



HyperFlex Upgrade Troubleshooting

- [HXDP Release 5.5\(1a\) and later Upgrade Error on M4 Servers, on page 1](#)
- [Upgrade to HXDP Release 6.0\(1a\) Fails or the hxupgrade_bundle persists, on page 2](#)
- [VMs Do Not Migrate During Upgrade, on page 2](#)
- [ESXi Host or HyperFlex Controllers In Lockdown Mode, on page 2](#)
- [Failed to Upgrade HyperFlex VIBs , on page 3](#)
- [HX Connect UCS Server Firmware Selection Dropdown Doesn't List the Firmware Version 4.1 or Above, on page 4](#)
- [Upgrade Fails in the Step - Entering Cluster Node into Maintenance Mode, on page 4](#)
- [Maintenance Mode Not Automatic for Cluster Containing VMs with vGPUs Configured, on page 5](#)

HXDP Release 5.5(1a) and later Upgrade Error on M4 Servers

Description

Starting with Cisco HyperFlex Release 5.5(1a), M4 servers are not supported. Attempts to upgrade clusters containing M4 or earlier HX generation servers to HXDP Release 5.5(1a) or later will fail in the pre-upgrade phase.

The upgrade and activity pages display an error in the **Bootstrap Upgrade** step. In some cases, the user is unable to view the error message and is shown a successful upgrade when in reality the upgrade has failed.

The fallback mechanism is to display a **ClusterUpgradeFailed** event along with a banner stating that the attempted upgrade is disallowed.

Symptom

An alert is generated and a banner appears stating that *One or more M4 platform nodes was detected in the cluster which is not supported starting HXDP 5.5(1a). Please follow the graceful node removal procedure to remove these nodes from the cluster or work with TAC to migrate from these nodes and retry the upgrade.*



Note The message center is populated with the same error message.

Action

Gracefully remove the unsupported nodes and retry the upgrade or contact TAC for further assistance.

Upgrade to HXDP Release 6.0(1a) Fails or the hxupgrade_bundle persists

Description

In the event the upgrade fails or the hxupgrade_bundle persists prior to doing kernel migration upgrade using the installer, manually remove the the hxupgrade_bundle from cvms and retry the kernel upgrade.

VMs Do Not Migrate During Upgrade

Description

Upgrading ESXi cluster fails with error "Node maintenance mode failed". This occurs in an online and healthy ESXi cluster where DRS and HA is enabled.

Action

Try the following workarounds in the following order:

1. If HA admission control policy is enabled and set to slot policy, change it to cluster resource percentage to tolerate one host failure and then retry upgrade.
2. Disable HA admission control policy or disable HA and then retry upgrade.
3. Try powering off a few VMs to make sure there is enough failover capacity in the cluster to tolerate at least one node failure and then retry upgrade.

ESXi Host or HyperFlex Controllers In Lockdown Mode

Description

If the ESXi host is in lockdown mode, pre-upgrade validation will fail with the error message *auth cancel*.

Action

Enable/Disable Lockdown Mode mode on the ESXi host and enable it after the upgrade is successful.

Using HyperFlex Controller VMs

1. Log into HX Connect.
2. On the Navigation pane, select **System Overview**.
3. On the **System Overview** tab, from the **Actions** drop-down list, you can enable or disable access to the controller VM using SSH as an administrator.

Using ESXi Hosts

1. Log into vSphere Web Client.
2. Browse to the host in the vSphere Web Client inventory.
3. Click the **Manage** tab and click Settings.
4. Under System, select **Security Profile**.
5. In the Lockdown Mode panel, click **Edit**.
6. Click **Lockdown Mode** and set the mode to Disabled.

Failed to Upgrade HyperFlex VIBs

Description

HXDP Upgrade to HX 4.5(1a) or above fails with error - "*Failed to upgrade HyperFlex VIBs . Reason: Some(System error)*".

The following error logs appear in the ESXi `esxupdate.log` file:

```
2020-12-01T11:59:22Z esxupdate: 333049: root: ERROR:
vmware.esximage.Errors.LiveInstallationError: ([], '([], "Error in running rm
/tardisks/scvmclie.v00:\nReturn code: 1\nOutput: rm: can\'t remove
\'/tardisks/scvmclie.v00\': Device or resource busy\n\nIt is not safe to continue. Please
reboot the host immediately to discard the unfinished update."
```

Action

Follow these steps to kill the process corresponding to `getstctlvmlogs` and retry the upgrade.

1. SSH to ESXi with root login.
2. Run the command `ps -c | grep -e cisco -e springpath` and note the process ID (PID). For example:

```
ps -c | grep -e cisco -e springpath
112056 112056 sh /bin/sh /opt/springpath/support/getstctlvmlogs
```
3. Kill the process using the command `kill -9 <PID from previous command>`. For example:

```
kill -9 112056
```
4. Go back to HX Connect or Intersight and retry the upgrade. If the issue still persists, please contact Cisco TAC for assistance.

HX Connect UCS Server Firmware Selection Dropdown Doesn't List the Firmware Version 4.1 or Above

Description

When you try to perform a combined upgrade from the HX Connect UI, the dropdown to select UCS server firmware doesn't show version 4.1 or later.

Action

Log into UCS Manager and confirm you have uploaded the UCS B and C firmware bundles to the Fabric Interconnect. If not, upload them and re-try the upgrade. If the UCS B and C firmware bundles are already uploaded to the Fabric Interconnect, apply below workaround to continue with upgrade.

1. From the HX Connect upgrade page, select **HX Data Platform** only.
2. Browse and select the appropriate HXDP upgrade package for your upgrade.¹
3. Enter your vCenter credentials.
4. Click **Upgrade**. This will bootstrap the management components. Refresh the UI screen.
5. Once the UI is refreshed, try the combined upgrade procedure. You should now be able to see the UCS server firmware version 4.1 or above listed in the dropdown menu.

Upgrade Fails in the Step - Entering Cluster Node into Maintenance Mode

Description

Failure at the **Entering Cluster Node into Maintenance Mode** step is caused because of a MTU mismatch in the vSwitch and port-groups. If the cluster has a node that was added at a later point using the node expansion method, the newly added node may have the MTU set to 9000 while other nodes are set to MTU 1500.



Note Below remediation is applicable only if your cluster has one or more nodes that were added as part of a cluster expansion, and they have the MTU set to 9000 while your original cluster nodes are set at a MTU of 1500. If this is not the scenario, please contact TAC for further assistance.

Action

- Log into vCenter.
- Check and confirm the MTU value set on all nodes.

¹ The version must be HXDP 4.5 or later.

- If the nodes that were part of the originally built cluster are set at a MTU of 1500 and some of the other nodes (nodes added later as part of cluster expansion) have the MTU set to 9000, change the MTU on all such nodes to 1500.
- Retry the upgrade.

Maintenance Mode Not Automatic for Cluster Containing VMs with vGPUs Configured

Description

For clusters that contain VMs with vGPUs configured, entering maintenance mode does not occur automatically even with DRS fully enabled. During rolling upgrades, it is necessary to manually handle these VMs to ensure that each ESXi host can enter maintenance mode and continue with the upgrade at the appropriate time.

Action

You can use one of these methods to proceed forward:

1. Manually vMotion the vGPU configured VMs to another ESXi host in the cluster.
2. Temporarily power off the vGPU configured VMs. They can be powered on again after the ESXi host reboots and rejoins the cluster



Note This is a limitation of DRS host evacuation and is documented, see "DRS fails to migrate vGPU enabled VM's automatically (66813) topic" on the VMware documentation site.
