



The bridge to possible

White paper
Cisco public

Cisco HyperFlex Fabric Interconnect Hardware Live Migration Guide

Version 1.0

October 2020

Contents

Document information	3
Intended use and audience	3
Legal notices	3
Introduction	4
Prerequisites	5
Overview	5
Prechecks	6
Before migration	7
Migration procedure	9
Optional: Upgrade HyperFlex server uplinks to 25Gbps with Cisco VIC 1457	15
Optional: Adding more links for higher network bandwidth	18
For more information	22

Document information

Document summary	Prepared for	Prepared by
Fabric Interconnect (FI) hardware online upgrade with HyperFlex™ cluster	Cisco field	Ganesh Kumar, technical marketing engineer

Intended use and audience

The audience for this document consists of system architects, system engineers, and any other technical staff who are responsible for planning or upgrading Fabric Interconnects with a Cisco HyperFlex cluster. Although every effort has been made to make this document appeal to the widest possible audience, the document assumes that readers have an understanding of Cisco UCS® hardware, terminology, and configurations.

Legal notices

All information in this document is provided in confidence and shall not be published or disclosed, wholly or in part, to any other party without Cisco's written permission.

Introduction

The Cisco 6400 Series Fabric Interconnects (FIs) are new Cisco UCS hardware platforms that use the latest Cisco® silicon Switch-on-Chip (SoC) ASIC technology. Because the hardware is new, you may need to make changes in port and software configurations to experience a successful migration. This document discusses the process for upgrading from Cisco UCS 6200 Series Fabric Interconnects to Cisco UCS 6454 Fabric Interconnects when Cisco HyperFlex clusters are connected to the fabric. The migration procedure remains the same for the 64108 platform as well. After reading this document, you should understand the upgrade process and any precautions you should consider for a successful migration.

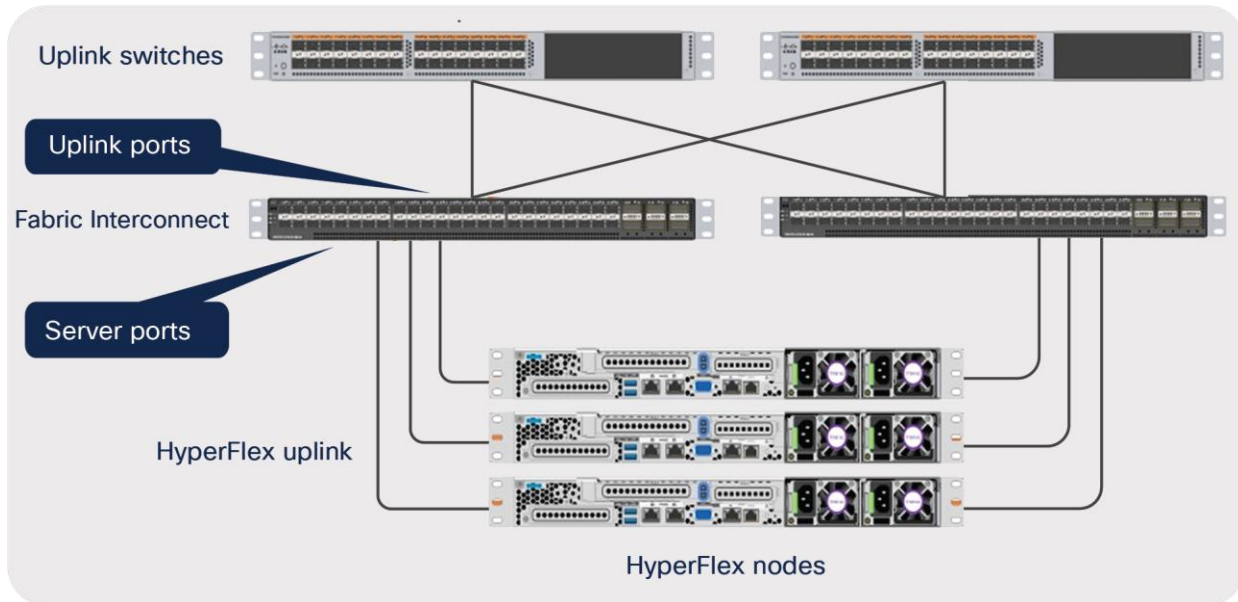


Figure 1.
Cisco HyperFlex network connectivity

Cisco HyperFlex clusters in the data center are deployed with a pair of Cisco UCS Fabric Interconnects (see Figure 1) for server uplink connectivity. The Fabric Interconnects are always Active/Active for data traffic. The Cisco UCS Fabric Interconnects are connected to a pair of network uplink switches for northbound traffic. As you can see in Figure 1 above, the HyperFlex network connectivity is recommended always to be redundant, and hence the Fabric Interconnect migration procedure is a complete online operation without any impact on data traffic.

Prerequisites

We recommend reviewing the Cisco HyperFlex HX Data Platform [release notes, installation guide, and user guide](#) before proceeding with any configuration. The HyperFlex cluster should be healthy and ONLINE. The UCS FI cluster state should be “HA READY” before you start the FI migration. Please contact Cisco support or your Cisco representative if you need assistance. Licenses from Cisco UCS 6200 Series Fabric Interconnects are not transferable to Cisco UCS 6400 Series Fabric Interconnects. You must obtain licenses for the Cisco UCS 6400 Series Fabric Interconnects before you upgrade if there are insufficient port licenses to cover the migration.

