation recommendations.
- Proactive messaging sent to your designated contacts by capturing and processing Call Home diagnostics and inventory alarms. These email messages contain links to the Smart Call Home portal and the TAC case if one was automatically created.
- Expedited support from the Cisco Technical Assistance Center (TAC). With Smart Call Home, if an alert is critical enough, a TAC case is automatically generated and routed to the appropriate support team through `https`, with debug and other CLI output attached.
- Customized status reports and performance analysis.
- Web-based access to all Call Home messages, diagnostics, and recommendations for remediation in one place; TAC case status; and up-to-date inventory and configuration information for all Call Home devices.

To ensure automatic communication among your HX storage cluster, you, and Support, see [Configuring Smart Call Home for Data Collection, on page 9](#).

## Configuring Auto Support Using HX Connect

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using the HX Connect user interface.

- 
- Step 1** Log into HX Connect.
- Step 2** In the banner, click **Edit settings (gear icon)** > **Auto Support Settings** and fill in the following fields.

UI Element	Essential Information
<b>Enable Auto Support (Recommended)</b> check box	Configures Call home for this HX storage cluster by enabling: <ul style="list-style-type: none"> <li>• Data delivery to Cisco TAC for analysis.</li> <li>• Notifications from Support as part of proactive support.</li> </ul>
<b>Send service ticket notifications to</b> field	Enter the email address that you want to receive the notifications.
<b>Terms and Conditions</b> check box	End user usage agreement. The check box must be checked to use the Auto-Support feature.
<b>Use Proxy Server</b> check box	<ul style="list-style-type: none"> <li>• <b>Web Proxy Server</b> url</li> <li>• <b>Port</b></li> <li>• <b>Username</b></li> <li>• <b>Password</b></li> </ul>

**Step 3** Click **OK**.

**Step 4** In the banner, click **Edit settings (gear icon) > Notifications Settings** and fill in the following fields.

UI Element	Essential Information
<b>Send email notifications for alarms</b> check box	If checked, fill in the following fields: <ul style="list-style-type: none"> <li>• <b>Mail Server Address</b></li> <li>• <b>From Address</b>—Enter the email address used to identify your HX storage cluster in Support service tickets, and as the sender for Auto Support notifications. Support information is currently not sent to this email address.</li> <li>• <b>Recipient List (Comma separated)</b></li> </ul>

**Step 5** Click **OK**.

## Configuring Notification Settings Using CLI

Use the following procedure to configure and verify that you are set up to receive alarm notifications from your HX storage cluster.



**Note** Only unauthenticated SMTP is supported for ASUP.

**Step 1** Log into a storage controller VM in your HX storage cluster using `ssh`.

**Step 2** Configure the SMTP mail server, then verify the configuration.

Email address used by the SMTP mail server to send email notifications to designated recipients.

Syntax: `stcli services smtp set [-h] --smtp SMTPSERVER --fromaddress FROMADDRESS`

Example:

```
# stcli services smtp set --smtp mailhost.eng.mycompany.com --fromaddress smtpnotice@mycompany.com
# stcli services smtp show
```

**Step 3** Enable ASUP notifications.

```
# stcli services asup enable
```

**Step 4** Add recipient email addresses, then verify the configuration.

List of email addresses or email aliases to receive email notifications. Separate multiple emails with a space.

Syntax: `stcli services asup recipients add --recipients RECIPIENTS`

Example:

```
# stcli services asup recipients add --recipients user1@mycompany.com user2@mycompany.com
# stcli services asup show
```

**Step 5** From the controller VM that owns the eth1:0 IP address for the HX storage cluster, send a test ASUP notification to your email.

```
# sendasup -t
```

To determine the node that owns the eth1:0 IP address, log into each storage controller VM in your HX storage cluster using `ssh` and run the `ifconfig` command. Running the `sendasup` command from any other node does not return any output and tests are not received by recipients.

**Step 6** Configure your email server to allow email to be sent from the IP address of all the storage controller VMs.

---

## Configuring Smart Call Home for Data Collection

Data collection is enabled by default but, you can opt-out (disable) during installation. You can also enable data collection post cluster creation. During an upgrade, Smart Call Home enablement is determined by your legacy configuration. For example, if `stcli services asup show` as enabled, Smart Call Home is enabled on upgrade.

Data collection about your HX storage cluster is forwarded to Cisco TAC through `https`. If you have a firewall installed, configuring a proxy server for Smart Call Home is completed after cluster creation.



---

**Note** Smart Call Home does not support the use of a proxy server in deployments where outgoing connections from an HX cluster require to go through a proxy server.

---

Using Smart Call Home requires the following:

- A Cisco.com ID associated with a corresponding Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service contract for your company.
- Cisco Unified Computing Support Service or Cisco Unified Computing Mission Critical Support Service for the device to be registered.

**Step 1** Log into a storage controller VM in your HX storage cluster.

**Step 2** Register your HX storage cluster with Support.

Registering your HX storage cluster adds identification to the collected data and automatically enables Smart Call Home. To register your HX storage cluster, you need to specify an email address. After registration, this email address receives support notifications whenever there is an issue and a TAC service request is generated.

Syntax:

```
stcli services sch set [-h] --email EMAILADDRESS
```

Example:

```
# stcli services sch set --email name@company.com
```

**Step 3** Verify data flow from your HX storage cluster to Support is operational.

Operational data flow ensures that pertinent information is readily available to help Support troubleshoot any issues that might arise.

**Note** Contact TAC to verify connectivity.

```
# asupcli [--all] ping
```

--all option runs the commands on all the nodes in the HX cluster.

**Step 4** (Optional) Configure a proxy server to enable Smart Call Home access through port 443.

If your HX storage cluster is behind a firewall, after cluster creation, you must configure the Smart Call Home proxy server. Support collects data at the url: <https://diag.hyperflex.io:443> endpoint.

a. Clear any existing registration email and proxy settings.

```
# stcli services sch clear
```

b. Set the proxy and registration email.

Syntax:

```
stcli services sch set [-h] --email EMAILADDRESS [--proxy-url PROXYURL] [--proxy-port PROXYPORT]
[--proxy-user PROXYUSER] [--portal-url PORTALURL] [--enable-proxy ENABLEPROXY]
```

Syntax Description	Option	Required or Optional	Description
	--email EMAILADDRESS	Required.	Add an email address for someone to receive email from Cisco support. Recommendation is to use a distribution list or alias.
	--enable-proxy ENABLEPROXY	Optional.	Explicitly enable or disable use of proxy.
	--portal-url PORTALURL	Optional.	Specify an alternative Smart Call Home portal URL, if applicable.
	--proxy-url PROXYURL	Optional.	Specify the HTTP or HTTPS proxy URL, if applicable.

Option	Required or Optional	Description
<code>--proxy-port PROXYPORT</code>	Optional.	Specify the HTTP or HTTPS proxy port, if applicable.
<code>--proxy-user PROXYUSER</code>	Optional.	Specify the HTTP or HTTPS proxy user, if applicable. Specify the HTTP or HTTPS proxy password, when prompted.

Example:

```
# stcli services sch set
  --email name@company.com
  --proxy-url www.company.com
  --proxy-port 443
  --proxy-user admin
  --proxy-password adminpassword
```

- c. Ping to verify the proxy server is working and data can flow from your HX storage cluster to the Support location.

**Note** Contact TAC to verify connectivity.

```
# asupcli [--all] ping
```

--all option runs the command on all the nodes in the HX cluster.

**Step 5** Verify Smart Call Home is enabled.

When Smart Call Home configuration is `set`, it is automatically enabled.

```
# stcli services sch show
```

**Step 6** Enable Auto Support (ASUP) notifications.

Typically, Auto Support (ASUP) is configured during HX storage cluster creation. If it was not, you can enable it post cluster creation using HX Connect or CLI.

If Smart Call Home is disabled, enable it manually.

```
# stcli services sch enable
```

## Cisco HX Data Platform Support Bundles

Beginning with Cisco HX Release 4.0(x), the recommended method of collecting support bundles is through the HX Connect user interface. HX Connect provides a simple easy to use interface enabling you to manage your HX environment. You can generate a support bundle that collects the logs from every selected controller VM and ESXi host in the HX storage cluster. The vCenter logs are not collected through HX Connect.

All support bundle timestamps are listed in the UTC timezone regardless of cluster timezone or server timezone settings.

After you generate a support bundle, you can upload it to the HX Data Platform FTP server for use by TAC. You can also download an existing support bundle.

Table 1: Support Bundle Types

Support Bundle Types	Details
All	Displays all logs.
Basic	Cisco HX Data Platform logs. Use this option to open TAC Cases. Starting with HX 5.0(2b), the basic support bundle includes license information with registration and expiration status. Does not contain roll over logs and the size is limited to 250Mb.
Detailed	Gathers Hyper-V logs and performance data for the environment in addition to the Basic support bundle. This is the default Support bundle as generated by storfs-support command.
Extended	When generated with Extended support bundle option alone, it only contains core files. When generated with Recommended Support bundle and the Extended Support bundle option, it includes core files and detailed support bundles.

To generate support bundles through the HX Connect user interface continue to the [Generating a Support Bundle Using HX Connect, on page 12](#) section.

In the event that HX Connect is offline, you can generate a support bundle through command line interface. To get started, continue to [Collecting Logs from Controller VMs, on page 16](#) or the [Collecting Logs from Controller VMs: HX Release 4.0\(2a\), on page 15](#) section.

## Generating a Support Bundle Using HX Connect

You can use the HX Connect user interface to generate a support bundle that collects the logs from every controller VM and ESXi host in the local HX storage cluster. If you are using replication to protect your virtual machines and their data, when you need to generate a support bundle, you also need to generate a support bundle from the remote HX storage cluster. The vCenter logs are not collected through HX Connect.





---

**Note** Creating logs from HX Connect may cause storfs to crash causing a Hyperflex cluster outage.


---

**Step 1** Log in to HX Connect.


**Step 2** In the banner, click **Edit Settings (gear icon) > Support Bundle**.



**Step 3** Under Support Bundle, select the tab you want to display. Options include: **All**, **Basic**, **Detailed**, and **Extended**. For more information about support bundles, see [Cisco HX Data Platform Support Bundles, on page 11](#).



**Step 4** Click **Generate**. The **Select bundle options** window appears.




**Step 5** Select the **Support bundle type** you will generate. The default option is **Recommended support bundle**.

The **Recommended support bundle** triggers creation of both **Basic support bundle** and **Detailed support bundle** from the same option.


a) All the nodes are selected by default. Uncheck the Nodes which are not required.

**Step 6** Click **Generate**.



Hover over the grayed out **Generate** button to view the **Basic support bundle generation in process** message.

**Step 7** If Detailed support bundle is selected, bundle generation begins when the Basic support bundle is complete. Confirm Detailed bundle generation by browsing to the Basic tab and hover over the grayed out **Generate** button to view the **Detailed support bundle generation in process** message.



Support bundle creation may take more than one hour.

**Step 8** Once the support bundle(s) are generated, the **Support Bundle Generation Completed** message appears.

**Step 9** Download your support bundle(s) as needed.

### What to do next

You can now upload the file to the HX Data Platform FTP server.

## Collecting Basic Support Bundles from Controller VMs

This feature was introduced in Cisco HX Release 5.0(2b).

**Step 1** Log in to each controller VM using `ssh`.

**Step 2** Run the command: `storfs-support --basic` to generate the basic support bundle.

**Step 3** Locate the generated support bundle file in the `/var/support` directory. For example:

```
storfs-support_2022-08-26-12-12-32_SpringpathControllerZ2RTF9ME88_WZP2216001X_basic.tar.gz
```

**Step 4** Untar the support bundle. For example:



```
tar -zxvf
storfs-support_2022-08-26-12-12-32_SpringpathControllerZ2RTF9ME88_WZP2216001X_basic.tar.gz
```

**Step 5** Verify that the file "smartlicensing.json" is present with the licensing information. For example:

```
hshell:/var/support/storfs-support_2022-08-26-12-12-32_SpringpathControllerZ2RTF9ME88_WZP2216001X_basic/var/support/asp-restapi-outputs$
ls
smartlicensing.json
```

## Collecting Logs from Controller VMs: HX Release 4.0(2a)

The **storfs-support** CLI is used to generate the default Detailed or Basic Support bundle for HX Release 4.0(2a). While the **storfs-support** CLI has the option for the basic support bundle, it does not offer the option to append the file name (basic vs detailed) to the support bundle file name.

To get started, select either **Option 1** to generate the Detailed Support bundle step or **Option 2** to generate the Basic Support bundle steps.



**Note** If you are generating a support bundle using the CLI, then the support bundle is displayed as **Other** type in HX Connect user interface.

### Step 1 Option 1 - Default Detailed Support Bundle

- Log in to each controller VM using ssh.
- Run the command: `storfs-support` to generate the default or detailed support bundle.

```
root@SpringpathController0FDF9RMJK:~# storfs-support
2017-04-28 05:24:18,505 - Storfs-Support - INFO -
2017-04-28 05:24:18,505 - Storfs-Support - INFO -
2017-04-28 05:24:18,505 - Storfs-Support - INFO - Initiating support generation...
2017-04-28 05:24:18,506 - Storfs-Support - INFO -
2017-04-28 05:24:18,506 - Storfs-Support - INFO -
2017-04-28 05:24:18,506 - Storfs-Support - INFO - Generating support archive. This may take some time...
2017-04-28 05:24:18,506 - Storfs-Support - INFO -
2017-04-28 05:31:57,692 - Storfs-Support - INFO - Support archive generated at: /var/support/storfs-support_2017-04-28--0
2017-04-28 05:31:57,692 - Storfs-Support - INFO - Removing directory... /var/support/cmds_output
```

Generating the logs takes about 2 minutes.

- Locate the `tar.gz` log file in the `/var/support` directory. For example:

```
storfs-support_2017-04-28--06-06-33_ucs--stctlvm-123-1.eng.storvisor.com.tar.gz
```

- Upload the `tar.gz` file from the controller VM to the HX Data Platform FTP server.

### Step 2 Option 2 - Basic Support Bundle

- Log in to each controller VM using ssh.
- Run the command: `storfs-support --basic` to generate the default or detailed support bundle.

```

root@hx-02-scvm-03:~# storfs-support --basic
/var/support
2020-02-19 12:33:01,315 - Storfs-Support - INFO -
2020-02-19 12:33:01,315 - Storfs-Support - INFO -
2020-02-19 12:33:01,315 - Storfs-Support - INFO - Initiating support generation...
2020-02-19 12:33:01,317 - Storfs-Support - INFO -
2020-02-19 12:33:01,317 - Storfs-Support - INFO -
2020-02-19 12:33:01,317 - Storfs-Support - INFO - Generating support archive. This may take some time...
2020-02-19 12:33:01,317 - Storfs-Support - INFO -
2020-02-19 12:33:01,317 - Storfs-Support - INFO -
2020-02-19 12:33:01,318 - Storfs-Support - INFO - STORFS_RUNTIMEIDIR: /tmp
2020-02-19 12:33:01,318 - Storfs-Support - INFO - STORFS_SOURCEDIR:

2020-02-19 12:35:34,446 - Storfs-Support - INFO -
2020-02-19 12:35:34,446 - Storfs-Support - INFO - Support archive generated at: /var/support/storfs-support_2020-02-19--1
2020-02-19 12:35:34,446 - Storfs-Support - INFO - Removing directory... /var/support/cmds_output

```

Note that the file name for the basic bundle is the same as the one generated for the detailed storfs-support in Option 2a. Generating the logs takes about 2 minutes.

