Stretched Cluster Upgrade

- Overview, on page 1
- Upgrade Guidelines for Stretched Cluster, on page 2
- Upgrading HyperFlex Stretched Cluster Using HX Connect, on page 2
- Upgrading a Witness VM for HXDP Release 5.0(x) and Earlier, on page 4
- Manually Upgrading ESXi for Cisco HyperFlex Stretched Cluster, on page 6
- Configuring Stretched Cluster for UCS FW Upgrade, on page 7

Overview

This section provides information related to upgrading a Cisco HyperFlex Stretched Cluster. The procedure for performing a Stretched Cluster upgrade is similar to the regular HyperFlex cluster upgrade procedure.

Cisco HyperFlex Stretched cluster upgrade process involves upgrading the below 3 components:

- Cisco HyperFlex Data Platform[1]
- VMware vSphere ESXi
- Cisco UCS Server Firmware

You can combine HyperFlex Data Platform and VMware ESXi upgrade in a single combined upgrade for HyperFlex Stretched clusters. Cisco recommends upgrading these 2 components in a combined upgrade from HyperFlex Connect. You can choose to upgrade one or two of these components at a time.

If you prefer to upgrade individual components one at a time, see Upgrade UCS Firmware, HX Data Platform and VMware vSphere - Individual Component Upgrade. The component upgrade process for standard cluster and HyperFlex Stretched cluster are the same.

This section describes the steps to perform a combined upgrade of HyperFlex Data Platform and VMware vSphere ESXi. In this process, the HyperFlex nodes will go through an optimized rolling reboot without any workload disruptions by use of VMware vMotion.

[1] Use the installer to upgrade to HXDP 6.0(1a). Upgrade to HXDP 6.0(1a) by HX Connect is not supported.

Upgrade Guidelines for Stretched Cluster

- Upgrade of UCS server firmware is not supported through HX Connect. UCS firmware upgrade must be done manually using Cisco UCS Manager. See Manage Firmware through Cisco UCS Manager.
- Upgrade of the HyperFlex Witness node is not required when upgrading stretched clusters but is strongly recommended. Refer to the Cisco HyperFlex Software Requirements and Recommendations for the latest witness version available.
- Health Check — Cisco recommends running this proactive health check on your HyperFlex cluster prior to upgrade. For more details, see Pre-Upgrade Intersight Health Check

Upgrading HyperFlex Stretched Cluster Using HX Connect

Before you begin

- Complete pre-upgrade validation checks.
- Download the latest *Cisco HX Data Platform Upgrade Bundle for upgrading existing clusters from previous releases*, from Software Download .
- Upgrade Cisco UCS Infrastructure.
- Upgrade UCS server firmware. For more information, see Manage Firmware through Cisco UCS Manager.
- Disable snapshot schedule, on the storage controller VM. SSH to HyperFlex cluster IP and run the command **stcli snapshot-schedule –disable** snapshot schedule.
- If DRS is *Enabled* and set to fully automatic mode, the VMs are automatically migrated to other hosts with vMotion.

Note If DRS is *Disabled*, vMotion the VMs manually to continue the upgrade process when prompted. For more information, see VMware Documentation for Migration with vMotion.

Step 1 Log in to HX Connect.
a) Enter the administrative username and password.
b) Click **Login**.

Step 2 In the Navigation pane, select **Upgrade**.

Step 3 On the **Select Upgrade Type** page, select **HX Data Platform** and **ESXi** and complete the following fields:

Step 4 On the **Select Upgrade Type** page, select **HX Data Platform** and complete the following fields:

UI Element	Essential Information
Drag the HX file here or click to browse	Upload the latest Cisco HyperFlex Data Platform Upgrade Bundle for upgrading existing clusters with previous release.tgz package file from Download Software - HyperFlex HX Data Platform. Sample file name format: storfs-packages-4.5.1a-31601.tgz.
Current version	Displays the current HyperFlex Data Platform version.
Current cluster details	Lists the HyperFlex cluster details like the **HyperFlex version** and **Cluster upgrade state**.
Bundle version	Displays the HyperFlex Data Platform version of the uploaded bundle.
(Optional) **Checksum** field	The MD5 Checksum number is available by hovering over the filename in the Cisco.com Software Download section. This is an optional step that helps you verify the integrity of the uploaded upgrade package bundle.

Step 5 Upload the VMware ESXi custom image offline upgrade bundle.

Step 6 Provide vCenter login credentials:

Essential Information	Essential Information
User Name field	Enter the vCenter <admin> username.
Admin Password field	Enter the vCenter <admin> password.