Cisco UCS Manager Release 4.1 is the minimum version that provides support for Cisco UCS 64108 Fabric Interconnects, and Release 4.0 is the bare minimum version that provides support for Cisco UCS 6454 Fabric Interconnects. To migrate from Cisco UCS 6200 Series Fabric Interconnects to Cisco UCS 6400 Series Fabric Interconnects:

- Cisco UCS 6200 Series Fabric Interconnects must be on Cisco UCS Manager Release 4.1(1) or a later release to migrate to Cisco UCS 64108 Fabric Interconnects, and Release 4.0(1) or a later release to migrate to Cisco UCS 6454 Fabric Interconnects.
- Cisco UCS 6400 Series Fabric Interconnects must be loaded with the same build version that is on the Cisco UCS 6200 Series Fabric Interconnect that it will replace.
 - Power on the Fabric Interconnects independently and connect to console.
 - Configure the new Fabric Interconnect in “[standalone](#)” mode, and check the version by running the command **show version**.
- The UCS 6400 Series Fabric Interconnect is intended as a replacement for the UCS 6200 Series Fabric Interconnect, but not as a replacement for the higher speed (or 40Gb) UCS 6332/6332-16UP Fabric Interconnects. Therefore, Cisco has not tested or published a plan to migrate from UCS 6332/6332-16UP Fabric Interconnects to UCS 6400 Series Fabric Interconnects. The high-density UCS 64108 Fabric Interconnect is an ideal upgrade from the high-density Cisco UCS 6296 Fabric Interconnect, while the UCS 6454 Fabric Interconnect is an ideal upgrade from the Cisco UCS 6248 Fabric Interconnect.
- HyperFlex stretched-cluster architecture requires the same model of Fabric Interconnect to be deployed in both sites. So, if you have a stretched cluster, make sure to upgrade/migrate all four Fabric Interconnects to the UCS 6400 platform. Specific information about HyperFlex stretched-cluster Fabric Interconnect requirements and limitations can be found in the [HyperFlex Stretched Cluster White Paper](#).

Overview

The fourth-generation Cisco UCS 6400 Fabric Interconnects provide both network connectivity and management capabilities for the data center HyperFlex system. The Cisco UCS 6400 Series offers line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. From a networking perspective, the Cisco UCS 6400 Series use a cut-through architecture, supporting deterministic, low-latency, line-rate 10/25/40/100 Gigabit Ethernet ports, and a switching capacity of 3.82 Tbps for the 6454 and 7.42 Tbps for the 64108 FI model.

For the purposes of this document, we use the Cisco UCS Fabric Interconnect 6454 model. However, the migration procedure remains the same for the 64108 model. It has 28 10/25-Gbps Ethernet ports, 4 1/10/25-Gbps Ethernet ports, 6 40/100-Gbps Ethernet uplink ports, and 16 unified ports that can support 10/25-Gbps Ethernet ports or 8/16/32-Gbps Fibre Channel ports (see Figure 2). All Ethernet ports are capable of supporting FCoE.




Figure 2.
Cisco UCS 6454 Fabric Interconnect ports (rear view)

Two hot swappable power supplies and four variable-speed fans insert into the front of the chassis with front-to-back cooling. Server-facing ports are at the rear of the chassis to simplify the connections to the servers. Serial, management, and USB ports are provided at the front of the chassis for easy access at the front of the rack (Figure 3).



Figure 3.
Cisco UCS 6454 Fabric Interconnect (front view)

Prechecks

The Cisco UCS 6200 Series Fabric Interconnects are the only Fabric Interconnects that can be migrated to the Cisco UCS 6454/64108 Fabric Interconnects while in service (Table 1). No other Fabric Interconnect series is supported.

Note: The Cisco UCS 6400 Series Fabric Interconnect is intended as a replacement for the UCS 6200 Series Fabric Interconnect, but not as a replacement for the higher speed (or 40Gb) UCS 6332/6332-16UP Fabric Interconnect. Therefore, Cisco has not tested or published a plan to migrate from UCS 6332/6332-16UP Fabric Interconnects to UCS 6400 Series Fabric Interconnects. The high-density UCS 64108 Fabric Interconnect is an ideal upgrade from the high-density Cisco UCS 6296 Fabric Interconnect.

Table 1. Migration support matrix for infrastructure hardware

From	To	Support
Cisco UCS 6100 Series Fabric Interconnects	Cisco UCS 6454 or UCS 64108 Fabric Interconnects	Not Supported
Cisco UCS 6200 Series Fabric Interconnects		Supported
Cisco UCS 6332 and 6332-16UP Fabric Interconnect		Not Supported
Cisco UCS 6324 Fabric Interconnects		Not Supported

Note: Even though the migration can be done live without impacting data traffic, it is recommended to perform the migration in a planned maintenance window. If you have the 6200 series Fabric Interconnects configured with Fiber Channel (FC) ports, you need to unconfigure those ports for the migration.

Before migration

For a successful Fabric Interconnect migration, it is important that the current pair of Fabric Interconnects and the new pair are running the same Cisco UCS Manager (UCSM) release. UCS FI 64108 requires that a minimum of UCSM Release 4.1, and FI 6454 requires that a minimum of UCSM Release 4.0, is running.