- c) Locate the `tar.gz` log file in the `/var/support` directory. For example:

```
storfs-support_2020-02-19--06-06-33_ucs--stctlvm-123-1.eng.storvisor.com.tar.gz
```

- d) Upload the `tar.gz` file from the controller VM to the HX Data Platform FTP server.

## Collecting Logs from Controller VMs

The `storfs-support` CLI is used to collect controller VM logs.



**Note** This method is not recommended for users using Cisco HX Release 4.0(2a).

To get started, perform the following steps.

**Step 1** Log in to each controller VM using `ssh`.

**Step 2** Run the command: `storfs-support`.

Generating the logs takes about 2 minutes.

**Step 3** Locate the `tar.gz` log file in the `/var/support` directory. For example:

```
storfs-support_2020-02-19--06-06-33_ucs--stctlvm-123-1.eng.storvisor.com.tar.gz
```

**Step 4** Upload the `tar.gz` file from the controller VM to the HX Data Platform FTP server.

## Generating an Audit Log

When you generate a support bundle through the HX Connect user interface, audit logs are automatically included.

---

**Step 1** Generate and download a support bundle. See [Generating a Support Bundle Using HX Connect, on page 12](#).

If you prefer, you can download an existing support bundle.

**Step 2** Unzip the support bundle and extract the files to your local computer.

**Step 3** In each controller VM, search for `/var/log/shell.log`.

**Step 4** Open each `shell.log` file and search for the audit trail records using keyword **stcli**.

This log contains an audit of `stcli` calls that are invoked in the shell.

**Example:**

```
2017-07-14T16:48:39.135+00:00 SpringpathControllerHOCBY4KNF1 shell: [pid=20396, uid=0] stcli cluster restart
```

**Where**

- `2017-07-14T16:48:39.135+00:00` is the time when the `stcli` command was invoked.
- `uid=0` specifies the user ID of the user who invoked the `stcli` command. In the example, the user ID is 0 and the user is root.
- `stcli cluster restart` command provides the `stcli` command that was executed.

**Step 5** In each controller VM, search for `/var/log/springpath/audit-rest.log`.

**Step 6** Open each `audit-rest.log` file and search for the audit trail records using keyword **audit**.

This log contains an audit of REST APIs.

**Example:**

```
2017-06-29-23:26:38.096 - Audit - 127.0.0.1 -> 127.0.0.1 - create /rest/datastores/00000000d8902473:00000000000100ef?action=mount; 200; administrator@vsphere.local 555ms
```

**Where**

- `2017-06-29-23:26:38.096` is the time when the REST API was invoked.
- `127.0.0.1` is the IP address from which the call arrived.
- `create` is the action that was performed.
- `/rest/datastores/00000000d8902473:00000000000100ef?action=mount` is the resource that is accessed with parameters.
- `200` is the HTTP status of this action.
- `administrator@vsphere.local` is the user who invoked this REST API.
- `555ms` is the time taken for this operation.

**Step 7** Collect all audit trail records from the previous steps and save them in separate files.

---

## Collecting Logs from ESXi Hosts

There are two options for collecting ESXi host logs.

### Step 1 Option 1

- a) Log in to each ESXi host using `ssh`.
- b) Run the command: `vm-support`

Generating the logs takes about 5 minutes.

- c) Locate the `.tgz` file in the `/var/tmp` directory. For example:

```
esx-localhost-2016-06-22--06.09.tgz
```

- d) Upload the `.tgz` file from the ESXi host to the HX Data Platform FTP server.

### Step 2 Option 2

- a) Log in to each controller VM using `ssh`.
- b) Run the command: `asupcli collect --type esx --subtype full`
- c) Upload the `tar.gz` file in the `/var/support/esx-asup-default` directory from each controller VM to the HX Data Platform FTP server.

The `tar.gz` file contains only ESXi logs. If TAC requested all logs, see [Collecting Logs from Controller VMs: HX Release 4.0\(2a\)](#), on page 15

## Out of Space Error When Generating a Support Bundle

An out of space error occurs when the storage controller VM does not have sufficient space left to generate the support bundle, typically due to the size of the core file or previously-generated log files consuming space. The following error displays when you are using the `vm-support` command to generate support bundles:

```
error = [Errno 28] No space left
```

To generate a support bundle when you receive this error:

**Step 1** Delete or move the core file and existing log files to a location outside of the storage controller VM.

**Step 2** Log in to the command line of the storage controller VM.

**Step 3** Generate a light support bundle.

```
# storfs-support
```

## Collecting Deployment Logs from the Cisco HX Data Platform Installer VM

**Step 1** Log in to the HX Data Platform Installer VM using `ssh` and the following credentials:

- Username: **root**
- Password (Default): **Cisco123**

**Note** Systems ship with a default password of `Cisco123` that must be changed during installation. You cannot continue installation unless you specify a new user supplied password.

**Step 2** Run the command: `deployment-support`  
Generating the logs takes less than a minute.

**Step 3** Locate the `.tar.gz` logs in the `/var/support` directory. For example:  
`storfs-support_2016-06-22--06-25-35_Cisco-HX-Data-Platform-Installer.tar.gz`

**Step 4** Upload the `.tar.gz` file to the Cisco FTP server.

---

## Collecting Logs from the vCenter Server

---

**Step 1** Log in to the vCenter server using `ssh`.

**Step 2** Run the command: `vc-support -l`  
Generating the logs takes about 10-20 minutes depending on the number of nodes that are running in the vCenter server.

**Step 3** Locate the logs in the `/storage/log` directory.

**Step 4** Upload the logs to the Cisco FTP server.

---

## Collecting Logs from the vSphere Web Client

You can selectively collect all or some of the logs for the Cisco HX Data Platform storage cluster ESXi hosts, controller VMs, and vCenter server through the vSphere Web Client.

---

**Step 1** Log in to the vSphere Web Client. Through the **Navigator**, select **vCenter Inventory Lists > Resources > vCenter Servers > server**.

**Step 2** Right-click the vCenter server for the HX Data Platform cluster and select **Export System Logs**.

**Step 3** From the **Source** panel, select the servers from which you want to collect the logs.

**Step 4** Optionally, to include the vCenter logs, select **Include vCenter Server and vSphere Web Client logs**, and click **Next**.

**Step 5** From the **Ready to Complete** panel, select the system logs to include.  
Select **Virtual Machines** to ensure the logs from the controller VMs for each selected server are included.

**Step 6** Click **Generate Log Bundle**.  
Generating the logs takes about 40-50 minutes.

**Step 7** When the logs are completed:  
a) Click **Download Log Bundle**.

- b) Specify a location to download the support bundle. For example, download the files to your local PC.
- c) Click **Finish**.
- d) Upload the files to the Cisco FTP server.

## Collecting Logs from the Cisco HXDP Plug-in

You can use the HX Data Platform Plug-in to collect the logs for just the HX storage cluster ESXi hosts and controller VMs.

- Step 1** Log in to vSphere Web Client. Through the **Navigator**, select **vCenter Inventory Lists > Resources > Clusters > cluster**.
- Step 2** Right-click an HX Data Platform cluster and select **Cisco HX Data Platform > Support**.
- Step 3** Click **Generate Log Bundle**.  
Generating the logs takes about 40-50 minutes.
- Step 4** When the logs are completed:
  - a) Click **Download Log Bundle**.
  - b) Specify a location to download the support bundle. For example, download the files to your local PC.
  - c) Click **Finish**.
  - d) Upload the files to the Cisco FTP server.

## Downloading an Existing Support Bundle



**Note** You can download a support bundle you generate up to midnight (12:00 a.m.) of the local controller VM time. HyperFlex stores two support bundle logs. When generating a new support bundle, the oldest one is deleted automatically.

- Step 1** Log in to vSphere Web Client. Through the **Navigator**, select **vCenter Inventory Lists > Resources > Clusters**.
- Step 2** Right-click the *cluster* and select **HX Data Platform Plug-in > Support** to display the **Support** dialog box.  
**Note** Select the *cluster* from the vSphere Web Client, not the HX Data Platform Plug-in.
- Step 3** Click **Download Support Bundle** (enabled).
- Step 4** Specify a download location (such as your local computer) to save the support bundle. You can then upload the file to the HX Data Platform FTP server.

## Uploading Support Bundles

After you generate a support bundle, you can upload it to the HX Data Platform FTP server using one of the following methods:

- ftp—Use port 21.
- sftp—Use port 22.
- https—Use port 443.

### Uploading Support Bundles using ftp or sftp

#### Before you begin

Generate a support bundle.

**Step 1** Open your FTP Client (such as Filezilla) and connect to the HX Data Platform FTP server using the following information:

Option	Description
Host	https://ftp.springpathinc.com
Port	ftp = 21 sftp = 22
Username	cisco
Password	cisco

**Step 2** After you connect to the HX Data Platform FTP server, create a folder with the Support case number.

**Step 3** Change directory to the new folder.

**Step 4** Upload the support bundle log files to this folder.

**Step 5** When the uploads are complete, contact Cisco Technical Assistance Center (TAC) and tell them the upload directory name.

**Step 6** Free up space in your HX storage cluster by deleting the content in `/var/support/`.

### Uploading Support Bundles using https

#### Before you begin

Generate a support bundle.

**Step 1** Open a browser window, navigate to the HX Data Platform FTP server URL hosted by Springpath, and log in using the following information:

Option	Description
URL	https://ftp.springpathinc.com
Username	cisco
Password	cisco

- Step 2** After you connect to the HX Data Platform FTP server, create a folder with the Support case number.
- Step 3** Change directory to the new folder.
- Step 4** Upload the support bundle log files to this folder.
- Step 5** When the uploads are complete, contact Cisco Technical Assistance Center (TAC) and tell them the upload directory name.
- Step 6** Free up space in your HX storage cluster by deleting the content in `/var/support/`.
-





## CHAPTER 4

# HX Data Platform Events

---


- [Understanding System Events, on page 23](#)
- [View HX Data Platform Plug-in Events , on page 24](#)
- [Automatically Acknowledged Events, on page 25](#)
- [Manually Acknowledging Alarms Using HX Connect, on page 26](#)
- [Cluster Access Policy Compliance Events, on page 26](#)
- [Cluster Events, on page 27](#)
- [Critical Infrastructure Events, on page 31](#)
- [Disk Events, on page 31](#)
- [Host Events, on page 34](#)
- [Node Events, on page 35](#)
- [Encryption Events, on page 37](#)
- [Replication and Recovery Events, on page 45](#)
- [Security Events, on page 47](#)
- [Resource Monitoring Events and Alarms, on page 48](#)
- [Smart Call Home Events, on page 53](#)
- [Smart License Events, on page 53](#)
- [Snapshot Events, on page 55](#)
- [Space Utilization Events, on page 56](#)
- [Storage Events, on page 59](#)
- [Upgrade Events, on page 59](#)

## Understanding System Events

HX Data Platform messages include the error, warning, and informational messages that the system displays during various events. These include HX storage cluster-wide events and events due to changes in HX storage cluster components.

Messages are initiated by activities in the HX storage cluster. They are distributed to assorted locations, including:

- VMware vCenter Events or Alarms pages—Some messages are directed to the vCenter Events and Alarms pages by the HX Data Platform. Some messages that are stored in HyperFlex log files are queried by vCenter and added to the vCenter Events and Alarm pages.

- HX Data Platform Plug-in Monitor > Events tab. See [View HX Data Platform Plug-in Events](#), on page 24.
- HX Data Platform Auto Support (ASUP) system—Auto Support must be enabled to send email notifications. Typically, Auto Support is configured during HX storage cluster creation by configuring the SMTP mail server and adding email recipients.
- Smart Call Home (SCH) notifications—SCH is an automated support capability that offers around-the-clock device monitoring, proactive diagnostics, real-time email alerts, service ticket notifications, and remediation recommendations for critical system events on your HX storage clusters.
- HX Connect user interface—In the header, the bell icon () displays an alarm count of your current errors or warnings. If there are both errors and warnings, the count shows the number of errors. For more detailed information, see the **Alarms** page or the **Events** page in the HX Connect user interface.

## View HX Data Platform Plug-in Events

The Monitor Events tab displays information about the state changes of the HX storage cluster. Events include user actions and system actions that occur on the HX storage cluster, hosts, or datastores. For example, adding a node to the HX storage cluster, removing a node from the HX storage cluster, or reconfiguring a VM resource.

You can perform the following tasks in the Events tab:

- Select an event to display the event details at the bottom of the tab.
- Use the filter controls above the list to filter the list. For example, type `memory` to display a subset of events.
- Click a column heading to sort the list.

From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex > Systems > Cisco HX Data Platform > cluster > Monitor > Events**. You have the following fields:

Field	Description
Description	Event message content. See the section for each event type.
Type	Type of message.
Date Time	Timestamp when the event occurred. Time is displayed in local browser time, calculated from UTC.
Target	Name of the target. Target type options include: storage cluster, host, datastore, or disk.
User	Consumer of the resource for the event.
VC Cluster Events	Link to vSphere storage cluster Events.

Field	Description
Event Detail	Same content for the event as the Event table. <b>Target link</b> - The Target object in the Event detail links to the vSphere target Summary page. For example, the HX storage cluster Summary page or node Summary page.

## Automatically Acknowledged Events

Storage cluster-wide events have an auto-acknowledge feature. When an event occurs, a message is published. When the condition for the message is corrected, for example, a node returns to healthy, the original alarm is acknowledged. Alarms can also be manually acknowledged.

The following events are automatically acknowledged when the corresponding event occurs.

Triggering Event	Resetting Event	Alarm Reset to
ClusterHealthCriticalEvent	ClusterHealthNormalEvent	Green
ClusterPolicyComplianceDegradedEvent	ClusterPolicyComplianceSatisfiedEvent	Green
ClusterPolicyComplianceFailedEvent	ClusterPolicyComplianceSatisfiedEvent	Green
ClusterPolicyComplianceImprovedEvent	ClusterPolicyComplianceSatisfiedEvent	Green
ClusterReadOnlyEvent	ClusterOnlineEvent	Green
ClusterUnhealthyEvent	ClusterHealthNormalEvent	Green
CriticalInfraComponentEvent	ClusterHealthNormalEvent	Green
FileSystemUsageAlertEvent	FileSystemUsageNormalEvent	Green
HighMemoryUsageAlertEvent	NormalMemoryUsageEvent	Green
HXHostUnreachable	HXHostReachable	Green
HXSshAccessToggleActionFailed	HXSshAccessToggleActionSuccess	Green
HXSyslogServerNotReachableWarningEvent	HXSyslogServerReachableEvent	Green
LowSystemMemoryAlertEvent	SystemMemoryNormalEvent	Green
NtpServerOfflineEvent	NtpServerOnlineEvent	Green
SmartLicenseEvalExpiringEvent	SmartLicenseEvalExpiringEvent	Green
SmartLicenseFeatureNotInLicenseEvent	SmartLicenseEvalExpiringEvent	Green
SpaceAlertEvent	SpaceRecoveredEvent	Green
SpaceCriticalEvent	SpaceRecoveredEvent	Green

Triggering Event	Resetting Event	Alarm Reset to
SpaceWarningEvent	SpaceRecoveredEvent	Green
VCCconnectionDown	VCCconnectionUp	Green

## Manually Acknowledging Alarms Using HX Connect

The **Alarms** page displays a list of HX storage cluster and ESXi alarms raised by the HyperFlex Data Platform. Acknowledging an alarm lets other users know that you are taking ownership of the issue. After an alarm is acknowledged, its alarm actions are discontinued. However, the alarm is still visible in the system. Alarms are neither cleared, nor reset when acknowledged.