Step 7 Click **Upgrade** to begin the combined upgrade process.

Step 8 The Validation screen on the **Upgrade Progress** page displays the progress of the checks performed. Fix validation errors, if any.

Note At this point, all pre-upgrade checks and validations are running, along with the initial upgrade stage. Within a few minutes, HX Connect returns and prompts you to confirm and start the second stage of the upgrade. The upgrade is not complete until both steps are performed in the UI. The system should never be left in a state where only the first step of the upgrade is complete.

Step 9 The HyperFlex Connect UI refreshes after the first step of the upgrade, and a banner pops up prompting you to provide the UCS and vCenter credentials and start the second stage of the upgrade process. Monitor the upgrade page and confirm that the upgrade is complete.

When upgrade is in progress, you may see an error message **Websocket connection failed. Automatic refresh disabled**. You can either refresh the page or log out and log back in to clear the error message. You can safely ignore this error message.

Note Perform post upgrade tasks once the upgrade is complete. If the upgrade fails, you can re-try the upgrade or contact Cisco TAC for further assistance.

Upgrading a Witness VM for HXDP Release 5.0(x) and Earlier

Before you begin

Attention This workflow is for use with existing Stretched Cluster installs (HXDP Release 5.0(x) and earlier). New Stretched Clusters installed using HXDP 5.5(1a) and later will auto-configure an Invisible Cloud Witness for site arbitration. Invisible Cloud Witness automatically runs the latest version, user maintenance of this component is not required.

- Select the Witness VM version that supports the HXDP version you are upgrading to.

  For supported versions see the *HX Data Platform Software Versions for HyperFlex Witness Node for Stretched Cluster* section of the Cisco HyperFlex Software Requirements and Recommendations

  for your respective upgrade.
- Upgrade HyperFlex Stretched Cluster.
- The upgraded HyperFlex Stretched Cluster must be in healthy state. To check the health state of Stretched Cluster after upgrade, run the following command:

  ```
  root@StCtlVM:~# hxcli cluster info | grep healthy
  ```

Step 1 Log into the witness VM using SSH and execute the following command to stop the service exhibitor.

```
root@WitnessVM:~# service exhibitor stop
```

Step 2 Copy the `exhibitor.properties` file available in the `/usr/share/exhibitor/` path to a remote machine from where you can retrieve the `exhibitor.properties` file.

```
scp root@<Witness-VM-IP>:
/usr/share/exhibitor/exhibitor.properties user@<Remote-Machine>:
/directory/exhibitor.properties
```

Step 3 Log out from the Witness VM. Power off and rename the Witness VM to WitnessVM.old.

Note Confirm that the IP address of the old Witness VM is unreachable, using the ping command.

Step 4 Deploy a new Witness VM and configure the same IP address as the old Witness VM.

Note If the IP address is not reachable, the Witness OVA deployment may contain stale entries in the `/var/run/network` directory. You must manually remove these entries and reboot the VM to have the assigned IP address become reachable on the network.

To reboot the VM, open the VM console in vCenter/vSphere and execute the following command:

```
rm -rf /var/run/network/*
reboot
```

Step 5 Log into the new witness VM using SSH and execute the following command to stop the service exhibitor.

```
root@WitnessVM:~# service exhibitor stop
```

Step 6 Copy the `exhibitor.properties` file from the remote machine (copied in Step 2) to the `/usr/share/exhibitor/` path of the new Witness VM.

```
scp /directory/exhibitor.properties root@<Witness-VM-IP>:
/usr/share/exhibitor/exhibitor.properties
```

Step 7 Verify if the following symlinks are preserved in the new Witness VM:

```
root@Cisco-HX-Witness-Appliance:~# cd /etc/exhibitor/
root@Cisco-HX-Witness-Appliance:/etc/exhibitor# ls -al
total 8
drwxr-xr-x 2 root root 4096 Sep 11 13:00 .
drwxr-xr-x 88 root root 4096 Sep 11 12:55 ..
lrwxrwxrwx 1 root root 41 Sep 11 13:00 exhibitor.properties
lrwxrwxrwx 1 root root 37 Jul 24 16:49 log4j.properties
```