Before you start the migration, make sure to perform the activities described below:

- Power on the new Fabric Interconnects and check the UCSM version running in them.
 - Connect the console.
 - Configure the new Fabric Interconnect in “standalone” mode, and check the UCSM version by running the command **show version**.
- Based on your new Fabric Interconnect model and the UCSM version that is running, download the required Cisco UCS Manager Infrastructure (Bundle A) package from the Cisco software downloads page, and upgrade the current 6200 Fabric Interconnects.
- If the new Fabric Interconnects were part of any other previous UCSM cluster, they may be running a different configuration. Therefore, clean up the new fabric interconnect by running the commands given below on each fabric interconnect:
 - `# connect local-mgmt`
 - `# erase configuration`
- Back up and export the Cisco UCS Manager configuration that can be used to restore the configuration in future if needed.
 - Connect to UCS Manager and navigate to **Admin menu -> Backup Configuration** and take a “Full State” backup (Figure 4).

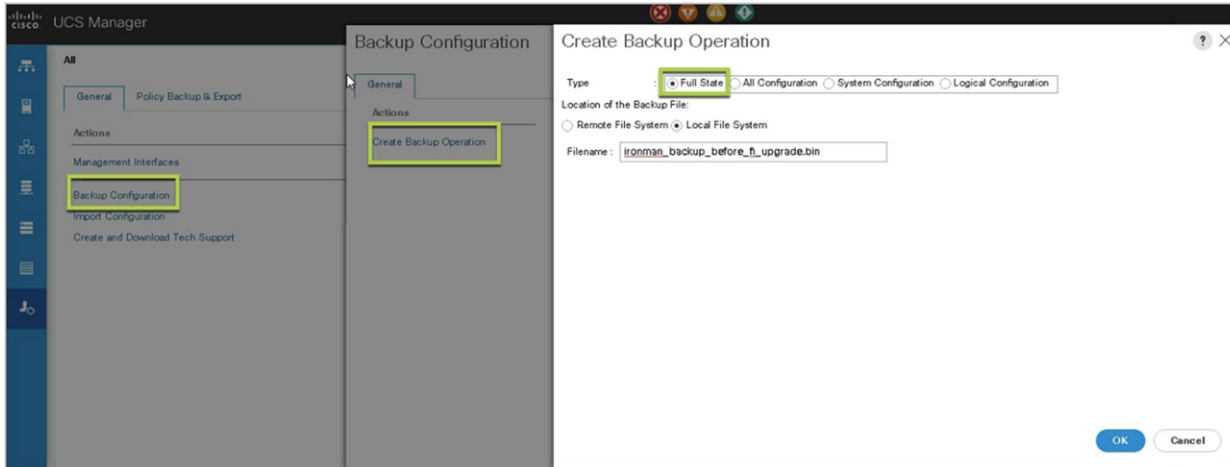


Figure 4.
Cisco UCS Manager configuration backup

- Note down the current Fabric Interconnect connection details, such as which server connects to which port in each FI and the uplink port connections.
- Note down the current Fabric Interconnect IP details (FI-A IP, FI-B IP, UCSM Cluster IP, Netmask, and the Gateway).
- Check and confirm that all connected servers have redundant data paths visible and online with appropriate failover policies configured. Servers not configured with network redundancy will experience a network outage during the migration.
- Check and confirm that the current UCSM cluster is in a healthy state by running the command **show cluster state**, and confirm that the HA status is "Ready."

```
ironman-A# show cluster state
Cluster Id: 0x7880a79ed73911ea-0x9eb3547feef634a1

A: UP, PRIMARY
B: UP, SUBORDINATE

HA READY

ironman-A#
```

Figure 5.
UCSM cluster status

- Login to HyperFlex Connect or SSH to the HyperFlex cluster IP and verify that the cluster is healthy and all the nodes are online.
 - Run the command **hxcli cluster state**, and confirm that the cluster state is “ONLINE” and the Resiliency Health is “HEALTHY.”

```
admin@SpringpathController26HVPFER95:~$ hxcli cluster health
Cluster State           : ONLINE
Resiliency Health       : HEALTHY
Policy Compliance       : COMPLIANT
Data Replication Factor : 3 Copies
# of node failures tolerable : 1
# of persistent device failures tolerable : 2
# of cache device failures tolerable : 2
Messages                : Storage cluster is healthy.
Data Replication Compliance : COMPLIANT
```

Figure 6.
HyperFlex cluster status

- Run the command **hxcli node list**, and verify that the storage controller and hypervisor status are “ONLINE” for all nodes in the cluster.

```
admin@SpringpathController84DL8LF46P:~$ hxcli node list
+-----+
| NODE NAME | HYPERVISOR ADDRESS | HYPERVISOR STATUS | CONTROLLER ADDRESS | CONTROLLER STATUS | MODEL | VERSION | # DISKS |
+-----+
| hxaf240-m5-node-2 | | ONLINE | | READY | HXAF240C-M5SX | 4.0(2b) | 7 |
| hxaf240-m5-node-3 | | ONLINE | | READY | HXAF240C-M5SX | 4.0(2b) | 7 |
| hxaf240-m5-node-1 | | ONLINE | | READY | HXAF240C-M5SX | 4.0(2b) | 7 |
+-----+
```

Figure 7.
HyperFlex node and controller status

Migration procedure

Now we are ready to start the migration. Perform the steps below in the order given for a successful online migration.