**Note** Starting with HX 4.0(2a), you must acknowledge and reset alarms from the HX Connect user interface and in vCenter. Resetting and acknowledging alarms from vCenter will not reflect in HX Connect and vice versa. For full context about the alarm, it is recommended that you view in HX Connect.

- Step 1** Log in to HX Connect.
- Step 2** In the menu, click **Alarms**.
- Step 3** Click the alarm you want to acknowledge, and then click **Acknowledge**.

vCenter recognizes a session with HX Connect, therefore system messages that originate with vCenter might indicate the session user; for example, **Acknowledged By** might list `com.springpath.sysmgmt.domain-c7`.

## Cluster Access Policy Compliance Events

### ClusterPolicyComplianceDegradedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster access compliance has degraded.		Event and Alarm			No	No	See the Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

### ClusterPolicyComplianceFailedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster access compliance has failed.		Event and Alarm		Yes	No	No	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

#### ClusterPolicyComplianceImprovedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster access compliance has improved.		Event and Alarm			No	No	None.

#### ClusterPolicyComplianceSatisfiedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster access compliance is satisfied.					No	No	None.

## Cluster Events

#### ClusterAddedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
New storage cluster has been added.	Online	Event			No	No	None

#### ClusterCapacityChangedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Effective physical capacity was changed with addition or removal of disks or nodes.	Online	Event		Yes	Yes	No	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**ClusterHealthCriticalEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
There are three simultaneous non-allocated HDD or SSD, or node failures.	Offline	Event and Alarm		Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**ClusterHealthNormalEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster returned to normal health state.  The system recovers from HDD or SSD, or node failure.  Or You replaced failed resources.	Online	Event		Yes	No	No	None

**ClusterOnlineEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
You created the storage cluster successfully.	Online	Event		Yes	No	No	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**ClusterReadOnlyEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster transitions to ReadOnly mode.	ReadOnly	Event and Alarm		Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.  If ASUP enabled, TAC ticket created.

**ClusterReadyEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster is online and serving I/Os. NFS datastores are reachable.	Online	Event			No	No	None

**ClusterShutdownEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
<p>The system triggered the Cluster Health Critical Event.</p> <p>The storage cluster is shutdown.</p> <p>All VMs are inaccessible.</p>	Offline	Event and Alarm	Yes	Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**ClusterUnhealthyEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster is unhealthy.	Online	Event and Alarm		Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**VmMetadataNotSyncedWithIntersight**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster {ClusterName} metadata is not synced with Intersight.	Online	Event and Alarm		Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

**VmMetadataSyncedWithIntersight**



Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster {ClusterName} metadata is synced with Intersight.	Online	Event and Alarm		Yes	Yes	Yes	See Cisco HyperFlex Data Platform Administration Guide for information about managing the HX storage cluster.

## Critical Infrastructure Events

### CriticalInfraComponentEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
Information message when controller VM fails or is shutdown.	Offline	Event and Alarm		Yes	Yes	Yes	See TAC.

## Disk Events

### DiskAddedEvent

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
You added a new disk.		Event		Event	No	No	None.

### DiskFailedEvent

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk failed.		Event and Alarm		Yes	Yes	No	Replace the disk.

**DiskHardBlacklistedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk has failed permanently.		Event and Alarm		Yes	Yes	Yes	

**DiskHealthEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disk health deteriorated.		Event	Yes	No	No	No	Customer to review the remaining life value or the disk and decide if it needs replacement. See the following Disk Life Values table for more details.

**Table 2: Disk Life Values**

Life Remaining %	Comments
0%	<ul style="list-style-type: none"> <li>• May be raised for any drive type - SSDs and HDDs.</li> <li>• The event happens when the internal SMART Health status of the drive reports a predictive failure or the drive fails to power up.</li> <li>• Drive needs Return Merchandise Authorization (RMA).</li> </ul>
1% to 14%	<ul style="list-style-type: none"> <li>• Applies to SSDs only.</li> <li>• SSD Wear Level reporting 14% or less life remaining.</li> <li>• Replacement drive needs to be ordered by Customer.</li> </ul> <p>Purchase a replacement SSD(s) to replace any drive(s) past the wear condition not covered by <a href="#">HW replacement contracts</a>.</p> <p><b>Note</b> It is possible for a SSD to wear out eventually and report 0% life remaining. However, the wearout may be identified by reviewing the nightly wear-level check logs; Look for logs and events from 14% remaining life onwards.</p>

**DiskPhysicalAddedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk was added to the storage cluster.		Event and Alarm		Yes	No	No	None.

**DiskPhysicalRemovedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk was removed from the storage cluster.		Event and Alarm		Yes	No	No	Replace the disk.

**DiskRepairingEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk is being repaired.		Event		Yes	No	No	

**DiskRepairSucceededEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk was successfully repaired.		Event		Yes	No	No	

**DiskRemovalCompletedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk removal succeeded.		Event		Yes	No	No	

**DiskRemovalFailedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk removal failed.		Event		Yes	No	No	

**DiskRemovalInProgressEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A disk removal is in progress.		Event		Yes	No	No	

**DiskRemovedEvent**

Description Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
You either physically removed a disk or it became offline.		Event and Alarm		Yes	No	No	Replace the disk.

## Host Events

**HXHostUnreachable**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when HX HyperV host is unreachable.	Online	Event		Yes		No	No	None

**HXHostOffline**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when Get-ClusterNode command fails when FOC is down	Online	Event		Yes		No	No	None

**HXUBRNotSupported**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when UBR is unsupported version or quorum configuration for the fail over cluster is using SMB access point.	Online	Event		Yes		No	No	None

## Node Events

**NodeDecommissionedEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node removed from storage cluster.		Events		Yes	No	No	

**NodeJoinedEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A new node joined the storage cluster.		Event		Yes	No	No	

**NodeMaintenanceEnteredEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node entered Maintenance Mode.		Event		Yes	No	No	

**NodeMaintenanceExitedEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node exited Maintenance Mode.		Events		Yes	No	No	

**NodePoweredDownEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
You powered down a storage controller VM node. The storage cluster operates with reduced capacity or might be in a Degraded state.		Event		Yes	No	No	

**NodeReadyForIOEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node is ready for I/O operations.		Event		Yes	No	No	

**NodeReadyForShutdownEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node prepared for shutdown.				Yes	No	No	

**NodeCriticalEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cannot access the node due to I/O failure to the node.	Offline	Event and Alarm		Yes	No	No	

**ComputeNodeAddedEvent**

Description	Cluster State	Reported in vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Compute node added to HX cluster successfully				Yes	No	No	

# Encryption Events

**Cluster Level SED Events**

- **EncryptionLocalConfigDisableCompletedEvent**
- **EncryptionRemoteConfigDisableCompletedEvent**
- **EncryptionLocalConfigEnableCompletedEvent**
- **EncryptionRemoteConfigEnableCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption configuration on the local or remote cluster successful.	Online	Events				No	No	None

- **EncryptionLocalConfigDisableFailedEvent**
- **EncryptionRemoteConfigDisableFailedEvent**
- **EncryptionLocalConfigEnableFailedEvent**
- **EncryptionRemoteConfigEnableFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption configuration on the local or remote cluster failed.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

- **EncryptionLocalConfigDisableStartedEvent**
- **EncryptionRemoteConfigDisableStartedEvent**

- **EncryptionLocalConfigEnableStartedEvent**
- **EncryptionRemoteConfigEnableStartedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption configuration on the local or remote cluster started.	Online	Events		Yes		No	No	None.

- **EncryptionLocalConfigRekeyCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring Rekey of a disk on the local cluster successful.  Message for ConfigRekey task precedes the ReKey task for local cluster encryption.	Online	Events				No	No	None

- **EncryptionLocalConfigRekeyFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring ReKey of a disk on the local cluster failed.  Message for ConfigRekey task precedes the ReKey task for local cluster encryption.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

- **EncryptionLocalConfigRekeyStartedEvent**



Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring ReKey of a disk on the local cluster started.  Message for ConfigRekey task precedes the ReKey task for local cluster encryption.	Online	Events		Yes		No	No	None.

- **EncryptionLocalRekeyCompletedEvent**
- **EncryptionRemoteRekeyCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring Rekey of a disk on the local or remote cluster successful.	Online	Events				No	No	None

- **EncryptionLocalRekeyFailedEvent**
- **EncryptionRemoteRekeyFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring ReKey of a disk on the local or remote cluster failed.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

- **EncryptionLocalRekeyStartedEvent**
- **EncryptionRemoteRekeyStartedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Configuring ReKey of a disk on the local or remote cluster started.	Online	Events		Yes		No	No	None.

- **EncryptionLocalDisableCompletedEvent**
- **EncryptionLocalEnableCompletedEvent**
- **EncryptionRemoteDisableCompletedEvent**
- **EncryptionRemoteEnableCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption on the local or remote cluster successful.	Online	Events				No	No	None

- **EncryptionLocalDisableFailedEvent**
- **EncryptionLocalEnableFailedEvent**
- **EncryptionRemoteDisableFailedEvent**
- **EncryptionRemoteEnableFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption on the local or remote cluster failed.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

- **EncryptionLocalDisableStartedEvent**
- **EncryptionLocalEnableStartedEvent**
- **EncryptionRemoteDisableStartedEvent**
- **EncryptionRemoteEnableStartedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disabling or enabling encryption on the local or remote cluster started.	Online	Events		Yes		No	No	None.

- **EncryptionLocalRemoteCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption task on the local or remote cluster successful.	Online	Events				No	No	None

• **EncryptionLocalRemoteFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption task on the local or remote cluster failed.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

• **EncryptionLocalRemoteStartedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption task on the local or remote cluster started.	Online	Events		Yes		No	No	None.

• **EncryptionLocalRemoteConfigCompletedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption configuration on the local or remote cluster successful.	Online	Events				No	No	None

• **EncryptionLocalRemoteConfigFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption configuration task on the local or remote cluster failed.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

- **EncryptionLocalRemoteConfigStartedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption configuration task on the local or remote cluster started.	Online	Events		Yes		No	No	None.

- **EncryptionOperationTimeoutEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Encryption task on the local or remote cluster timed out.	Online	Events		Yes		No	No	Identify and correct the issue and retry the task.

### Node Level SED Events

- **EncryptionNodeDisableStartEvent**

- **EncryptionNodeEnableStartEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disable or enable encryption on a node in the HX cluster started.	Online	Events				No	No	None.

- **EncryptionNodeDisableSuccessEvent**

- **EncryptionNodeEnableSuccessEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disable or enable encryption on a node in the HX cluster successful.	Online	Events				No	No	None

- **EncryptionNodeRekeyStartEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
ReKey a node in the HX cluster started.	Online	Events				No	No	None.

- EncryptionNodeRekeySuccessEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
ReKey a node in the HX cluster successful.	Online	Events				No	No	None

#### Disk Level SED Events

- EncryptionDiskDisableEvent
- EncryptionDiskEnableEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disable or enable encryption on a disk on a node in the HX cluster task.	Online	Events				No	No	None

- EncryptionDiskDisableFailedEvent
- EncryptionDiskEnableFailedEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Disable or enable encryption on a disk on a node in the HX cluster failed.	Online	Events				No	No	None

- EncryptionDiskEraseEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Erase an encrypted disk on a node in the HX cluster task.	Online	Events				No	No	None

- **EncryptionDiskEraseFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Erasing an encrypted disk on a node in the HX cluster failed.	Online	Events				No	No	None

- **EncryptionDiskRekeyEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
ReKey an encrypted disk on a node in the HX cluster task.	Online	Events				No	No	None

- **EncryptionDiskRekeyFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
ReKey an encrypted disk on a node in the HX cluster failed.	Online	Events				No	No	None

- **EncryptionDiskSecureDriveEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Secure an encrypted disk on a node in the HX cluster task.	Online	Events				No	No	None

- **EncryptionDiskSecureDriveFailedEvent**

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
An encryption secured disk on a node in the HX cluster failed.	Online	Events				No	No	None

- EncryptionDiskSecureDriveUnlockFailedEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Unlocking a failed encryption secured disk on a node in the HX cluster failed.	Online	Events				No	No	None

- EncryptionDiskUnlockForeignEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Unlocking an unknown disk on a node in the HX cluster task.	Online	Events				No	No	None

- EncryptionDiskUnlockForeignFailedEvent

Description	Cluster State	Reported in vCenter	Reported in HXDP Plug-in	Reported HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Unlocking an unknown disk on a node in the HX cluster failed.	Online	Events				No	No	None

## Replication and Recovery Events



**Note** Events that are reported in vCenter are listed in the respective *object* > **Events** page. For example, virtual machine events or resource pool events.

**AddedDpGroup**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine protection group created.	Online	Event		Yes		No	No	None

**AddedDpVm**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine protection initiated.	Online	Event		Yes		No	No	None

**CreateDpVmSnapshotFailed**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine snapshot creation failed.	Online	Event		Yes		No	No	None

**FailoverFailed**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when virtual machine recovery fails.	Online	Event		Yes		No	No	None

**RecoveryFailed**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine recovery failed.	Online	Event		Yes		No	No	None

**RecoveryInitiated**



Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Initiated recovery of Virtual Machine.	Online	Event		Yes		No	No	None

**RecoverySucceeded**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Successfully recovered Virtual Machine.	Online	Event		Yes		No	No	None

**ReplicationFailed**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine replication failed.	Online	Event		Yes		No	No	None

**ReplicationPeriodExceeded**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Virtual Machine replication did not complete in the specified interval.	Online	Event		Yes		No	No	None

**TestFailoverFailed**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when virtual machine test recovery fails.	Online	Event		Yes		No	No	None

## Security Events

**LockdownModeEnabledAlarm**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when lock down mode is enabled on one or more nodes.	Online	Event		Yes		No	No	None

**HXSyslogServerNotReachableWarningEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when syslog server is not reachable."	Online	Event		Yes		No	No	None

v

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported in HX Connect	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Triggered when controller VM SSH access toggle action failed.	Online	Event		Yes		No	No	None

## Resource Monitoring Events and Alarms

This table describes the possible event severities in increasing order of severity:

Severity Level	Description
info	A basic notification or informational message that is possibly independently insignificant.
warning	A potential or impending service-affecting fault that currently has no significant effects in the system. An action should be taken to further diagnose, if necessary, and correct the problem to prevent it from becoming a more serious service-affecting fault.
error	A nonservice-affecting fault condition that requires corrective action to prevent a more serious fault from occurring. If the issue can not be resolved, contact support at <a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a> .
critical	A service-affecting condition that requires immediate corrective action. If the issue can not be resolved, contact support at <a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a> .

Name/Description	Cluster State	Reported to vCenter	Shown on HX Connect UI as Event	Send ASUP	Send SCH	SCH report results in an SR	Action
<b>Info Severity Level</b>							
<b>CoreFileInfoEvent-HX</b> Controller VM has one or more core files		Yes	Yes	Yes	No	No	No user action required
<b>Warning Severity Level</b>							
<b>FileSystemUsageAlertEvent</b> Not enough free capacity on one or more partitions of HX Controller VM		Yes	Yes	Yes	No	No	Login to the Controller VM and check the File System usage.
<b>NtpClockDriftEvent-HX</b> Controller VM clock not in sync with configured NTP servers		Yes	Yes	Yes	No	No	Confirm whether your NTP source is reachable from the Controller VM or from the Storage cluster. Run the <b>show ntp</b> command to determine the issue.

