Troubleshooting Topics

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. If that cluster is migrated to a new vCenter, however, the EAM integration will not be configured. For more information about removing EAM dependencies for upgraded clusters, see the HyperFlex ESXi Agent Manager Removal Process tech note.

- 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-vpxd 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-vpxd status
# service vmware-vpxd start
```

Step 2 Retry the upgrade.

Enter Maintenance Mode should succeed.

Upgrade Fails at vMotion Compatability 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 if **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.

a) Login to vCenter using `ssh`.
b) 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.
c) 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

1. Recreate the datacenter cluster. Add the hosts to the HX vCenter cluster, one at a time.
2. Re-register the vSphere Replication virtual machine from the datastore.
3. From the vSphere Replication appliance web front end, recreate the vSphere Replication plug-in. Verify the vSphere Replication plug-in is available in vCenter.
4. 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

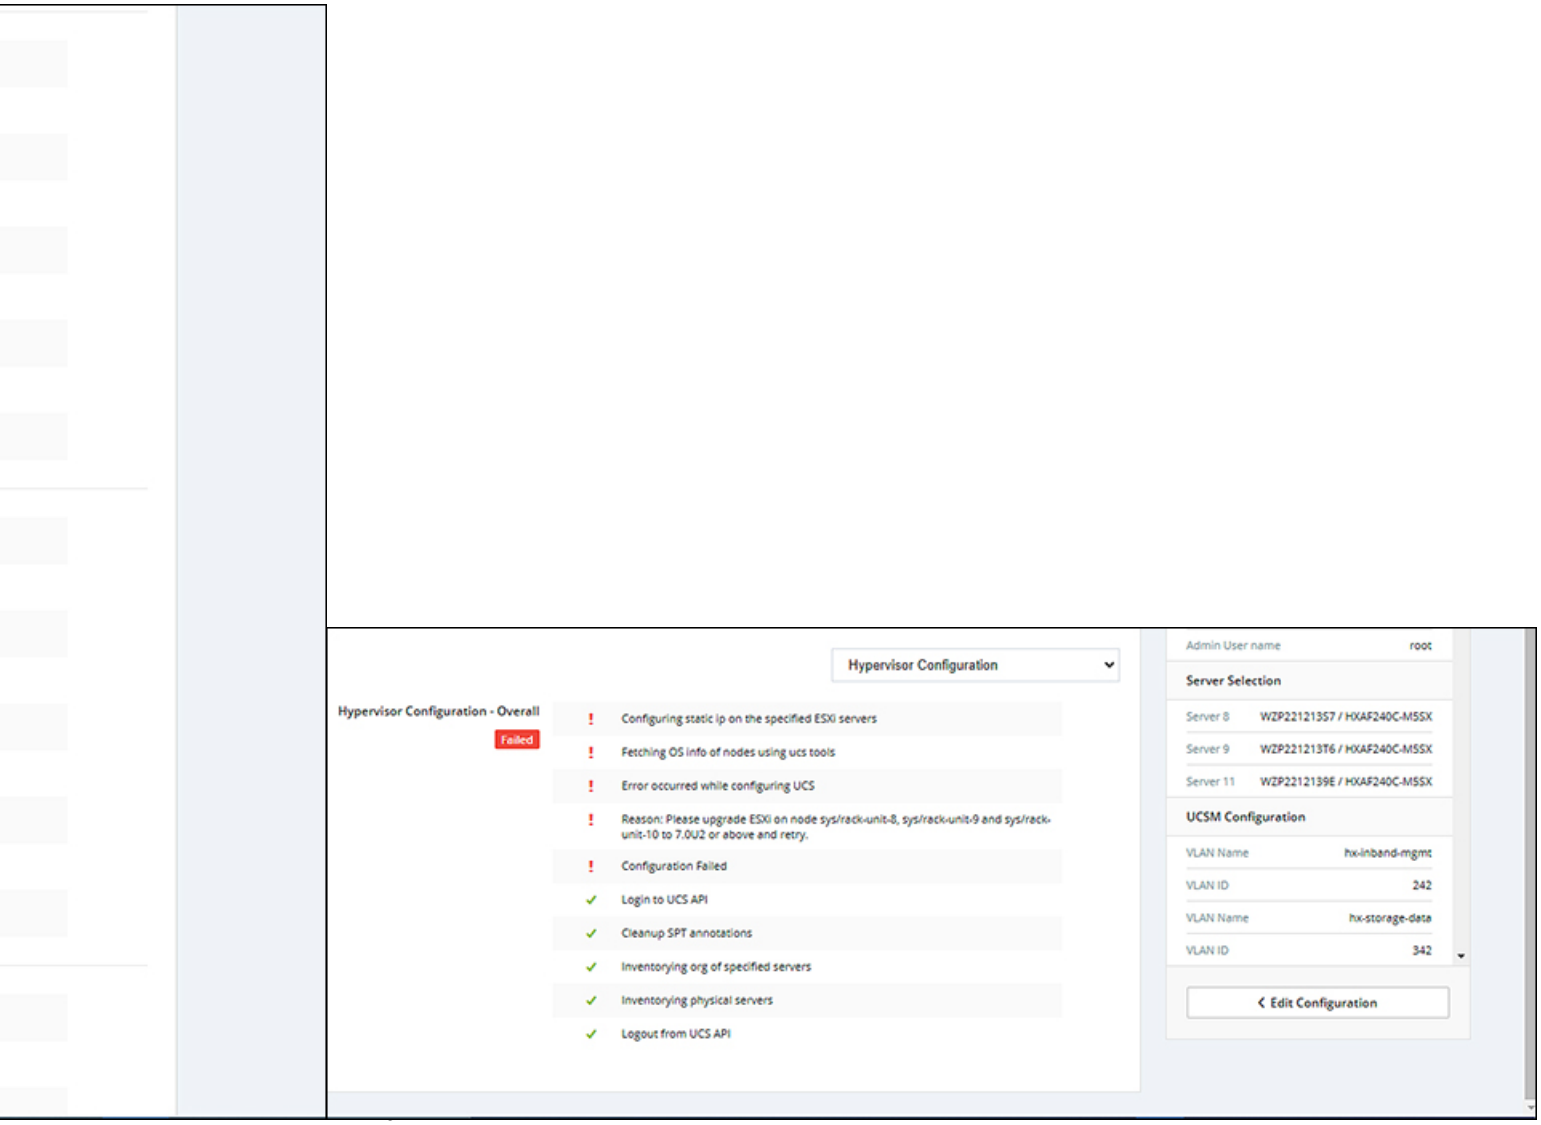

Action

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 Login to 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.

a) From the DirectPath I/O Configuration page, select **Configure Passthrough**.
b) From the Mark devices for passthrough list, select the LSI card for the pass through.
c) 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.

a) 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**.

b) 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 stCtlvm
```

```
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: Login to 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.sysmgmt.cluster_domain_id` and `com.springpath.sysmgmt.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: Login to 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 Controll (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 is 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.

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-ds1/ | 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**.**

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 Redlolog 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.

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.
- 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.

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 vShpere 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 datstore, 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 followig 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.

a) From vSphere Web Client, select ***host*** > **Manager** > **System** > **Time Configuration** > **User Network Time Protocol**.
b) From the **NTP Service Startup Policy** field, select **Start and stop with host**. Click **OK**.
c) 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.

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 Firebox 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.