Note: This Fabric Interconnect migration is an online process with very quick failover of physical NICs. The migration can be done without powering down any workloads running on the HyperFlex cluster. Even though the migration can be done live without impacting the data traffic, it is recommended to perform the migration in a planned maintenance window.

1. Log in to UCS Manager and identify the subordinate Fabric Interconnect, then disable the server ports connected to each HyperFlex node.

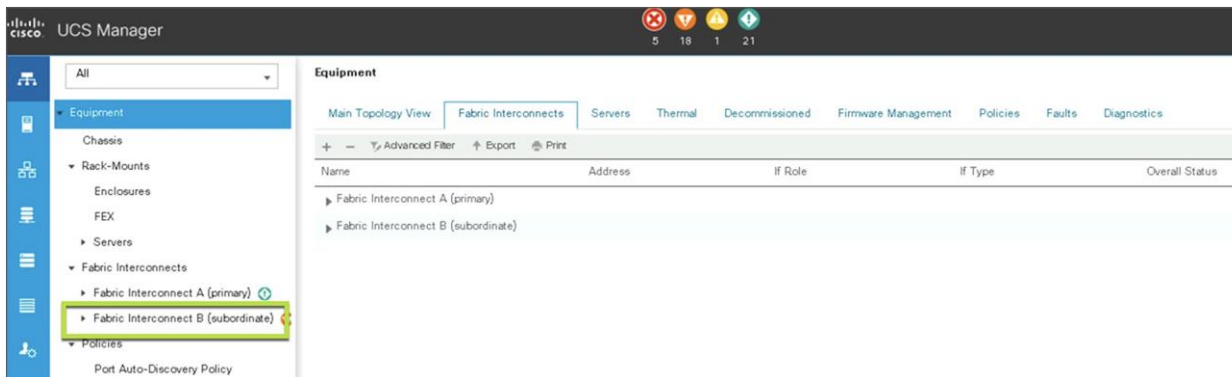


Figure 8.
Identifying subordinate Fabric Interconnect

Note: If you have a Cisco UCS B-Series chassis connected to the same pair of Fabric Interconnects, follow the Cisco UCS–recommended process to evacuate the traffic from the Fabric Interconnect before you disable the server ports. Refer to [UCS Fabric Interconnect Traffic Evacuation](#) for more details.

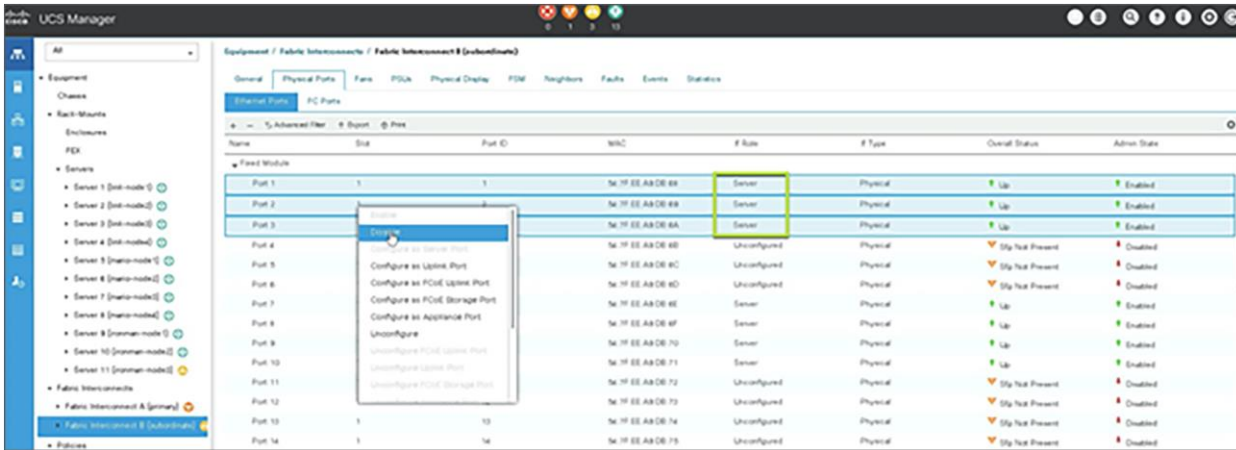


Figure 9. Disable subordinate Fabric Interconnect server ports

2. Confirm that the data path is failed over and that data traffic is not being disrupted from the hypervisor side.

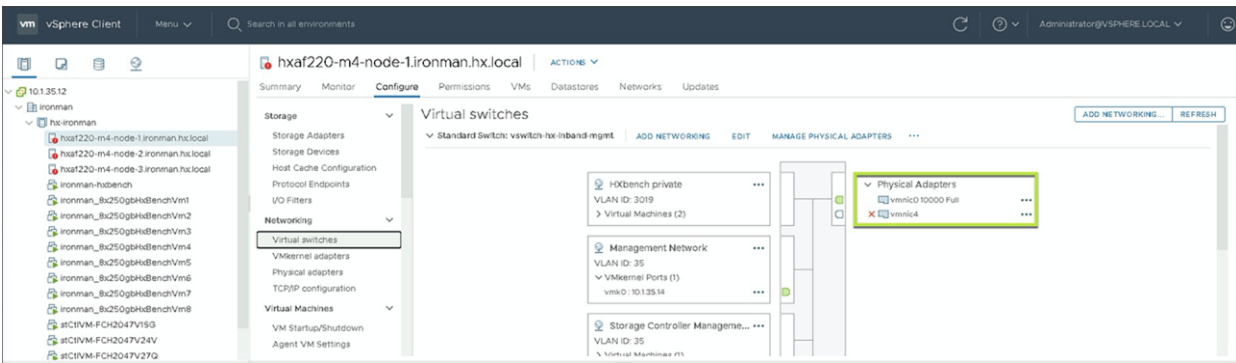


Figure 10. Hypervisor traffic failover

3. Disable all the uplink interfaces. Go to LAN -> Fabric A/B (subordinate fabric) -> click Disable All Uplink Interfaces.