Name/Description	Cluster State	Reported to vCenter	Shown on HX Connect UI as Event	Send ASUP	Send SCH	SCH report results in an SR	Action
<b>NtpClockOfflineEvent</b> -One or more NTP servers configured on HX controller VM are not responding		Yes	Yes	Yes	No	No	<p>Confirm whether your NTP source is reachable from the Controller VM or from the Storage cluster. Run the <b>show ntp</b> command to determine the issue.</p> <p><b>Note</b> If this alarm persists for more than 24 hours then contact support.</p>
<b>VCConnectionDown</b> -Node has lost connection to vCenter		Yes	Yes	Yes	No	No	<p>Confirm whether the vCenter is reachable from the Controller VM.</p> <p><b>Note</b> If this alarm persists for more than 24 hours then contact support.</p>
<b>Error Severity Level</b>							

Name/Description	Cluster State	Reported to vCenter	Shown on HX Connect UI as Event	Send ASUP	Send SCH	SCH report results in an SR	Action
<b>ClusterUnhealthyEvent</b> -Cluster is unhealthy		Yes	Yes	Yes	Yes	Yes	See list of possible alarms.
<b>Critical Severity Level</b>							
<b>NodeCriticalEvent</b> -Cannot access the node due to I/O failure to the node		Yes	Yes	Yes	No	No	Requires immediate corrective action. Contact support for assistance to resolve the issue.

#### HX Data Platform Alarms

Name/Description	Alarm Code	Message	Severity	Send SCH	SCH Generates an SR
<b>DiskRemovedEvent</b> -Triggered when a disk is removed.	HXA-DSK-0001	A disk {UUID} was removed. {DISKDETAILS}	Warning	No	No
<b>DiskFailedEvent</b> -Triggered when a disk fails in a host.	HXA-DSK-0002	A disk failed. {DISKDETAILS}	Warning	Yes	No
<b>DiskPhysicalRemovedEvent</b> -Triggered when a physical disk is removed in a host.	HXA-DSK-0003	Physical disk {UUID} ({DISKPATH}) on controller host {HOSTNAME} at slot {DISKSLOT} was removed. {DISKDETAILS}	Warning	No	No

Name/Description	Alarm Code	Message	Severity	Send SCH	SCH Generates an SR
<b>DiskHardBlacklistedEvent</b> -Triggered when disk in a HX cluster is in a hard black list and cannot be recovered.	HXA-DSK-0004	{DISKMEDIUM} Disk {DISKPATH} on node {HOSTNAME} has failed permanently. {DISKDETAILS}	Warning	Yes	Yes
<b>DiskHealthEvent</b> -Triggered when a disk deteriorates in a cluster.	HXA-DSK-0005	Disk health has deteriorated. SSD with UUID: {UUID} and S/N: {SN} on host {HOSTNAME} ({IPADDRESS}) has {LIFE_LEFT}% remaining life. SSD must be replaced. {DISKDETAILS}	Warning	No	No
<b>VCConnectionDown</b> -Triggered when node has lost connection to VCenter.	HXA-NOD0012	{NAME} has lost connection to VCenter.	Warning	No	No
<b>ClusterUnhealthyEvent</b> -Triggered when the health status of the HX cluster is unhealthy. This alarm is reset to green when the cluster health status is back to normal.	HXA-CLU-0002	Cluster [{NAME}] is unhealthy.	Critical	Yes	Yes
<b>NodeCriticalEvent</b> -Triggered due to I/O access failure to the node.	HXA-NOD0001	Cannot access the node due to I/O failure to the node.	Critical	No	No

## Smart Call Home Events

### CallhomeEndpointConnectionFailedEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node is unable to reach Smart Call Home endpoint.	Online	Event			No	No	None

### CallhomeEndpointConnectionOKEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node is able to reach Smart Call Home endpoint.	Online	Event			No	No	None

## Smart License Events

### SmartLicenseEvalEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Smart License for the cluster is in Evaluation Mode.	Online	Event			No	No	None

### SmartLicenseEvalExpiringEvent

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Smart License Evaluation period is expiring soon	Online	Event		Yes	No	No	Check the system information page in HX Connect for details on license expiry register the cluster with Cisco Smart Software Manager by going to <a href="#">Cisco Software Central</a> and clicking “Smart Software Licensing”

**SmartLicenseEvalExpiredEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Smart License Evaluation period for the cluster has expired.	Online	Event		Yes	No	No	Register the cluster with Cisco Smart Software Manager by going to <a href="#">Cisco Software Central</a> and clicking “Smart Software Licensing”

**SmartLicenseFeatureNotInLicenseEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
The feature used in the cluster isn't available in the HX license tier.	Online	Event		Yes	No	No	The feature will reset when the license meets compliance.



**SmartLicenseInComplianceEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
The number of licenses allocated to the cluster is compliant with Smart Licensing.	Online	Event			No	No	None.

**SmartLicenseOutOfComplianceEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Insufficient number of Smart Licenses in the Virtual Account makes the cluster non compliant.	Online	Event		Yes	No	No	Purchase or transfer applicable licenses to your Cisco Smart Software Manager Virtual Account by going to <a href="#">Cisco Software Central</a> and clicking “Smart Software Licensing”

## Snapshot Events

**ScheduledSnapshotRPREachedMaxLimitEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A VM resource pool scheduled snapshot failed.  Scheduled snapshot failed since the virtual machine has reached maximum number of supported snapshots.		Event and Alarm			No	No	Adjust the retention policy on scheduled snapshot or delete user created snapshots.

**ScheduledSnapshotVMFolderReachedMaxLimitEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A VM folder scheduled snapshot failed.  Scheduled snapshot failed since the Virtual Machine has reached maximum snapshots supported.		Event and Alarm			No	No	Adjust the retention policy of scheduled snapshot or delete user created snapshots.

**ScheduledSnapshotVMReachedMaxLimitEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
A VM scheduled snapshot failed.  Scheduled snapshot failed since the virtual machine has reached maximum number of supported snapshots.		Event and Alarm			No	No	Adjust the retention policy on scheduled snapshot or delete user created snapshots.

## Space Utilization Events

This section describes the messages that appear related to the space used by the storage cluster.




---

**Note** When you add more nodes to the storage cluster, the HX Data Platform plug-in does not immediately reflect the storage cluster capacity.

---

**SpaceAlertEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
<p>Space capacity usage remains at error level. This alert is issued after storage capacity has been reduced, but is still above the warning threshold.</p> <p>The overall cluster storage capacity consumed is beyond the healthy threshold.</p> <p>The storage cluster is online and can perform write operations.</p> <p>HX Summary Health bar: Orange</p>	Online	Event and Alarm	Event	Yes	Yes	Yes	<p>Add storage or delete files.</p> <p>Continue to reduce the amount of storage capacity used, until it is below the warning threshold.</p>

**SpaceCriticalEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
<p>Space capacity usage is at error level.</p> <p>The cluster storage is full and offline.</p> <p>Your system has exceeded the functional amount of allocated capacity.</p> <p>The storage cluster transitions to read-only mode. It does not accept any write operations.</p> <p>This is a critical threshold. The system triggers a Critical alarm in vCenter.</p> <p>HX Summary Health bar: Red</p>	Read Only	Event and Alarm	Event	Yes	Yes	No	<p>Add storage or delete files.</p> <p>Reduce the amount of storage capacity used to below the warning threshold.</p> <p>Reclaimed space released after deleting files within few hours of deletion, based on system workload.</p> <p>The storage controller VM monitors and reclaims space until the system becomes operational (transitions out of read-only state).</p>

**SpaceRecoveredEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
Used capacity is below the warning threshold. Usage is back to normal.  HX Summary Health bar: Blue	Online		Event	Yes	No	No	None.

**SpaceWarningEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
Space capacity usage is at error level. This is the first level warning you are approaching a critical state.  The overall cluster storage capacity consumed is beyond the healthy threshold.  This is a warning threshold. The system triggers an Error alarm in vCenter.  The storage cluster can perform write operations.  HX Summary Health bar: Yellow	Online	Alarm	Event	Yes	Yes	No	Add storage or delete files.  Do not continue the storage cluster operations until you reduce the amount of storage capacity used to below this warning threshold.

**VirtualSpaceWarnEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
Cluster virtual usage is above warning threshold.		Event and Alarm		Yes	Yes	No	

**VirtualSpaceWarnClearEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH report results in an SR	Action
Cluster virtual space usage is back to normal.		Event		Yes	No	No	

## Storage Events

**StorageApdTimeoutEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
On an ESX server and APD state was entered and was not exited within defined number of seconds.				Yes	Yes	Yes	

## Upgrade Events

**ClusterUpgradeCompletedEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
The storage cluster upgraded successfully.	Online	Event		Yes	No	No	

**ClusterUpgradeFailedEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
One or more nodes in the storage cluster failed to upgrade.	Offline	Event		Yes	No	No	See TAC.

**ClusterUpgradeStartedEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Cluster upgrade to version <i>&lt;version&gt;</i> started.	Online	Event		Yes	No	No	

**NodeUpgradeCompletedEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node upgrade to version <i>&lt;version&gt;</i> for <i>node_ID</i> completed successfully. <i>node_ID</i> is the ID of the node that appears in the message.	Online	Event		Yes	No	No	

**NodeUpgradeFailedEvent**

Trigger	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node upgrade for <i>node_ID</i> failed. <i>node_ID</i> is the ID of the node that appears in the message.	Degraded	Event and Alarm		Yes	No	No	See TAC.

**NodeUpgradeStartedEvent**

Description	Cluster State	Reported vCenter	Reported HXDP Plug-in	Reported ASUP	Reported SCH	SCH Generates an SR	Action
Node upgrade to version <i>&lt;version&gt;</i> started.	Degraded	Event		Yes	No	No	



## CHAPTER 5

# Troubleshooting Topics

---

- [Install and Upgrade Issues, on page 61](#)
- [Host Issues, on page 73](#)
- [Disk Issues, on page 77](#)
- [VM Issues, on page 79](#)
- [Datastore Issues, on page 84](#)
- [ReadyClone, Snapshot, and Replication Issues, on page 89](#)
- [Cluster Issues, on page 96](#)
- [Interface Issues, on page 99](#)

## Install and Upgrade Issues

### Deploy IP Addresses Page Lists Duplicate Servers

#### Description

During HX Data Platform deploy, the IP addresses page lists the same servers twice.

#### Action: Select only one from pair

This might occur when UCS Manager configuration is skipped and HX Data Platform references both UCS Manager and the imported JSON file. Select only one of each pair of IP addresses.

### Installation Fails when Manually Reboot FIs

#### Description

Installation fails when FIs are manually rebooted during deploy.

#### Action: Reboot the HX Data Platform Installer

- 
- Step 1** Reboot the HX Data Platform installer VM.
- Step 2** Restart the deployment.
-

## During UCS Manager Only Upgrade, Controller VM Might Not Power On

### Description

During UCS Manager only upgrade, Controller VM might not power on after exiting the node out of maintenance mode.

### Action: Restart the EAM service on vCenter

VMware VCenter EAM service will not automatically power on the controller VM. The controller VM will be outside of the EAM resource pool.




---

**Note** Newly deployed HX clusters starting with HyperFlex Release 4.0(1a) no longer leverage the vSphere ESX Agent Manager (EAM) for the HyperFlex Storage Controller VMs. HX clusters built prior to HX 4.0(1a) will continue to utilize EAM. However, if that cluster is migrated to a new vCenter the EAM integration will not be configured. Users should contact TAC for help removing EAM.

---

- Restart the EAM service on the VCenter by running `/etc/init.d/vmware-eam restart`.

EAM should re-scan all the EAM agent VMs, and resolve all issues on these VMs, including powering on the Controller VM.

## Deploy or Upgrade Fail with Error: "NoneType" object has no attribute 'scsiLun'

### Description

Deploying or upgrade fails with error: "NoneType" object has no attribute 'scsiLun"

### Action: Disconnect and reconnect

This is a VMware issue. Disconnect the hosts from vCenter and reconnect them.




---

**Important** Do not remove the node from the cluster. This is a disconnect only.

---

## Upgrade Fails to Enter Maintenance Mode

### Description

Upgrade fails because a node failed to enter maintenance mode.

### Actions: Restart vmware-vmtoolsd service

If all other validations are successful then this might be a VMware issue, where the VMware VPXD crashed.

**Step 1** Ensure that VPXD restarted, and if not, manually restart it from the ESX command line.

```
# service vmware-vmtoolsd status
# service vmware-vmtoolsd start
```

**Step 2** Retry the upgrade.



Enter Maintenance Mode should succeed.

---

## Upgrade Fails at vMotion Compatibility Validation

### Description

Retry upgrade fails at validation on vMotion compatibility.

### Action: Rescan storage system from host

This is due to a sync issue between vCenter and ESXi.

Rescan the storage system on the ESX host using a vCenter client.

See VMware article, *Perform Storage Rescan in the vSphere Client*, at

<https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.hostclient.doc/GUID-FA49E8EF-A3DC-46B8-AA5B-051C80762642.html>

## Upgrade VM Power On Error: No compatible host was found

### Description

While attempting an upgrade, VM fails to power on with error: No compatible host was found

### Action: Manually power on the VM

---

**Step 1** From ESX command line, power on the VM.

**Step 2** Using controller VM command line, run command

```
# stcli cleaner start
```

---

## During Upgrade when two Nodes Fail, Controller VMs Power On Fail

### Description

If, during an upgrade, two nodes fail, the upgrade fails because controller VMs do not power on.

### Action: Restart EAM Service

---

**Step 1** Restart the vCenter EAM service.

From the ESX command line:

```
# /etc/init.d/vmware-eam restart
```

**Step 2** Proceed with the upgrade.

---

## Upgrade with Pre-6.5 vCenter Groups Some Controller VMs

### Description

After upgrading HX Data Platform using a vCenter version older than 6.5, some controller VMs are listed in a resource pool labeled, ESX Agents.

### Action: None required

No action is required. There is no functionality impact. All the virtual machines, including controller VMs are EAM registered and remain in the HX Cluster. All HX Cluster operations work as expected.

If you need to perform group operations, from the vCenter interface, drag and drop the controller VMs into ESX Agents resource pools.

## A Node Fails to Upgrade due to vCenter Issues

### Description

Sometimes during an online upgrade the vCenter daemon crashes on a node. When this happens, the node cannot enter HX maintenance mode. Without entering HX maintenance mode, the node cannot complete the upgrade. All other nodes, with properly functioning vCenter, complete the upgrade.

### Action: Re-run the Upgrade on the Affected Node

---

**Step 1** Correct the vCenter issue.

**Step 2** Restart the upgrade from any node in the cluster.

HX Data Platform skips any node that is already upgraded and moves on to complete the upgrade on any node that was previously missed.

---

## HX Data Platform Installer Shows Host Managed by Different vCenter

### Description

HX Data Platform Installer shows that a host is managed by a different vCenter.

When a host removed from vCenter, typically this removes the managementServerIP from the host summary information.

If host services were not running when the host was removed, vCenter continues to show the host.

### Action: Reboot vCenter

Reboot vCenter and the host should no longer be listed with vCenter.

## Configuration Settings Differ Between HX Data Platform and UCS Manager

### Description

During the installation, upgrade, and expand the storage cluster processes, HX Data Platform installer verifies the configuration settings entered with the settings in UCS Manager. Mismatches can occur, for example, in the following scenarios:

- Sometimes by the time the task is ready to apply the validations and configurations, a previously unassociated server is no longer unassociated. These servers need to be disassociated.
- You are using servers that were previously associated with an HX Data Platform storage cluster. These servers need to be disassociated.
- Manually entered existing storage cluster configuration information is prone to errors. The information, such as VLAN IDs and LAN configuration, needs to match the information listed in UCS Manager. Use previously saved configuration file to import the configuration.