If the symlinks are not available, execute the following command:

```
root@Cisco-HX-Witness-Appliance:/etc/exhibitor# ln
 -s /usr/share/exhibitor/exhibitor.properties exhibitor.properties
root@Cisco-HX-Witness-Appliance:/etc/exhibitor# ln -s /usr/share/
exhibitor/log4j.properties log4j.properties
root@Cisco-HX-Witness-Appliance:/etc/exhibitor# ls -al total 8
drwxr-xr-x 2 root root 4096 Sep 11 13:00 .
drwxr-xr-x 88 root root 4096 Sep 11 12:55 ..
lrwxrwxrwx 1 root root 41 Sep 11 13:00 exhibitor.properties
-> /usr/share/exhibitor/exhibitor.properties
lrwxrwxrwx 1 root root 37 Jul 24 16:49 log4j.properties
 -> /usr/share/exhibitor/log4j.pr
```

Step 8 **Important** Skip this step if your upgrade is an HX Release prior to version 4.5(2a).

Get the cluster UUID value from `/etc/springpath/clusteruuid` file on any Controller VM and then edit `/usr/share/zookeeper/conf/server.jaas.conf` and replace `HX_PLACEHOLDER` with Cluster UUID.

Example:

```
On any Storage Controller VM, get the contents of clusteruuid
 file -cat /etc/springpath/clusteruuid 5b4f718033594dad:663273adab09777a
On the Witness VM, replace the HX_PLACEHOLDER value with the cluster
UUID -Before -QuorumServer {
       org.apache.zookeeper.server.auth.DigestLoginModule required
user_hyperflex="HX_PLACEHOLDER";};QuorumLearner {
       org.apache.zookeeper.server.auth.DigestLoginModule required
       username="hyperflex"
       password="HX_PLACEHOLDER";};
After -QuorumServer {org.apache.zookeeper.server.auth.DigestLoginModule
 required user_hyperflex="5b4f718033594dad:663273adab09777a";};
 QuorumLearner { org.apache.zookeeper.server.auth.DigestLoginModule
 required
       username="hyperflex"
       password="5b4f718033594dad:663273adab09777a";};
```

Step 9 **Note** This step is required for users that are moving to Witness VM Node version 1.1.1 and later, if the Witness VM being upgraded is a version previous to 1.1.1.

Run the `/usr/share/springpath/storfs-misc/setexhibitorconfig.sh` command to upgrade the Witness exhibitor configuration.

Note The `setexhibitorconfig.sh` automates the process of editing the `exhibitor.properties` file, and replaces all of the data IP addresses with the management IP addresses for each corresponding controller VM.

Note It is normal for this command to not show any output when upgrading from a Witness VM that is older than 1.1.1.

Step 10 Start the service exhibitor by executing the following command:

```
root@Cisco-HX-Witness-Appliance:~# service exhibitor start
exhibitor start/running, process <ID>
```

Manually Upgrading ESXi for Cisco HyperFlex Stretched Cluster

Step 1 Select one of the hosts and put it in HX maintenance mode using the vSphere Web Client. After the host enters maintenance mode, complete the following steps.

Step 2 Copy files using SCP, start the SSH service in the destination ESXi hosts as well.

Note
- On HX240, you can use the local SpringpathDS datastore or a mounted HX datastore.
- On HX220, you can use either a mounted HX datastore or create a temporary RAM disk.

```
scp local_filename user@server:/path/where/file/should/go
```

Step 3 Log into ESXi, and execute the following command to query the list of available image profiles and for profile name verification.

```
esxcli software sources profile list -d <location_of_the_esxi_zip_bundle_on_the_datastore>
```

Attention Full path must be used when running the `esxcli` software command.

Example:

```
[root@localhost:~] esxcli software sources profile list -d /vmfs/volumes/5d3a21da-7f370812-ca58-0025
b5a5a102/HX-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10-upgrade-bundle.zip
Name                                        Vendor  Acceptance Level  Creation Time       Modification
 Time
------------------------------------------  ------  ----------------  ------------------
------------------
HX-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10  Cisco   PartnerSupported  2019-04-02T00:14:56
2019-04-02T13:38:34
```

Step 4 Run the following command to perform the upgrade.

```
esxcli software profile update -d <path_to_profile_ZIP_file> -p < profile name>
```

Example:

```
[root@HX-ESXi-01:/vmfs/volumes/1a234567-89bc1234] esxcli software profile update -d
/vmfs/volumes/1a234567-89bc1234/HX-Vmware-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10.zip
-p HX-ESXi-7.0U3-20328353-Cisco-Custom-7.3.0.10
```

Step 5 After the ESXi host comes up, verify that the host has booted up with the correct version.

```
vmware -vl
```

Step 6 Exit maintenance mode using the vSphere Web Client.