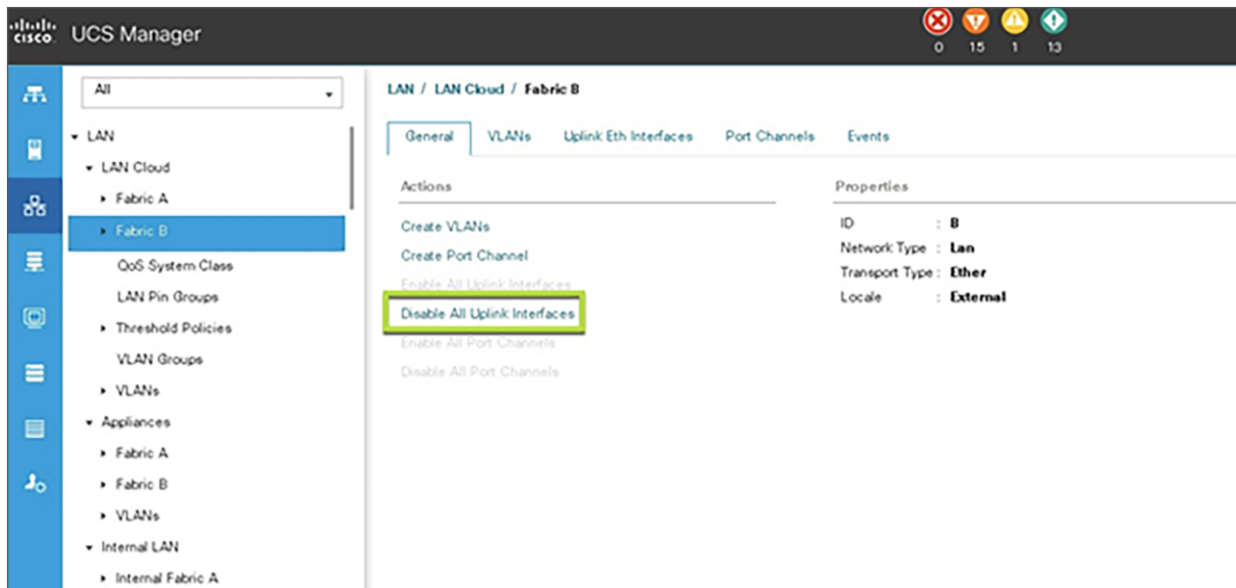


Figure 11.
Disable subordinate Fabric Interconnect uplink ports

4. Power down the subordinate Fabric Interconnect by unplugging it from the power source.
5. Mount the replacement Cisco UCS 6400 Series Fabric Interconnect into either the same rack or an adjacent rack. Refer to the [Cisco UCS 6400 Series Installation Guide](#) for details.
6. Disconnect the cables from the HyperFlex nodes to the subordinate Fabric Interconnect ports on the UCS 6200 Series Fabric Interconnect.
7. Connect these cables to the corresponding ports on one of the new Cisco UCS 6400 Series Fabric Interconnects, using the connection records to preserve the port mapping. The servers may need to be reacknowledged (that is, may require a reboot) if the new server port numbers are changed.
8. Connect the L1/L2 cables that were disconnected from the old Fabric Interconnect to the new Cisco UCS 6400 Series Fabric Interconnect. L1 connects to L1, L2 connects to L2.
9. Connect the power to the new Cisco UCS 6400 Series Fabric Interconnect. It will automatically boot and run POST tests. If it reboots itself, this is normal behavior.

Important: Directly connect the console port to a terminal and observe the boot sequence. You should at some point see the Basic System Configuration Dialog, where you will configure the switch as a subordinate interconnect. If you do not see this dialog, either you have different builds of software on your old primary and new subordinate Fabric Interconnect, or the new subordinate was previously part of a cluster and will now need to have all configuration information wiped before it can be added to a cluster as a subordinate. In either case, immediately disconnect the L1 and L2 connections, bring the Fabric Interconnect up as standalone, then correct the issue before proceeding further. The primary and subordinate Fabric Interconnects must be running the same software build.

10. Enter the IP address details for the Fabric Interconnect. After this step, the installer will perform all the required configurations and sync with the primary Fabric Interconnect.