**Action: Import existing configuration**

When installation, upgrade, or expand the storage cluster is completed, an option to Save Configuration is available. Use this option to save the cluster configuration information, then import the file with the saved configuration details when you need to make changes to the storage cluster.

**Action: Disassociate the Server**

See the Cisco HyperFlex Systems Getting Started Guide for steps on unassociating a server through UCS Manager. Briefly, the steps include:

- 
- Step 1** From the UCS Manager, select **Associated tab > node > Actions > Disassociate Server**.
  - Step 2** Confirm the node is disassociating, select **Unassociated tab > node > Assoc State**. The transition state is **removing**.
  - Step 3** Confirm the node completes the disassociation. Wait until the **Assoc State** is **none**. Do not select a node that has an Assoc State, removing.
- 

## Cluster Creation Fails with DNS Error due to FQDN

**Description**

Sometimes when the Fully Qualified Domain Name (FQDN) is provide to identify objects in the storage cluster, cluster creation fails. This is typically because the Domain Name Service (DNS) server specified is unavailable.

This applies to all possible domain name objects that are entered for any HX Data Platform Installer object that is identified by the domain name or IP address. This can include: vCenter server, ESX servers, controller VM addresses, storage cluster management or data network addresses, DNS servers, NTP servers, mail servers, or SSO servers.

**Action: Verify the DNS server**

- 
- Step 1** Login to the command line of the HX Data Platform Installer VM. For example, use `ssh`.
  - Step 2** Verify the DN servers provided, work.
    - a) Identify the DNS server IP addresses.

```
# nslookup <dns_server>
```

b) Repeat for every DNS server configured for use by the storage cluster.

**Step 3** Verify each object needed to create the cluster can be resolved from the provided DNS server. These objects are provided through either a JSON file or the HX DP Installer GUI fields.

a) Identify the cluster objects' IP addresses.

```
# nslookup <object> <dns_server>
```

<object> is the FQDN or IP address of all the possible objects you provide to the HX Data Platform Installer for cluster configuration: vCenter server, ESX servers, controller VM addresses, storage cluster management or data network addresses, DNS servers, NTP servers, mail servers, or SSO servers.

b) Repeat for every cluster configuration object.

**Step 4** If either Step 2 or Step 3 cannot be verified, then use IP addresses only instead of Fully Qualified Domain Names (FQDN) in the HX Data Platform Installer GUI.

## Offline Upgrade Cluster Start Command Error: Node Not Available

### Description

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`.

**Action: Manually power on the controller VMs, then start the storage cluster.**

**Step 1** Manually power on the controller VMs.

- a) Login to the vSphere Web Client.
- b) Locate the controller VMs that are not powered on.

From the Navigator select, **vCenter Inventory Lists > Virtual Machines > vm**.

Storage controller VMs, have the prefix, `stCtlVM`.

- c) From the right-click or Actions menu select, **Power > Power On**.
- d) Repeat until all the storage controller VMs are powered on.

**Step 2** Restart the storage cluster.

- a) Login to the command line of any controller VM.
- b) Run the command.

```
# stcli cluster start
```

## vSphere Replication Plug-in Fails after HX Plug-in Deployed

### Description

This error occurs when the vSphere Replication plug-in is installed after the HX Data Platform plug-in. Recommended order is to install the vSphere Replication plug-in first, then install the HX Data Platform plug-in.

**Action: Unregister the HX Data Platform plug-in**

This task removes the HX extensions from the vCenter Managed Object Browser (MOB).

**Before you begin**

1. Remove the vSphere Replication plug-in from the vCenter MOB.
2. Remove the vSphere Replication virtual machine from the vCenter inventory.
3. Remove the HX vCenter cluster from the vCenter datacenter.

---

**Step 1** Download the vSphere ESX Agent Manager SDK, if you have not already done so.

**Step 2** Remove the HyperFlex cluster object from vCenter.

**Step 3** Login to the vCenter server MOB extension manager.

- a) In a browser, enter the path and command.

```
https://vcenter_server/mob/moid=ExtensionManager&doPath=extensionList
```

- b) Enter the login credentials.

**Step 4** From the vCenter server MOB extension manager, view the mob and the extension associated with removed cluster.

- a) Locate `rootFolder` and click the `(Datacenters)` link.

From the Data Object Type: ServiceContent page, scroll through the Name column, click the link in the Value column.

- b) Locate the `childEntity` and click the `(datacenter_name)` link.

From the Managed Object Type: ManagedObjectReference:Folder page, scroll through the Name column, click the link in the Value column.

- c) Locate the `hostFolder` and click the `(host)` link.

From the Managed Object Type: ManagedObjectReference:Datacenter page, scroll through the Name column, click the link in the Value column.

- d) Locate the `childEntity` and note the corresponding Value for `(datacenter_name)`. This is the domain ID of the cluster to be unregistered.

From the Managed Object Type: ManagedObjectReference:Folder page, scroll through the Name column, click the link in the Value column.

**Step 5** Unregister the extension from the ExtensionManager page.

- a) In a browser, enter the path and command.

```
https://vcenter_server/mob/moid=ExtensionManager&method=unregisterExtension
```

- b) Enter the extension key Value or `(datacenter_name_link)`, and click **Invoke Method**.

**Step 6** If the removed cluster was the CIP that vCenter used for communicating with the HX Data Platform plug-in, restart the vsphere-client services.

- a) From the vCenter server MOB extension manager, view the mob and the extension associated with removed cluster.  
b) Locate `extensionManager` and click the `ExtensionManager` link.

From the Data Object Type: ServiceContent page, scroll through the Name column, click the link in the Value column.

- c) Locate the `extensionList["com.springpath.sysmgmt"]` link.

From the Managed Object Type: `ManagedObjectReference:ExtensionManager` page, scroll through and click link in the Value column. Click `(more...)` if needed to show the full list.

- d) Locate `server` and click the `server` link.

From the Data Object Type: `Extension` page, scroll through the Name column, click the link in the Value column.

- e) Locate a URL that ends with `/plugins/stGui-1.0.zip`. This is the cip that is used for the HX Data Platform plug-in. For example, `"https://cs002-cip-m.eng.storvisor.com/plugins/stGui-1.0.zip"`

From the Data Object Type: `ExtensionServerInfo[]` page, scroll through the list of line items to locate the Name: url, Type: string, and Value with `/plugins/stGui-1.0.zip`.

**Step 7** If the cip located in the previous Step is associated to the cluster that you removed from vCenter, the extension needs to be cleaned up.

- Login to vCenter using `ssh`.
- Clean up the HX Data Platform plug-in extension folder. This might happen if another cluster is running and older, yet compatible, version of the HX Data Platform plug-in.
- Restart the vSphere services. Run the command:

```
# service vsphere-client restart
```

**Step 8** Log out of all sessions and log back in.

### What to do next

- Recreate the datacenter cluster. Add the hosts to the HX vCenter cluster, one at a time.
- Re-register the vSphere Replication virtual machine from the datastore.
- From the vSphere Replication appliance web front end, recreate the vSphere Replication plug-in. Verify the vSphere Replication plug-in is available in vCenter.
- From the HX Data Platform Installer, reinstall the HX Data Platform plug-in and recreate the storage cluster.

## Upgrade Fails but Reports All Nodes Up-to-Date

### Description

This is due to `RemoteException` sent by vCenter and is most likely due to intermittent network connectivity between the HX storage cluster and vCenter.

**Action:** Retry the upgrade

## Restarting Online Upgrade Fails

### Description

In some rare cases, restarting online upgrade on a HX storage cluster with the previous upgrade in failed state might fail again. Even though the HX cluster recovered from failure and is in a healthy state.

**Action: Retry the upgrade again**

When retrying upgrade using CLI, please use `-f` or `--force` option command `stcli cluster upgrade` or use HX Data Platform Plug-in to retry the upgrade.

## Controller VM Fails to Power On During Cisco UCS Upgrade

**Description**

Sometimes when vSphere is exiting maintenance mode, all the VMs on the server do not power on. This can include the storage controller VM.

**Action: Manually restart the controller VM**

This is a known VMware issue. For more information, see VMware KB article - [Auto-Start Is Not Run When Manually Restarting a Host in Maintenance Mode](#).

## Firmware Upgrade Fails from Server Storage Controller with Unsupported Board

**Description**

Upgrading the UCS firmware failed. Possible reason due to use of an unsupported board in the HX server.

**Action: Decommission then recommission the board.**

- 
- Step 1** Decommission and then recommission the referenced board
  - Step 2** Verify that the server is healthy.
  - Step 3** Retry the firmware upgrade.
  - Step 4** If this does not resolve the issue, contact Cisco TAC for more assistance.
- 

## A Node Fails to Upgrade due to vCenter Issues

**Description**

Sometimes during an online upgrade the vCenter daemon crashes on a node. When this happens, the node cannot enter HX maintenance mode. Without entering HX maintenance mode, the node cannot complete the upgrade. All other nodes, with properly functioning vCenter, complete the upgrade.

**Action: Re-run the Upgrade on the Affected Node**

- 
- Step 1** Correct the vCenter issue.
  - Step 2** Restart the upgrade from any node in the cluster.  
HX Data Platform skips any node that is already upgraded and moves on to complete the upgrade on any node that was previously missed.
-

## Upgrade Stalls Waiting for Node to Return to Healthy State

### Description

If your LSI version is older than version 9, sometimes the disks are not found during an upgrade on the node. If the node is not healthy, the upgrade cannot proceed.

LSI version 9 is associated with UCS firmware version 2.2(6f) and 2.2(7c).

**Action: Reboot the node manually.**

**Step 1** Login to the controller VM command line. For example, using `ssh`.

**Step 2** Verify the disks are showing. Run the `lsscsi` command.

```
# lsscsi
[2:0:0:0]   disk      ATA          INTEL SSDSC2BB12 CS01 /dev/sdb
[2:0:1:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sdc
[2:0:2:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sdd
[2:0:3:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sde
[2:0:4:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sdf
[2:0:5:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sdg
[2:0:6:0]   disk      SEAGATE     ST1200MM0088    N003 /dev/sdh
[2:0:7:0]   disk      ATA          INTEL SSDSC2BX48 CS01 /dev/sdi
[3:0:0:0]   disk      VMware     Virtual disk     1.0  /dev/sda
```

**Step 3** Reboot the node manually.

## Cluster Expansion Error: No cluster found

### Description

From the HX Data Platform Expand Cluster wizard, the HX storage cluster was not discovered.

**Action: Manually enter cluster IP address**

Manually enter the HX storage cluster management IP address in the Management IP Address field of the Expand Cluster wizard.

To locate the cluster IP address:

**Step 1** From the vSphere Web Client, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform**.

**Step 2** Click-select the storage cluster name. From the **Action Menu** at the top of the panel, select **Summary**.

**Step 3** Locate the Cluster Management IP Address in the Summary display.

## Cluster Expansion Fails when DNS Server is Not Found

### Description



Expanding a storage cluster requires a DNS server, even if you specify the new node using an IP address and are not using a FQDN. The HX Data Platform Installer checks for any DNS servers that were provided during cluster creation.

- If any of the previously provided DNS servers are not reachable, cluster expansion fails.
- If you did not specify a DNS server when you installed HX Data Platform, cluster expansion fails.

If either of these conditions apply, perform the corrective action.

**Action: Identify and Provide Correct DNS Servers**

**Step 1** Login to the command line of any HX controller VM. For example, use `ssh`.

**Step 2** Identify and DNS servers configured for the storage cluster.

```
# stcli services dns show
```

Sample Response

```
10.64.1.8
```

```
10.64.1.9
```

If no DNS addresses are listed, skip to Step 4.

**Step 3** Remove all DNS servers that are no longer available to the storage cluster.

```
# stcli services dns remove --dns <dns_server>
```

**Step 4** Add any DNS servers that are new to the storage cluster.

If you did not specify a DNS server when you created the storage cluster, add a fake DNS server.

```
# stcli services dns add --dns <dns_server>
```

**Step 5** Verify each object needed to create the cluster can be resolved from the provided DNS server. These objects are provided through either a JSON file or the HX DP Installer GUI fields.

a) Identify the cluster objects' IP addresses.

```
# nslookup <object> <dns_server>
```

*<object>* is the FQDN or IP address of all the possible objects you provide to the HX Data Platform Installer for cluster configuration: vCenter server, ESX servers, controller VM addresses, storage cluster management or data network addresses, DNS servers, NTP servers, mail servers, or SSO servers.

b) Repeat for every cluster configuration object.

c) If an object fails to resolve, add an entry to the DNS server file.

**Step 6** Verify the DN servers provided, work.

a) Identify the DNS server IP addresses.

```
# nslookup <dns_server>
```

b) Repeat for every DNS server configured for use by the storage cluster.

**Step 7** Repeat Step 5 and Step 6 to verify every added DNS server is valid and every HXDP object can be resolved through each DNS server.

**Step 8** Return to the HX Data Platform Installer and proceed with the storage cluster expansion.

---

## Expand Cluster Fails with Stale HX Installer

### Description

An expand cluster node is added to an incorrect cluster. This happens when the same HX Data Platform Installer is used to create multiple clusters, then that same HX DP installer is used to expand one of the clusters. The HX DP installer defaults to adding the node to the most recent cluster.

**Action: Redeploy HX Data Platform Installer OVA**

---

**Step 1** Redeploy the HX Data Platform Installer OVA.

**Step 2** Use the new HX Data Platform Installer to expand the cluster.

---

## Installation Fails when Secure Boot is Enabled

### Description

Installation fails when ESXi is redeployed on nodes with secure boot enabled. During install or expansion if the nodes are a mix of ESXi version 7.0 U2 and earlier then the hypervisor configuration phase may fail with

the following

The screenshot displays a 'Hypervisor Configuration' interface. On the left, under 'Hypervisor Configuration - Overall', there is a 'Failed' status indicator. The task list includes:

- ❌ Configuring static ip on the specified ESXi servers
- ❌ Fetching OS info of nodes using ucs tools
- ❌ Error occurred while configuring UCS
- ❌ Reason: Please upgrade ESXi on node sys/rack/unit-8, sys/rack/unit-9 and sys/rack/unit-10 to 7.0U2 or above and retry.
- ❌ Configuration Failed
- ✅ Login to UCS API
- ✅ Cleanup SPT annotations
- ✅ Inventorying org of specified servers
- ✅ Inventorying physical servers
- ✅ Logout from UCS API

On the right side, configuration details are shown:

- Admin User name: root
- Server Selection**
  - Server 8: WZP22121357 / HXAF240C-M55X
  - Server 9: WZP221213T6 / HXAF240C-M55X
  - Server 11: WZP2212139E / HXAF240C-M55X
- UCSM Configuration**
  - VLAN Name: hx-inband-mgmt
  - VLAN ID: 242
  - VLAN Name: hx-storage-data
  - VLAN ID: 342

An 'Edit Configuration' button is located at the bottom right of the configuration panel.

**Action**

Upgrade all nodes in your cluster to ESXi version 7.0 U2 or higher and retry.

# Host Issues

## Post Manual ESX Installation statsd restart

**Description**

After manually re-installing ESX on your HX Data Platform servers, reset the stats daemon to ensure performance statistics display correctly.

**Action: restart stats daemon**

- 
- Step 1** Login to the command line of the controller VM of the ESX host.
- Step 2** Run the restart command.
- ```
# /etc/init.d/statsd restart
```
- Step 3** Repeat Step 1 and Step2 on the controller VM of every ESX host in the storage cluster.
- 

## scvmclient Management Services Restarted when services.sh restart Issued

**Description**

Running `services.sh restart` restarted `scvmclient` management services.