Step 7 Ensure that the cluster becomes healthy between each ESXi upgrade.

```
hxcli cluster info [flags]
```

Step 8 Repeat this process for all hosts in the cluster in a sequence.

Note Make sure that the cluster becomes healthy between each ESXi upgrade.

Configuring Stretched Cluster for UCS FW Upgrade

During upgrade, the following customized UCS policies are validated and adjusted for HyperFlex:

- HFP (Host Firmware Package) - Host Firmware Packages provide consistent firmware files for the multiple components of a HyperFlex node. This includes CIMC, BIOS, HBA and SAS Expander firmware, VIC and other components. Unlike typical UCS Host Firmware Packages, they also control disk firmware, due to the criticality of this to HyperFlex Data Platform. Note that Self Encrypting Drives (SED) firmware is controlled by HyperFlex Data Platform directly and not UCS Manager policies.
- VNIC Templates - Virtual NIC (VNIC) templates provide consistent configuration of VNIC's between UCS fabrics. HyperFlex VNIC Templates are configured as redundancy pairs to ensure changes to HyperFlex VNIC's on one UCS fabric is applied to the other.
- Ethernet Adaptor Policies - Ethernet Adaptor Policies provide performance related properties for HyperFlex VNIC's.
- BIOS Policies - BIOS policies control configuration and performance of key hardware resources on a HyperFlex node, such as CPU and Memory. HyperFlex uses specific configuration to provide consistent high performance.
- VNIC/VHBA Placement Policies - VNIC/VHBA placement policies determine the PCI addresses presented to the HyperFlex node for a given VNIC/VHBA. HyperFlex sets this in a consistent manner so further configuration can proceed successfully.

Step 1 SSH to any CVM on a site and change directory into /tmp

Step 2 Use the `su diag` Enter into the diag shell on a storage controller VM.[2]

Step 3 Run the following command: `/usr/local/bin/hx.py --upgrade-cluster-config`. This generates a file called "`customer_site_config.json`" and saves it in the /tmp directory.

Step 4 Edit the `customer_site_config.json` file to change the firmware version and the org name appropriately. For example:

Example:

```
{
    "id": "Advanced",
    "collapse": true,
    "label": "Advanced",
    "groups": [
      {
        "id": "firmware",
        "label": "UCS Firmware",
        "items": [
```

[2] HXDP Release 5.0(2a) and later support the `su diag` command. For more information, see the Diag User Overview.

```
          {
            "id": "version",
            "label": "UCS Firmware Version",
            "type": "text",
            "description": "UCS Firmware Version to be used on the HX servers",
            "placeholder": "ex: 3.2(2d)",
            "defaultValue": "3.2(2d)",
            "value": "4.1(1d)"  #<<<<<----------------- Change this
          },
          {
            "id": "version-m5",
            "label": "UCS Firmware Version",
            "type": "text",
            "description": "UCS Firmware Version to be used on the M5 HX servers",
            "placeholder": "ex: 3.2(2d)",
            "defaultValue": "3.2(2d)",
            "value": "4.1(1i)" #<<<<<----------------- Change this
          }
        ]
      },
      {
        "id": "org",
       "items": [
                  {
                        "id": "name",
                        "label": "Hyperflex Org name",
                        "type": "text",
                        "value": "Faridabad", #<<<<<----------------- Change this
                       "description": "The name of the org in ucsm which is to be used for creation
 of all the policies and profiles for this Hyperflex cluster"
                  }
        ]
      }
    ]
```

Step 5 Execute the command again and enter the UCSM IP and credentials.

For example:

```
/usr/local/bin/hx.py --upgrade-cluster-config
```

Example:

```
[root@SpringpathControllerVP0RX5DWTC:/# /usr/local/bin/hx.py --upgrade-cluster-config
                [UCS Manager] [in_progress][  0.00%][ETA:  0:18:00] Login to UCS API
                 UCS host name or virtual IP address: 10.42.17.11
                 Connecting to admin@10.42.17.11...
                Password:
```

Step 6 Ensure that the command runs without any error. If there is an error, contact Cisco TAC.

Note Note that this command (hx.py) is being run for the first site FI domain. You need to run the same steps for the second site FI domain later.

Step 7 Perform the following steps in vCenter and UCSM:

a) Verify that Pending reboot appears in the pending activities of the UCSM.
b) Put one host in maintenance mode.
c) Reboot the server and then wait for the server to come online and cluster to be online/healthy.
d) Perform the same steps for the remaining nodes.

Step 8 Repeat Steps 4, 5 and 6 again for the other site.