Note: If the installer detects any unsupported configuration, it will show the error in the console. You need to resolve any displayed configuration issues, then hit any key to restart the installer.

```
Enter the configuration method. (console/gui) ? console
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect will be added to the cluster. Continue (y/n) ? y
Enter the admin password of the peer Fabric interconnect:
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address:
Peer Fabric interconnect Mgmt0 IPv4 Netmask: 255.255.255.0
Cluster IPv4 address :
Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric Interconnect Mgmt0 IPv4 Address
Physical Switch Mgmt0 IP address :
Local fabric interconnect model(UCS-FI-6454)
Peer fabric interconnect model (UCS-FI-6248UP)
Installer is checking for compatibility with the peer fabric interconnect. Please wait...
=====
Multicast Optimize is enabled in the following QoS classes. Proceeding with migration will disable it
silver
=====
Local and peer fabric interconnect are compatible to migrate. Continuing with the installer...
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
```

Figure 12.
Subordinate Fabric Interconnect IP configuration

11. Once the configuration is complete, you will see the login prompt in the console. Log in to the new subordinate Fabric Interconnect and check the cluster status by running the command **show cluster extended-state**. Check and confirm that the management services are up on both Fabric Interconnects. Also, check and confirm that L1 and L2 (eth3 and eth4, as shown in Figure 13, below) report status as “UP.”

Note: An HA status of “NOT READY” is expected at this point until we enable the server ports that enable the Fabric Interconnects to form a cluster.

```
ironman-B#
ironman-B# show cluster extended-state
Cluster Id: 0x7880a79ed7391lea-0x9eb3547feef634a1

Start time: Fri Aug 21 16:03:55 2020
Last election time: Fri Aug 21 16:03:57 2020

B: UP, SUBORDINATE
A: UP, PRIMARY

B: memb state UP, lead state SUBORDINATE, mgmt services state: UP
A: memb state UP, lead state PRIMARY, mgmt services state: UP
  heartbeat state PRIMARY_OK

INTERNAL NETWORK INTERFACES:
eth3, UP
eth4, UP

HA NOT READY
No device connected to this Fabric Interconnect
ironman-B#
```

Figure 13.
Cluster status after replacing the subordinate Fabric Interconnect

12. Enable the uplink ports from the UCS Manager UI on the subordinate Fabric Interconnect.

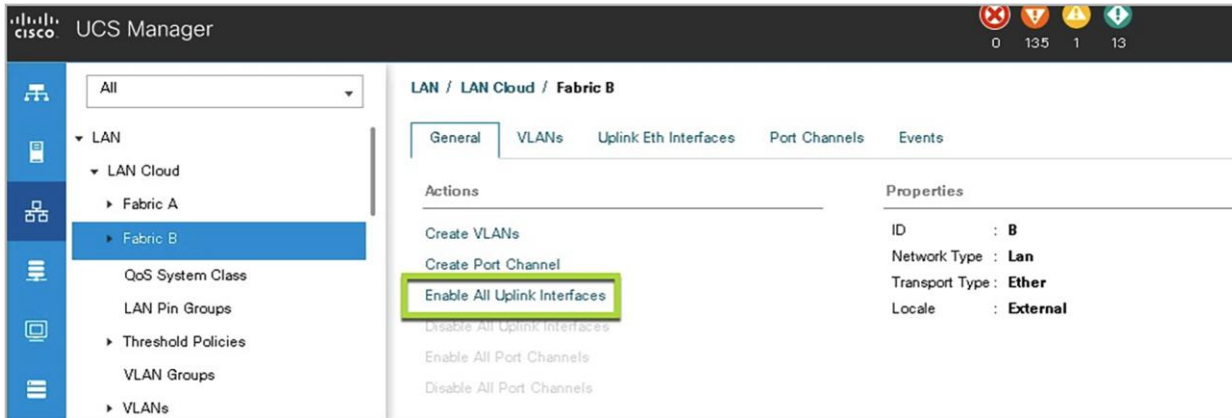


Figure 14.
Enable uplink interfaces

13. Confirm the uplink ports are up by running the commands below using the Fabric Interconnect CLI:

- # connect nxos
- # show interface status | grep -i uplink

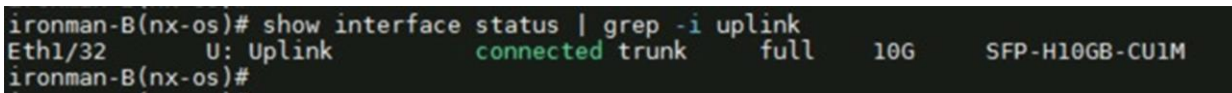


Figure 15.
Uplink interface status

14. Enable the previously disabled server ports from the UCS Manager UI.

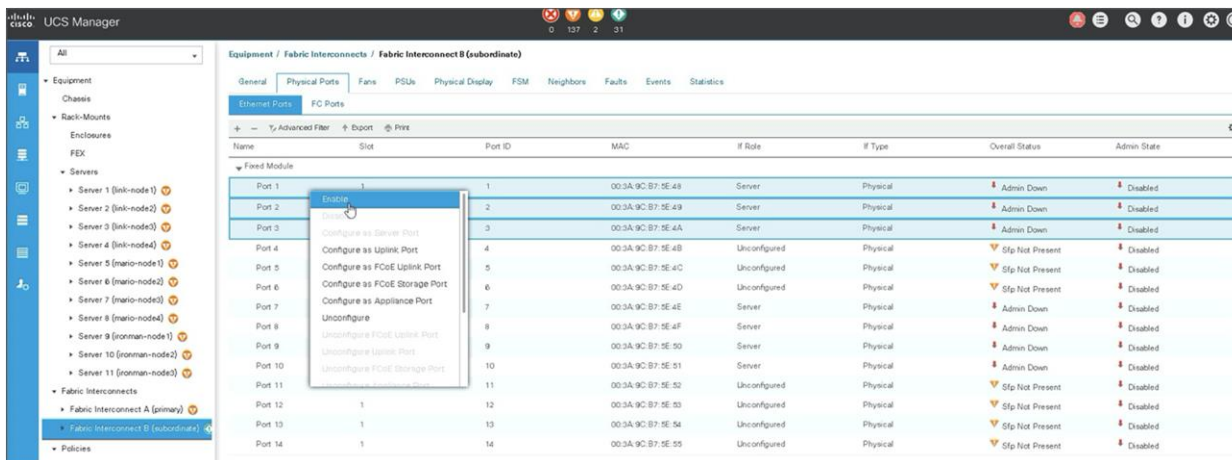


Figure 16.
Enable server interfaces

Note:

- If you have a Cisco UCS B-Series chassis connected to this pair of Fabric Interconnects and you had enabled Evacuation Mode, you can disable the evacuation mode now.
- Be patient. Once the server interfaces are enabled, the server will transition through a shallow discovery process (no reboot or traffic interruption will occur) and then the interfaces will be up. This process will take time to complete.

15. Verify the UCS cluster status and confirm that the HA status is ready.

```
ironman-A# show cluster extended-state
Cluster Id: 0x7880a79ed73911ea-0x9eb3547feef634a1

Start time: Thu Aug 6 19:00:03 2020
Last election time: Fri Aug 21 16:03:58 2020

A: UP, PRIMARY
B: UP, SUBORDINATE

A: memb state UP, lead state PRIMARY, mgmt services state: UP
B: memb state UP, lead state SUBORDINATE, mgmt services state: UP
  heartbeat state PRIMARY_OK

INTERNAL NETWORK INTERFACES:
eth1, UP
eth2, UP

HA READY
Detailed state of the device selected for HA storage:
Server 1, serial: WZP22330UVL, state: active
Server 2, serial: WZP22330UP4, state: active
Server 4, serial: WZP223315Z0, state: active
ironman-A#
ironman-A#
```

Figure 17.
UCS cluster HA status after subordinate Fabric Interconnect replacement

16. Verify that the data path is back up and redundant now from the hypervisor side.

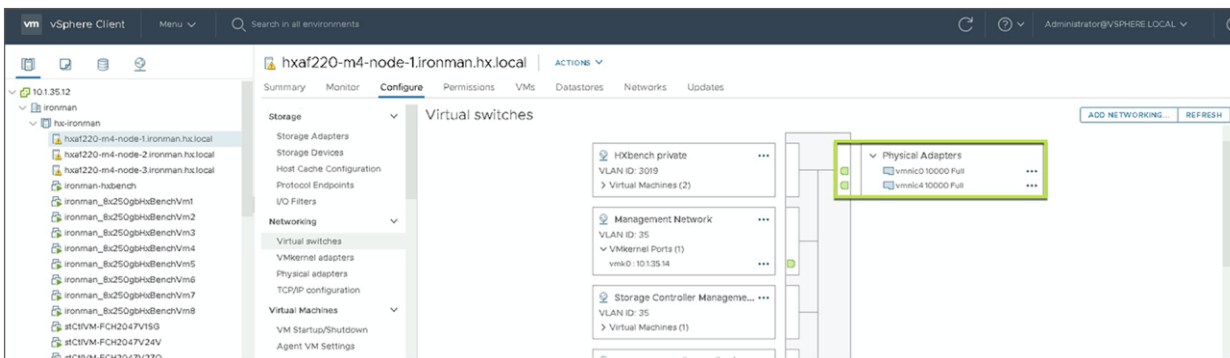


Figure 18.
Data path recovered

17. SSH to the UCSM cluster IP and run the commands given below to promote the subordinate Fabric Interconnect to primary. (Note: Verify that the cluster status is HA Ready before proceeding to the next step)

- # connect local-mgmt
- # cluster lead <a/b (current subordinate)>

```
ironman-A# connect local-mgmt
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2009, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
ironman-A(local-mgmt)# cluster lead b
If the system is at 'infrastructure firmware' auto-install 'pending user Ack' stage, this action will result in upgrading and rebooting c
urrent primary. Please check the outstanding faults (scope monitoring <enter> show new-faults) and make sure the data-paths on FI-B are
established properly before making it primary to ensure there is no data outage as the cluster lead change will acknowledge the reboot a
nd proceed with the upgrade.
Do you want to continue? (yes/no):yes
```

Figure 19.
Promoting the subordinate Fabric Interconnect

Note: You may lose connectivity to UCS Manager when this failover happens. If so, connect to the UCSM virtual IP again.

18. Check and verify that the UCS cluster status is healthy and HA status is “READY.”
19. Cable the second new Fabric Interconnect identically to the Fabric Interconnect that it is replacing and allow the configuration to be applied to the second new Fabric Interconnect.
20. Repeat the process from steps 1 through 16 for the second new Cisco UCS 6400 Series Fabric Interconnect.