---

**Caution** Running this command will cause the HX datastores to disconnect from the particular host.

---

1. Put the node into maintenance mode.
2. Login to the ESX command line.
3. Restart the services.
 

```
# services.sh restart
```
4. Restart the ESX host daemon, vCenter Agent services, and the controller VM.
 

```
# /etc/init.d/hostd restart
```

```
# /etc/init.d/vpxa restart
```

## ESX Server Reboot on Upgrade Exited Maintenance Mode

**Description**

A power reset of the ESX server during upgrade exited upgrade and put the server into maintenance mode.

**Action: Manually exit Maintenance Mode**

Manually exit the server from maintenance mode and upgrade will continue.

## EAM Did Not Start on Compute Node

**Description**

EAM did not automatically restart on a compute node.

**Action: Manually restart EAM**

## Remove Node Fails when a Node is Down.

### Description

Remove node is not allowed if only 3 nodes are up.

### Action: Add a replacement node first

Replacing a node in a 3 node cluster requires TAC assistance. If failed nodes reduce the cluster to 3 nodes, replacing the node requires TAC assistance.

## Rebooting an HA Enabled ESX Host

### Description

If you enable HA on a host in the storage cluster that the system cannot access, then when you reboot the ESX host, the storage controller VM is powered off.

This is an artifact of interactions between how VMware handles HA failures and ESX Agent Manager (EAM) configurations. It might cause storage controller VMs to not power on after recovery.

### Action: Power on the storage controller VM on an HA enabled ESX host

---

**Step 1** Reconfigure HA on the host on which it initially failed.

**Step 2** Manually power on the storage controller VM.

---

## Node Failure While Simultaneously Adding Another Node to Cluster

### Description

When you add a node to an existing storage cluster, the storage cluster continues to have the same HA resiliency as the originating storage cluster until a rebalance completes.

For example, if you have a 3 node storage cluster and you add 2 converged nodes to the storage cluster. Until rebalance completes, the storage cluster behaves as a 3 node storage cluster and not a 5 node storage cluster. So, if a node fails before rebalance completes, the storage cluster status is degraded.




---

**Note** Rebalance typically occurs under the following conditions:

- Per the 24 hour rebalance schedule, two hours after a node fails, or if the storage cluster is out of space.
  - When single disk usage exceeds 50% or disk cluster aggregate usage is less than 50%.
- 

### Action: Manually initiate storage cluster rebalance

---

**Step 1** From a storage controller VM command line.

```
# stcli rebalance start --force
```

**Step 2** To monitor rebalance status, use command.

```
# stcli rebalance status
```

## Configure PCI Passthrough After Changing vNIC or vHBAs

### Description

After vNIC or vHBA are manually added to a Cisco HyperFlex (HX) service profile or service profile template, the PCI devices are re-enumerated, and the VMware directpath I/O configuration is lost. When the service profile is changed, the host hardware is updated and the PCI passthrough must be reconfigured. Perform the following steps on each ESX host with a modified service profile

Perform the following steps on the storage controller VM of the modified ESX host:

**Action: Update the vSphere Service Profile on the ESX Host**

- 
- Step 1** Put the ESX host into HX Maintenance mode.
- Step 2** Make or confirm the changes, such as adding hardware, in the Service Profile.
- Step 3** Reboot the ESX host.
- This host loses the direct path configuration.
- Step 4** Log into vCenter and select the DirectPath I/O Configuration page.
- From vCenter Client: Select the *ESX host* > **Configuration tab** > **Hardware pane** > **Advanced Settings** > **Edit**.
- From vCenter Web Client: From the **vCenter Inventory**, select **Resources** > **Hosts** > *ESX host* > **Manage** > **Settings** > **Hardware** > **PCI Devices** > **Edit**.
- Step 5** Select the LSI card for passthrough.
- From the DirectPath I/O Configuration page, select **Configure Passthrough**.
  - From the Mark devices for passthrough list, select the LSI card for the pass through.
  - Click **OK**.
- Step 6** Reboot the ESX host.
- Step 7** Re-map the PCI device to the HX storage controller VM (StCtlVM), by editing the storage controller VM settings.
- Locate and remove the unknown PCI Device.
- From vCenter Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **PCI device 0** > **Remove** > **OK**.
- From vCenter Web Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **Remove PCI device 0** > **OK**.
- Locate and re-add the LSI Logic PCI device.
- From vCenter Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **Add** > **PCI Device** > **LSI Logic PCI device** > **OK**.
- From vCenter Web Client: Right-click the *HX storage controller VM*, select **Edit Settings** > **PCI Device** > **Add** > **LSI Logic PCI device** > **OK**.
- Step 8** Remove the ESX host from HX Maintenance mode.

When the host is active again, the HX storage controller VM properly boots and rejoins the storage cluster.

---

## Secure Boot Cannot Be Enabled After Upgrade

This is a corner case for secure boot enablement. It may be necessary to manually remove and reinstall certain vSphere Installation Bundles (VIBs) in version ESXi 7.0 if originally upgrading from an ESXi 6.0 release (then to 6.5/6.7 before 7.0). To enable Secure Boot without interruption to the ESXi network after an upgrade:

---

**Step 1** Login to the node and uninstall NENIC driver.

**Note** Do not reboot the node.

**Step 2** Re-install the VIB (e.g. Cisco\_bootbank\_nenic\_1.0.33.0-1OEM.670.0.0.8169922.vib).

**Step 3** Repeat the first two steps on all the nodes and reboot nodes in a rolling fashion.

**Step 4** Enable Secure Boot. For detailed steps, see [Enabling Secure Boot Mode](#).

---

## Disk Issues

### Cannot Allocate Data when All or Most Disks Fail on a Node in a 3 Node Storage Cluster

#### Description

If all the hard disks on a node fail, HX Data Platform cannot allocate data to the node. If this occurs in a 3 node storage cluster, HX Data Platform cannot maintain the three minimum copies of data required to maintain data integrity. The result is a virtual ENOSPC state.

If up to a few hard disks on a node fail, the storage cluster becomes imbalanced as it attempts to write to the node and consumes the space on the remaining disks. For example, if there were 10 HDDs on all three nodes, and 9 of the HDDs fail on the third node, the imbalance results as the disk on the third node restricts the cluster size to 10% of the actual cluster size. This is a physical ENOSPC state. This might also cause an all paths down (APD) state.

**Action: Physically balance the storage on all the nodes in the storage cluster**

---

**Step 1** Replace the damaged disks with good disks.

**Step 2** Add another node to the storage cluster.

**Step 3** Adjust the storage capacity of the nodes to ensure they match.

**Step 4** If the storage cluster does not automatically recover, you might need to restart the storage cluster.

- a) Login to the command line of the storage controller VM on any node in the cluster.
- b) Shutdown the storage cluster.

```
stcli cluster shutdown
```

- c) Restart the storage cluster.

```
stcli cluster start
```

---

## Removing Disks Causes Rebooting Storage Controller VM to Fail

### Description

If you removed disks, then rebooted the storage controller VM host before an automatic rescan could complete, the storage controller VM might not power on.

**Action: Power on the storage controller VM after removing disks**

---

**Step 1** Ensure storage controller VM is powered off.

**Step 2** Run the script.

```
/opt/springpath/support/rescanLuns.sh
```

**Step 3** Power on the storage controller VM.

---

## Cluster Management IP Fails after NVME Drive is Replaced

### Description

Occasionally when the NVME drive is replaced, the cluster service may not run.

**Action: Start the service with the `start cip-monitor` command.**

Use the `status cip-monitor` command to check the status of `cip-monitor`.

---

## Recover Failed SSD Hosting the Storage Controller VM

### Description

If the SSD hosting the storage controller VM fails, it must be recovered.

**Action: Recover the failed SSD**

---

**Step 1** Log in to the command line of the host with the failed SSD.

**Step 2** Verify the status of the SSD is dead timeout.

```
esxcli storage core device list -d SSD_ID | grep 'Status:'
```

```
Status: dead timeout
```

**Step 3** Kill the `vmx` of the storage controller VM.

```
ps | grep vmx | grep -i stCtIvm
```



```
kill -9 process_id_of_controller_vm
```

- Step 4** Re-scan the storage adapter.
- ```
esxcli storage core adapter rescan -a
```
- Step 5** Replace the disk with a new SSD of same specifications.
- Step 6** Restart the *hostd*.
- Step 7** Power on the storage controller VM.
- 

## How to synchronize SCVM clock after an installation

Follow the steps below to synchronize SCVM clock after an installation

1. Power off user VMs before stopping storfs.
2. On each node,run the following commands:

```
Stop storfs  
ntpsync -yfs.
```

3. Wait for all the nodes to synchronize with NTP.
4. Start storfs.

## vNode Scrubber Detection

### Description

The vNode scrubber periodically scans the storage disks to detect latent disk corruptions and repair bad blocks.

Use `log messages/stats/commands` to verify its operation and progress.

```
Starting Vnode Scrubber  
Vnode Scrubber Ends  
Vnode Scrubbing FT <FT Number>
```

### Action

The vNode scrubber is initiated once per month but can be disabled.

## VM Issues

### Controller VM Prevented from Powering On

#### Description

vSphere EAM was not able to power on the controller VM due to insufficient resources.

This occurs when vSphere HA is turned on and admission control is set to **Reserved failover capacity to be at 1 host**. With this setting, HA admission control reserves enough resources to ensure one host can be fully failed over.

**Action: Adjust vSphere Admission Control**

Review the VMware document, *Best Practices for Admission Control*, <https://docs.vmware.com/en/VMware-vSphere/6.0/com.vmware.vsphere.avail.doc/GUID-BD6D9434-84C8-4937-BC76-04852F5EA136.html>.

Make appropriate vSphere adjustments.

## Automatic vMotion Migrations Fail with Timeout Error

### Description

This has been observed in 16 + 16 node clusters. It is a known VMware issue. For more information, see VMware KB article - [vMotion of a VM fails with the error: "Timed out waiting for migration data" \(2143834\)](#).

**Action: Verify network connectivity**

## Storage Controller VM Power On Fails with Two Node Failures

### Description

This is due to a VMware EAM (ESX Agent Manager) issue. The EAM did not mark the VM on the host correctly.

**Action: Re-register the storage cluster**

Re-register the storage cluster to sync the vCenter view. From the storage controller VM, run the command:

```
# stcli cluster reregister
```

## Fail Adding VM to Host with HA and DRS Enabled

### Description

When a user VM power on fails with `All required agent virtual machines are not currently deployed on host 'hostname' and the controller VM on the same ESX host is down.`

HA does not allow VMs to be powered on in a host if any VM marked as agent, in this case, the controller VM, is not powered on.

If DRS places a user VM on this host, the VM does not power on.

**Action: Complete the steps.**

- 
- Step 1** From the vCenter cluster, change the DRS Automation setting to Manual.
  - Step 2** From the ESX host, power on the VM directly.
  - Step 3** In the Power On dialog box, select an ESX host that has a powered on controller VM.
-

## Degraded Performance on VM with Disk Limit Shares

### Description

Powered on VM with disk limit shares set, degrades the performance on the respective datastore.

**Action:** Per VMware expected behavior.

---

**Step 1** Disable the mclock scheduler.

**Step 2** Move to the default scheduler.

---

## DRS Migrates VMs when Storage Cluster in Read Only State

### Description

When a storage cluster is in Read Only state, the VMware DRS process migrates VMs to datastores even though they are also in the Read Only storage cluster. The result is the VMs do not start.

**Action:** Manually disable DRS when storage cluster in Read Only state.

---

**Step 1** Select the HX Data Platform storage cluster.

From the vSphere Web Client Navigator, select **vCenter Inventory Lists > vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster**.

**Step 2** Select Summary tab, then click the VC Cluster link to toggle to the **vCenter Summary** tab. Click **Manage > Services > vSphere DRS**. Click **Edit**, then uncheck **Turn ON vSphere DRS** and click **OK**.

---

## VM Power On Fails Due to Stale EAM Extension

### Description

If you have partially installed or uninstalled HX Data Platform, sometimes a stale ESX Agent Manager (EAM) for the HX Data Platform extension remains. This can prevent virtual machines from powering on after a completed HX Data Platform installation. Remove stale extensions using the Managed Object Browser (MOB) extension manager.

**Action:** Remove stale EAM HX Data Platform extension

---

**Step 1** Download the vSphere ESX Agent Manager SDK, if you have not already done so.

**Step 2** Remove the datacenter from your vSphere cluster.

**Step 3** Identify the HX cluster UUID.

Every agency has a field `cluster_domain_id` which refers to the underlying vSphere extension. This extension ID uses a Managed Object ID (moid).

From a storage controller VM command line, run the command:

```
# stcli cluster info | grep vCenterClusterId:
vCenterClusterId: domain-c26
```

**Step 4** To unregister the storage cluster extension: Log into the vCenter server MOB extension manager  
First unregister the HyperFlex cluster.

a) In a browser, enter the path and command.

```
https://vcenter_server/mob/?moid=ExtensionManager
```

*vcenter\_server* is the IP address of the vCenter where the storage cluster is currently registered.

b) Enter administrator login credentials.

**Step 5** Locate the HX storage cluster extensions with the cluster IDs. Scroll through the **Properties > extensionList** to locate the storage cluster extensions:

```
com.springpath.sysgmt.cluster_domain_id and com.springpath.sysgmt.uuid.cluster_domain_id.
```

Copy each of these strings into your clipboard. Exclude the double quotes (") on either end of string, if there are any.

**Step 6** Unregister each storage cluster extension.

a) From the **Methods** table click **UnregisterExtension**.

b) In the **UnregisterExtension** popup, enter an extension key value, `com.springpath.sysgmt.cluster_domain_id`.

For example: `com.springpath.sysgmt.domain-26`

c) Click **Invoke Method**.

**Step 7** To remove stale EAM extensions: Log into the vCenter server MOB ESX agencies extension manager.  
Second remove stale EAM extensions that were associated with the HyperFlex cluster.

a) In a browser, enter the path and command.

```
https://vcenter_server/eam/mob/
```

*vcenter\_server* is the IP address of the vCenter where the storage cluster is currently registered.

b) Enter administrator login credentials.

**Step 8** Locate the stale HX storage cluster ESX agency extensions with the cluster IDs.

a) Scroll through the **Properties > agency > Value**.

b) Click an agency value.

c) In the **Agency** window, check the **Properties > solutionID > Value** extension. Verify has the correct *cluster\_domain\_id*.

For example: `com.springpath.sysgmt.domain-26`

**Step 9** Remove stale ESX agency extensions.

a) From the **Agency** window, **Methods** table select a method.

Stale ESX agencies can be removed using either the `destroyAgency` or `uninstall`.

b) In the *method* popup, click **Invoke Method**.

**Step 10** Refresh the **ExtensionManager** tab and verify that the **extensionList** entry does not include `com.springpath.sysgmt.cluster_domain_id` extensions.

**Step 11** Restart the vSphere Client services.

The HX Data Platform extensions are removed when the vSphere Client services are restarted. Restarting the vSphere client service temporarily disables access to vCenter through the browser.

**Step 12** Run and complete the HX Data Platform installation.

For additional options on removing stale EAM extensions, see Technical Assistance Center (TAC).

---

## Deleting VM Folder or File Taking Very Long Time

### Description

If user VMs become inaccessible to vSphere and the ESX \*.lck files created for those user VMs remain in the file system, then deleting the VM files and/or folders might take a very long time.

**Action: Clear ESX server VM lock files**

---

**Step 1** Locate all the VM lock files in the storage cluster.

```
# cd /vmfs/volumes/my_datastore  
# find . -name .lck* | xargs -n1 rm
```

**Step 2** Retry deleting the VM file or folder.

---

## VM Disk Usage and vCenter VM Committed Size Mismatch

### Description

When a VM has snapshots or ReadyClones, the VM disk usage does not match vCenter's virtual machine committed size.

**Action: None**

## Migrating a VM Task Fails and Results in Replication Error 10001

### Description

You can run a maximum of 64 replication jobs concurrently. Where, 48 replication slots are reserved for protection and 16 replication slots are reserved for migration. If the replication slots reserved for protection are available to use, a maximum of 64 migration jobs would succeed to start sync replication.

When you initiate a migration of a VM, a new replication job is triggered. The new replication job could fail and result in a replication error (1001) in one of the following scenarios:

- When the data replication is configured with a large number of VMs with an aggressive replication interval and if the scheduled replication jobs are constantly running. If you initiate a new replication for protection and/or migration when the ongoing replications replicate 128 VMDKs, the new replication job will fail even though the slots are available.

- When migration of more than 16 VMs has been initiated at a time during the progress of 48 replications for protected VMs.

**Action:**

Re-initiate the migrate VM task when there are fewer than 64 replication jobs running in the system. You can also increase the replication interval of the existing replication schedules to create a window and then revert to the original replication interval after the migration task is complete.

## VM Migration Results in an Error

**Description**

Migrating a VM takes a new snapshot and will initiate a new replication task. If an existing replication task is currently in progress, and if another snapshot has not yet been replicated, the migration task fails and the following error message is displayed:

**PrepareFailover failed. Most recent not failed snapshot for VM: xx-xx-xx has not been replicated.**

**Action:**

Re-initiate migration of the VM when all ongoing replication tasks are complete. If there is a stale snapshot, delete it manually using the following command:

```
stcli dp vm snapshot delete --vmid xxx --snapshot-id xxx
```

## VM Migration BadVersionException Error

**Description**

After protecting VMs, while scheduler replication in-progress, executed migrate operation fails with a "BadVersionException" error.

**Action:**

If a "BadVersionException" error appears during VM migration, retry the migrate operation.

## Datastore Issues

### Removing Compute Node Did Not Remove Datastores

**Description**

Removing a compute node from the storage cluster did not remove the associated datastores.

**Action: Manually remove the datastores**

## Adding Multiple Datastores Error: mountpoint not found on host

### Description

Due to a VMware issue, adding multiple datastores at the same time sometimes does not mount all the datastores.

### Action: Remount datastore

- Mount fewer datastores at once.

Use the HX plug-in to remount any datastores that are not initially mounted.

## NFS All Paths Down with Message File Locked by Consumer on Host

### Description

Due to a VMware issue with Serial I/O Control (SIOC), an NFS all paths down (APD) is seen with a message containing the following:

```
NFSLock: 2210: File is being locked by a consumer on host host_name with exclusive lock.
```

### Action: Toggle Storage I/O Control

---

**Step 1** From vCenter, datastore view, select *datastore* > **Configuration** > **Properties**.

**Step 2** Toggle **Storage I/O Control** to its opposite state.

If it is enabled, then disable it. If it is disabled, then enable it.

**Step 3** Return **Storage I/O Control** to its original state.

If it is enabled, then disable it. If it is disabled, then enable it.

**Step 4** Verify the NFS lock is removed.

---

## Maximum Queue Depth Limit of NFS is not Optimal

### Description

When you upgrade Hyperflex cluster from Hyperflex 2.5 or previous releases, the maximum queue depth limit of NFS datastore per host is set to 256 by default. This setting might lead to performance issues.



---

**Note** This section is not applicable for VDI specific deployments.

---

### Action

For every node where the queue depth is less than 1024, execute the following procedure to check the maximum NFS Queue Depth and increase the maximum limit if needed:

```
esxcli system settings advanced list -o /NFS/MaxQueueDepth  
maxQDepth:256 <- Low value
```

Place the node in the Hyperflex Maintenance mode and run the following command. Reboot the node and ESXi host for the following changes to take effect.

```
esxcli system settings advanced set -o /NFS/MaxQueueDepth -i 1024
esxcli system settings advanced list -o /NFS/MaxQueueDepth
maxQDepth:1024 <- Optimal value
```

## Mounting Datastore Fails After Changing VLAN ID

### Description

If you change the VLAN ID after you created the storage cluster, mounting a datastore to the storage cluster fails. Also existing datastores can become unmounted from the storage cluster.

**Action: Reload the ESX server firewall.**

See your VMware ESX documentation for directions on reloading an ESX server firewall.

## Datastore fails to mount if data vSwitch has an existing vmkernel port configured with an IP address

### Description

Per VMware configuration requirements, duplicate IP addresses or duplicate rules result in loss of connectivity.

**Action: Ensure that your traffic is utilizing the intended VMkernel interface.**

Configure the following:

- Only one VMkernel port per IP subnet.  
If you are using vSphere 5.x, then this does not apply for iSCSI multi-pathing, or multi-NIC vMotion.
- A dedicated non-routable VLAN or dedicated physical switch for vMotion purposes.
- A dedicated non-routable VLAN or dedicated physical switch for IP Storage purposes.
- A dedicated non-routable VLAN or dedicated physical switch for Fault Tolerance purposes.

## Datastore Mount Fails after Motherboard Updated with ESXi 7.0 U3 Node

### Description

If the datastore mount fails due to the Motherboard update with an ESXi version 7.0 U3 node then perform the following workaround:

### Action

```
[root@ucs2479:~] vi /etc/hosts
[root@ucs2479:~] /bin/auto-backup.sh
Files /etc/vmware/dvsdata.db and /tmp/auto-backup.2110399//etc/vmware/dvsdata.db differ
Saving current state in /bootbank
Creating ConfigStore Backup
Locking esx.conf
Creating archive
Unlocked esx.conf
```



```
Using key ID
62baf055-3b7c-4584-b4b5-a412a020484c to encrypt
Clock updated.
Time: 05:51:06   Date: 12/14/2021   UTC
```

## Remounting Datastore after Cluster Restart

### Description

Sometimes after a storage cluster returns to a healthy state, existing datastores might not be automatically remounted. This could happen when the storage cluster is rebooted while one or more nodes are down, or when it takes a long time for the storage cluster to reboot.

### Action: Mount the datastore.

Choose a method:

- Using the HX Data Platform plug-in.
- Using the command line.

- 
- Step 1** Using the HX Data Platform plug-in.
- a) Select From the vSphere Web Client Navigator, select **vCenter Inventory Lists > Cisco HyperFlex Systems > Cisco HX Data Platform > cluster > Manage > Datastores.**
  - b) Select a **datastore.**
  - c) Click the mount datastore icon or select from the Actions menu.
  - d) Confirm to mount the datastore, click **OK.**
- Step 2** Using the command line.
- a) Login to a storage controller VM of a node in the storage cluster.
  - b) Run the command.
- ```
# stcli datastore mount datastore_name
```

---

## Datastore Does Not Unmount When Storage I/O RM Daemon Running

### Description

If the VMware Storage I/O RM feature is enabled, there is a VMware bug where-in Storage I/O RM writes tracking files even on datastores on which Storage I/O RM is not enabled. These tracking files prevent HX Data Platform datastores from being unmounted.

### Action: Retry unmount.

- 
- Step 1** Retry unmount datastore.
- Step 2** Stop the Storage I/O RM daemon on every ESX host on which the HX Data Platform datastore is mounted.
- Step 3** Unmount the datastore.

**Step 4** Restart the Storage I/O RM daemon, if needed.

---

## Datastore Delete Fails with error: com.vmware.vim25.PlatformConfigFault

### Description

There is a known VMware issue where vSphere selects random datastores and uses it for heartbeating. This blocks HX Data Platform attempts to delete the datastore. See VMware KB, [Unmounting or removing a datastore in a HA cluster fails with the error: The vSphere HA agent on host failed to quiesce file activity on datastore \(2055137\)](#).

**Action: Check the ESXi host and vCenter, then retry deleting the datastore.**

---

**Step 1** Ensure VMs are not running on the datastore.

**Step 2** From ESX host, check to see if the HX Data Platform datastore is being used by VMware service, **storageRM**.

```
# ls -ltra /vmfs/volumes/stfs-dsl/ | grep -i iorm
-rwxr-xr-x 1 root root 16511 Jan 20 20:05 .iormstats.sf
drwxr-xr-x 1 root root 1125 Jan 20 20:06 .iorm.sf
```

**Step 3** Check the **storagerm** status.

```
# /etc/init.d/storageRM status

storageRM is running
```

**Step 4** Stop the **storagerm** service.

```
# /etc/init.d/storageRM stop
watchdog-storageRM: Terminating watchdog process with PID 34096
storageRM stopped
```

**Step 5** Alternatively, disable vSphere HA.

- a) From vSphere client Home, select **vCenter Inventory Lists > Resources > Clusters > storage cluster > Manage > Settings > Services**.
- b) Click **vSphere HA**. Click **Edit**.
- c) Unselect **Turn on vSphere HA**, if it is selected.

**Step 6** Try to delete the datastore again.

**Step 7** If you disabled vSphere HA, re-enable it.

**Step 8** This is one possible solution, if this doesn't resolve the issue, contact Technical Assistance Center (TAC).

---

## Datastore Not Accessible Due to Clock Skew

### Description

In a Hyper-V cluster, when a couple of converged nodes are shutdown for a few hours and then powered back on, the cluster may show healthy, but the converged nodes show as not accessible to the HX datastore in HX Connect.

If you run the command: **ntpd -gq**, the controller comes back with a message indicating:

```
no ntp servers found
```

This indicates that the NTP is not functional between the controller and Windows DC being used as an NTP server.

**Action:** Edit the `ntp.conf` file, and add `"tos maxdist 30"`. For more information, see [Synchronizing ESXi/ESX time with a Microsoft Domain Controller](#).

## Datastore not synced results in errors during Disaster Recovery

### Description

When the datastore is running out of space, the following error appears: "Module MonitorLoop Power on Fail". For more information, see ["Module MonitorLoop power on failed error when powering on VM on vSphere"](#).

When the datastore capacity is increased, the free space may not be synced leading to different errors in DR-operations, for example, the following error appears "A specified parameter was not correct: path".

**Action:** These need to be resolved on the VC/ESXi environment for example, by running a Rescan Storage from the vSphere Client for each of the ESXi hosts to make sure that the capacity is correctly reflected. For more information, see ["VMware ESXi datastore free space not updating after volume size increase from the vCenter Server"](#).

## ReadyClone, Snapshot, and Replication Issues

### Replication Fails when using VMware Tools to Quiesce VMs

#### Description

Sometimes replication fails when the option to **Use VMware Tools to quiesce the virtual machine** is selected.

Sometimes replication fails if the VM is undergoing any change in its layout or other guest tool related activity when the replication starts. These changes include: VMDK addition; HX Native Snapshot or Redolog Snapshot creation, deletion, or consolidation; VM reconfiguration; or vMotion. This is a transient state.



**Note** VM disks total count limit is 128 for all the VMs which are replicating concurrently.

Maximum number of VMs that can replicate for protection at any given time concurrently is 48.

If replication fails, it is automatically retried. If replication fails after multiple attempts, the failure might not be due to a transient VM layout change. Consider unselecting the **Use VMware Tools to quiesce the virtual machine** option.

**Action:** Unselect `Use VMware Tools to quiesce the virtual machine option`

**Step 1** Login to HX Connect.

**Step 2** Select **Edit Schedule** for the protection method used for the virtual machine.

- Protected independently—Select **Replication > Protected Virtual Machines > vm\_name > Edit Schedule**.

- Protected independently—Select **Replication** > **Protection Groups** > *protection\_group* > **Edit Schedule**.

**Step 3**    Unclick Use VMware Tools to quiesce the virtual machine and click **Save Changes**.

---

## Replication Errors with VMware Guest Provider (quiesce)

### Description

For all scheduled down-time, all VMs must be powered-down or shut-down gracefully and brought back up.

In case of any unexpected Storage and File system events such as:

- All paths down (APD) or PANIC
- Latency
- Deliberate/On-demand Cluster/File System Shutdown or Start

I/O will be lost to storage, guest and guest service usually have time-outs and go into read-only or read-write mode.

### Action: Retry recovery

After the storage is back up and running, determine the state of the guest file system and perform the associated recovery action:

- **Read-Write** - If the guest system is in read-write mode the services running inside the guest will likely be out-of-sync. Restart the VMware guest tools service.
- **Read-Only** - If guest system is in read-only mode, reboot of the VM.

## Reprotect Operation Fails on a Recovered VM when Another Replication is In-progress on the Same VM

The reprotect operation fails on a recovered VM, when another replication is in-progress on the same VM.

**Action: Retry the reprotect operation when replication is not in-progress for the VM.**

---

**Step 1**    Reprotection of a VM operation from Target to Source when Replication is in-progress from Source to Target is not supported, and results in an error.

**Step 2**    Retry the re-protect operation on Target, when replication is not in-progress for the VM from Source.

---

## Migrate Task Fails When Target Datastore Is Renamed

### Description

When the target datastore of a protected VM is renamed from the vCenter and migrate is performed on the VM, the migrate task fails at recovery with the error response:

Failover: Failover failed for VM: KD-3. Reason: Clone files failed with error 10031 for <VMID>

#### Action

Retry migrate on the same VM.



**Note** If the latest snapshot was created as scheduled after renaming the datastore, migrate will work without any issue.

## Backup Workflows Fail with an error message

### Description

In some instances, during a snapshot workflow or a backup workflow, delta disks are not accurately cleaned. At this point, the workflow cannot execute completely and fails. The following error message is also displayed:

```
Non-native snapshots exists - Error
```

### Action:

Delete the snapshots that point to the delta disks, or delete all snapshots and re-initiate the backup workflow.

## SRM Recovery Fails with vSphere Cluster Service (vCLS)

### Description

Recovery operations fail due to an incompatibility between vSphere 7.0 U1 feature vSphere Cluster Service (vCLS) and Protect the vCLS agent VMs with SRM 8.3.1

### Action: Move vCLS VMs to a non-replicated datastore

In the event SRM recovery fails. Verify that the vCLS VMs in the cluster are placed on a non-replicated datastore. If they are not, perform the following tasks:

1. Storage vmotion the VM(s) to a non-replicated datastore.
2. Unprotect the vCLS VMs using either the **stcli** or **webcli** command. This will relinquish the storage occupied by snapshots for these VMs on both source and target datastores.
3. Run the **stcli dp vm list -brief** command, to get a list of the biosuids for all VMs protected for disaster recovery.
4. Run the **stcli dp vm delete -vmid <vmbiosuuid>** command on every vCLS VM that was storage vmotioned out of the HX replicated datastores.

For more information on how the Site Recovery Manager interacts with vSphere Cluster Services, see <https://docs.vmware.com/en/Site-Recovery-Manager/8.4/com.vmware.srm.admin.doc/GUID-531FB787-8B30-401B-A921-C15C21D0BAA2.html>.

## Backup Software Fails when SSLv3 Disabled

### Description

Backup software sometimes fails when SSLv3 is disabled due to VMware bugs on vSphere 5.5 and 6.0 u1.

**Action:** See VMware KB articles

---

Click the links to the related VMware articles.

- vSphere 6.0u1, See Enabling support for SSLv3 in ESXi (2121021) at [http://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2121021](http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2121021).
  - vSphere 5.5, See Enabling support for SSLv3 on vSphere 5.5 (2139396) at [https://kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2139396](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2139396).
- 

## Recovery Fails after Renaming Datastore

### Description

If a recovery is performed immediately after renaming a target datastore, the recovery might fail with error:  
Datastore not found.

**Action:** Retry recovery

Wait a few minutes after renaming the datastore, then re-try the recovery.

## Recover, Migrate, or Test Recovery fail when Silent Mode notification is enabled

### Description

If the recover, migrate, or test recovery operations fail in conjunction with the silent mode notification, review the HX Connect activity messages within the recovery tab. If the activity messages indicate a failure due to cleanup of an existing VM, performing the following steps may resolve the issue:

**Action:** Manually Identify and Deleted the VM that Remains from a Prior Test Recovery Operation

- 
- Step 1** Manually inspect the destination HyperFlex cluster using HX Connect or the VMware vSphere client user interface. Search for an existing VM with the same name as the VM on which the operation is being attempted. If the VM is located, it may exist because of a prior recovery operation that has not been cleaned up.
- Step 2** Manually delete the VM from disk on the destination cluster using the VMware vSphere client.  
Exercise caution, make sure you have identified the correct VM.
- Step 3** Retry the recover, migrate, or test recovery operation.
- 

## Rebooting the Node Stops Recovery, Restore, and Clone

### Description

Rebooting the node while running the following commands, stops the command action and causes subsequent retries of the commands to fail.

- `stcli dp vm recover`
- `stcli dp vm restore`
- `stcli dp vm clone`

**Action: Clean up the virtual machine**

---

**Step 1** Clean up the virtual machine from vCenter.

Choose an appropriate option.

- If the virtual machine is registered, remove it from vCenter.
- If the virtual machine is not registered, delete the VM folder from the datastore.

**Step 2** Rerun the `stcli dp vm` command.

---

## Rerunning stcli vm recover Command Fails

**Description**

An initial attempt to recover a virtual machine, using `stcli vm recover`, did not complete. Rerunning the command also does not complete.

The partial recovery might have left the virtual machine registered with vCenter. This needs to be removed.

**Action: Clean up residual files**

---

**Step 1** Delete the virtual machine from vCenter.

**Step 2** Delete the `/vmfs/volumes/<volume>/<vmname>` directory.

**Step 3** Rerun the recovery command.

```
# stcli vm recover
```

---

## REST API Protection Group Filter Returns All Groups

**Description**

When using the REST API to locate a protection group, filtered REST calls returns to all protection groups.

**Action: None**

The filter parameters, name and type, for `groups:get` are not supported.

## VM Stunned During Snapshot Consolidation

### Description

If you have any redo log snapshots for VMs in the HX storage cluster, edit the ESXi host configuration where the redo log snapshots reside. If this step is not completed, VMs might be stunned during snapshot consolidation.

Redo log snapshots are snapshots that are created through the VMware Snapshot feature and not through the HX Data Platform Snapshot feature.

**Action:** Set `snapshot.asyncConsolidate="TRUE"` on ESXi host

- 
- Step 1** Login to the ESXi host command line
- Step 2** Locate and open for editing the file, `/etc/vmware/config`
- Step 3** Set the `snapshot.asyncConsolidate` parameter to `TRUE`.
- ```
snapshot.asyncConsolidate="TRUE"
```
- 

## Native Snapshots with Quiesce Option

### Description

Native Snapshots with Quiesce option, on Windows 2008 or Windows 2012 server, when the VMs are powered on, is not supported.

**Action:** Use non-quiesce option

Power off the VM, then take the snapshot or use the non-quiesce default option.

## Quiesce Based Snapshots Fail without an Error Message

### Description

When a quiesced based snapshot fails without a standard error message in HX Connect or vCenter.

**Action:**

Review the `vmware.log`. If you observe `VIX_E_TOOLS_NOT_RUNNIN` in the `vmware.log`, run the VMware tools using vSphere client and retry the operation.

## vMotion Fails to Move a Native Snapshot Datastore

### Description

vMotioning a native snapshot fails to move the associated datastore. Though using vMotion for native snapshot VMs is supported, selecting storage vMotion only is not supported on VMs with native snapshots.

**Action:** vMotion original VM only

If the VM needs to be moved to a different datastore, delete the snapshots from the source datastore, then vMotion the original VM.



## Reprotect option is unavailable for VMs in Protecting State

Reprotect option is unavailable for VMs in protecting state.

**Action: Unprotect and protect the VMs again.**

---

**Step 1** Select VM (Local VMs at site B), click **unprotect**. A message displays that shows the virtual machine as unprotected and that it has also disappeared local for secondary site, that is, remote in primary site and appeared in the Virtual machine section at site B.

**Step 2** Select VM and protect again.

A full copy of the VM to the target site can be accomplished using this method. A copy of the VM exists at target site.

---

## Site Recovery Manager Failover and Reprotect Issues

### Failover

Running Site Recovery Manager (SRM) recovery, with approximately 750 VMs, could fail for a few of the VMs if the VM are not accessible or disconnected from VC during the run. The SRM error **Prepare failover failed for VM <VM name>** may occur.

If SRM recovery fails, perform the following steps:

1. Unregister the VM from placeholder datastore at the protected site.
2. Re-register the VM from source datastore at the protected site.
3. Restart the Failover **PrepareFailover** operation.



---

**Note** It is recommended that when using 700 or more VMs, you should deploy the VMs on 2 or more other DRO (SRA) datastore pairs so that there are multiple protection groups.

---

### Reprotect

Running Site Recovery Manager (SRM) recovery, with approximately 750 VMs, and after implementing the steps for **Failover**, reprotect may fail with the error **Protection group PG2-BTOA has protected VMs with placeholders which need to be repaired**. This error is due to SRM having affected/faulty VMs in its protection group but the protection can not be configured. The unprotected VMs need to be removed from SRM protection group to proceed.

1. Navigate to **SRM > Protection Groups**, select the protection group and **Virtual Machines**.
2. Locate the affected VM and select **Remove protection** to remove the VM from protection group.
3. Re-run the reprotect procedure.

# Cluster Issues

## After Cluster reregister Controller VMs not EAM Agents

### Description

Controller VMs are not listed as EAM agents after `stcli cluster reregister`.

**Action: recreate the cluster**

---

**Step 1** Delete the vCenter cluster.

**Step 2** Recreate the vCenter cluster.

**Step 3** Reregister the HX cluster.

```
# stcli cluster reregister
```

---

## Cluster Becomes Unhealthy after Multiple Reregisters

### Description

Performing multiple cluster re-registers might cause the cluster to become unhealthy.

**Action: Recreate the cluster**

The HX Cluster lost the vCenter information, the virtCluster and HX Connect status indicates the cluster offline. However, the HX Data Platform cluster indicates it was healthy throughout.

Recreate the cluster.

```
# stcli cluster recreate
```

## ClusterNotConfigured Error after Node Removed

### Description

After removing a node from a cluster, on one of the controller VMs, the `stcli cluster info` command lists `ClusterNotConfigured`.

**Action: Refresh the cluster**

From the controller VM command line, run:

```
# stcli cluster refresh
```

## Cluster Capacity Higher than Individual Disks

### Description

Total Cluster usage shown might be higher than the usage shown for individual disks.

For example, cluster usage can be 80%, yet highest utilized disk might show only 76% usage.

**Action: None**

The difference can be attributed to management layer handling. Use the cluster usage value to make all utilization related decisions.

## Re-registering a Cluster Does Not Re-register Compute Nodes with EAM

### Description

This can occur in a variety of scenarios. Possible scenarios include:

#### Scenario 1

1. Start from an older HX version, prior to 2.1.x.
2. Add a compute node.
3. Re-register the cluster.
4. Upgrade the cluster. Task fails to include the compute nodes.

#### Scenario 2

1. Start from an older HX version, prior to 2.1.x.
2. Add a compute node.
3. Upgrade the cluster. Task completes.
4. Re-register the cluster. Task fails at the EAM level.

#### Scenario 3

1. Start with a new HX version, 2.1.x or later.
2. Add a compute node.
3. Re-register the cluster. Task fails at the EAM level.

**Action: Remove compute nodes before re-register**

- 
- Step 1** vMotion any VMs off of the compute nodes and remove the compute nodes from the HX cluster.
- Step 2** Re-register the HX cluster.
- Step 3** Add the compute nodes to the HX cluster.
- 

## Latency Spikes Seen for Workloads with Large Working Sets

### Description

Large working set workloads require accessing data from the capacity tier. As of HX Data Platform version 2.1(1b) backend access is optimized to significantly reduce the magnitude and frequency of high latency spikes.

- For hybrid clusters – When this symptom is present, the upgrade requires a longer maintenance window. Also, the default upgrade process does not automatically enable this optimization. Contact Cisco TAC to enable this performance enhancement during the upgrade process.
- For All Flash clusters - The upgrade times are not significantly affected and the default upgrade path automatically enables this performance enhancement.

**Action: Upgrade to 2.1(1c) or greater**

## Cluster Health Status Remains Unhealthy after Rebalance

### Description

In any three node cluster, including ROBO storage clusters, a single node in maintenance mode or failure causes the cluster to become unhealthy. Rebalance does not correct this.

**Action: Return node to healthy state**

Check that a node or component within a node is not failing. The cluster remains unhealthy as long as a component or node is failed. When the component or node returns to a healthy state, the cluster recovers and becomes healthy again.

## NTP Not Configured on ESXi Hosts

### Description

Sometimes if the ESXi host is power cycled, fails, or enters and exits maintenance mode the NTP server does not sync.

**Action: Manually configure NTP on ESXi host**

---

Enable the NTP client.

- From vSphere Web Client, select *host* > **Manager** > **System** > **Time Configuration** > **User Network Time Protocol**.
  - From the **NTP Service Startup Policy** field, select **Start and stop with host**. Click **OK**.
  - Repeat for each ESXi host in the storage cluster.
- 

## Cluster Capacity Different Than Provisioned

### Description

Sometimes in the HX Data Platform plug-in, Cluster Capacity in the Summary tab and Provisioned in the Manage tab show a different amount of storage allocated to the storage cluster. This occurs under the following conditions.

- **Cleaner not completed yet.** VMs have been removed, but the cleaner has not been run. The cleaner is an automatic process, after it completes the Cluster Capacity and Provisioned amounts should match. See the Cisco HX Data Platform Command Line Interface Reference guide for information on the cleaner command.

- **Thick provisioning or Thick clones.** If thick disks or clones are created, then HX Data Platform does not reserve the space. A soft reservation is used and datastores show space used, but the space is not used in the storage cluster. This is by design to help administrators with not over-provisioning their datastore.

**Action:** None.

## Connectivity to Storage Controller VM when using vShield

### Description

vShield interferes with HX Data Platform activity. Installing vShield in the HX Data Platform cluster is not recommended.

### Action: Exclude selected HX components

If you need to install vShield, exclude the HX storage controller VMs and vCenter from vShield protection. See VMware vCloud Networking and Security documentation, at [https://www.vmware.com/support/pubs/vshield\\_pubs.html](https://www.vmware.com/support/pubs/vshield_pubs.html).

- 
- Step 1** Install the vShield Manager.
- Step 2** Exclude HyperFlex Storage Controller VM's and vCenter Server from the vShield App Protection.
- From vCenter select, **Host & Clusters > Settings & Reports > vShield App > Exclusion List > Add**. Select each controller VM, `stCtlVM<name>`.
- Step 3** Ensure network connectivity to the storage controllers (ping, ssh etc.).
- Step 4** Install and configure vShield components.
- Step 5** To validate the configuration is working, reboot all the ESXi hosts simultaneously to get the datastore offline. Then repeat step 3 after the system is back up.
- 

## Storage Cluster Missing from vCenter Cluster after Cluster Node Powered Off

### Description

A node in the vCenter cluster was powered off. The storage cluster is fine within number of tolerated down nodes. However, the storage cluster cannot be managed through vSphere.

Known VMware vSphere 6.0 bug. See <https://communities.vmware.com/thread/514117?start=0&tstart=0>.

### Action: Reset the node.

Power on the node or disconnect the powered off node from the vCenter cluster.

## Interface Issues

### Multiple VM Power Operations Causes Task Queue to Error Out

#### Description

Multiple VM power operations causes task queue to error out.

**Action: Clean the queue**

Power operations can be initiated through HX Connect, but are performed through vCenter. The maximum vCenter task collector is 32. This is not modifiable.

---

**Step 1** Clean the queued tasks.

See the related article, *VCS vSphere – Check new notifications stuck on Queued – VMware vCenter Update Manager Check Notification*, at <http://www.natestiller.com/2011/02/vcs-vsphere-check-new-notifications-stuck-on-queued-vmware-vcenter-update-manager-check-notification/>

**Step 2** Logout and log back in to HX Connect.

**Step 3** Retry the power operations.

Do not exceed 32 simultaneous operations.

---

## HX Connect Data Does Not Refresh

### Description

Sometimes the HX Connect status fields do not refresh the data displayed.

### Action: Clear Browser Cache

#### • Microsoft Internet Explorer

1. From your IE browser select, **Settings (gear) > Safety > Delete browsing history**
2. Click the appropriate checkboxes.
3. Click **Delete**.

#### • Google Chrome

1. From your Chrome browser select, **Menu (3 vertical dots) > More tools > Clear browsing data**
2. Click the appropriate checkboxes.
3. Click **CLEAR BROWSING DATA**.

#### • Mozilla Firefox

1. From your Firefox browser select, **Menu (3 vertical bars) > Options (gear) > Advanced > Network**
2. In the **Cached Web Content** section, click **Clear Now**.

## Performance Charts Show a Gap while Node Rebooting

### Description

Sometimes events, such as a node rebooting, on the HX cluster affect system performance. The performance charts might show a data gap for the duration of the event.

**Action: None**

When the event completes, the performance chart reporting continues.

## Cannot See the HX Data Platform Plug-in Through vSphere Clients

### Description

Cannot see Cisco HyperFlex Systems or Cisco HX Data Platform in vSphere client or web client. There are a few possible situations when this might occur. Complete the action appropriate to your situation.

**Action: Select an option**

- Restart vCenter Service after HX storage cluster creation
- Restart vCenter Service after an upgrade
- Restart vCenter Service after adding another cluster to a vCenter with an existing cluster
- Install latest Adobe FlashPlayer in Firefox browser

---

### Step 1

Restart vCenter Service.

- a) Login to the vCenter server command line.
- b) Restart the vCenter service.  

```
# ssh root@vc_address # service vsphere-client restart
```
- c) Wait for vCenter to restart. This typically takes a few minutes.
- d) Logout and re-login to vCenter to ensure the vCenter interface refreshes.

### Step 2

Install latest Adobe FlashPlayer in Firefox browser.

- a) View the Shockwave Flash version.  
From a Firefox browser address bar, enter `about:addons`.
  - b) Check the version, then download and install the latest Flash Player from <https://get.adobe.com/flashplayer/>.
  - c) View the Shockwave Flash version again.
  - d) If more than the latest Flash version is listed, disable the older versions.
  - e) Reload vSphere web client.
- 

## Performance Charts Display Not Formatted at 100%

### Description

The performance chart display is not formatted at 100% zoom.

Selecting an optional metric and a smaller resolution at the same time shows a chart that is not formatted correctly.

**Action: Change the zoom in the chart**

## HX Data Platform Plug-In Feature Not Performing

### Description

Sometimes this occurs after a new cluster is created on an existing vCenter that also has different versions of the HX Data Platform.

### Action: Cycle vSphere Login

Log out of the vSphere client, then log back in.