Optional: Upgrade HyperFlex server uplinks to 25Gbps with Cisco VIC 1457

The fourth-generation Cisco UCS 6400 Fabric Interconnects offer line-rate, low-latency, lossless 10/25/40/100 Gigabit Ethernet, Fibre Channel over Ethernet (FCoE), and Fibre Channel functions. The new 6400 series Fabric Interconnect along with Cisco VIC 1457 adapters can be combined to achieve up to 100Gbps (50Gbps per Fabric Interconnect) network bandwidth from Fabric Interconnects toward the servers. Cisco VIC 1457 is a 10/25-Gbps quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M5 generation of Cisco HyperFlex and UCS servers. If your HyperFlex nodes have Cisco VIC 1457, you can upgrade the current southbound 10Gbps connectivity to 25Gbps without any downtime. New 25Gbps cables are required for the interfaces to negotiate at this higher speed. Refer to [Transceiver and Cable Support](#) section of the Cisco UCS Virtual Interface Cards (VICs) 1400 series data sheet for supported transceiver and cable details.

Note: Although this upgrade is an online upgrade with vNIC failover, it is recommended to plan the activity during a maintenance window.

Procedure:

Before you start the server uplink upgrade, make sure that the vNIC failover policies are correctly configured and that redundant paths are up.

1. Check and confirm that the redundant paths are up from the hypervisor side.
2. Note down the exact server port numbers on both Fabric Interconnects for each HyperFlex node.
3. SSH to the UCSM Management IP. Then do the following:
 - Run the command `show cluster state` to identify the subordinate Fabric Interconnect.
 - Run the command `connect nxos <a/b>` connect to the NX-OS shell of the subordinate Fabric Interconnect.
 - Run the command `show interface status` and identify the server ports.

```
ironman-B(nx-os)# show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--
Eth1/1	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/2	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/3	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/4	--	xcvrAbsen	1	auto	auto	--
Eth1/5	--	xcvrAbsen	1	auto	auto	--
Eth1/6	--	xcvrAbsen	1	auto	auto	--
Eth1/7	S: Server, Port-ch	connected	1	full	10G	SFP-H10GB-CU2M
Eth1/8	S: Server, Port-ch	connected	1	full	10G	SFP-H10GB-CU2M
Eth1/9	S: Server, Port-ch	connected	1	full	10G	SFP-H10GB-CU2M
Eth1/10	S: Server, Port-ch	connected	1	full	10G	SFP-H10GB-CU2M

Figure 20.
Fabric Interconnect server interfaces before the link upgrade

Note: As you see above, in Figure 20, the server ports are configured as port channels by default with VIC 1457.

4. Log in to UCS Manager and navigate to **Equipment -> Rack-Mounts -> Servers**, and select the first HyperFlex node. In the right-hand pane, go to VIF paths and expand all the paths. Check and confirm that all the virtual NICs are up and no “unpinned” virtual NICs are present.

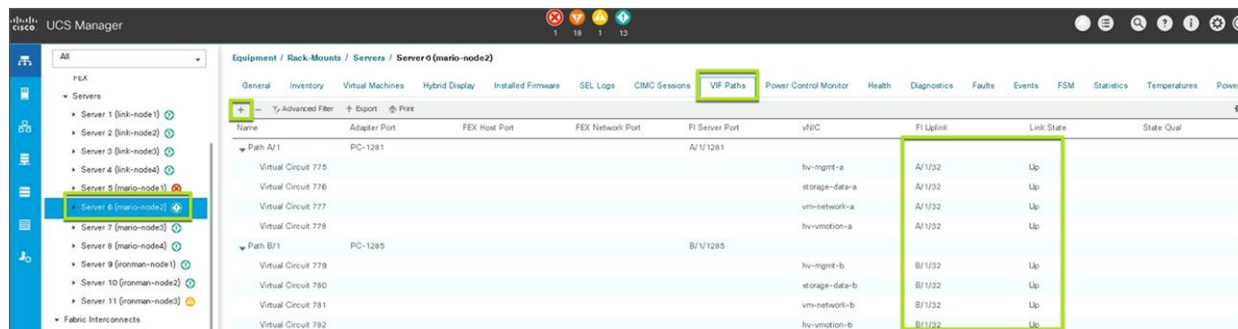


Figure 21.
HyperFlex node VIF paths

5. Check and confirm that the HyperFlex cluster state is online and healthy by running the command **hxcli cluster health**.
6. Disconnect the link from the HyperFlex node to the subordinate Fabric Interconnect.
7. Connect the new 25G-capable cable to the same Fabric Interconnect port that was previously connected.
8. From the UCS Fabric Interconnect CLI, verify that the physical link is up and is negotiated at 25G speed.

```
ironman-B(nx-os)# show interface status
```

Port	Name	Status	Vlan	Duplex	Speed	Type
mgmt0	--	connected	routed	full	1000	--

Port	Name	Status	Vlan	Duplex	Speed	Type
Eth1/1	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/2	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/3	S: Server	connected	1	full	10G	SFP-H10GB-CU1M
Eth1/4	--	xcvrAbsen	1	auto	auto	--
Eth1/5	--	xcvrAbsen	1	auto	auto	--
Eth1/6	--	xcvrAbsen	1	auto	auto	--
Eth1/7	S: Server, Port-ch	connected	1	full	25G	SFP-H25GB-CU3M
Eth1/8	S: Server, Port-ch	connected	1	full	25G	SFP-H25GB-CU3M
Eth1/9	S: Server, Port-ch	connected	1	full	25G	SFP-H25GB-CU3M
Eth1/10	S: Server, Port-ch	connected	1	full	25G	SFP-H25GB-CU3M

Figure 22.
Fabric Interconnect server interfaces after link upgrade

9. Wait for the virtual NICs to be pinned to the appropriate uplink ports as shown in Figure 22, above.

Note: It may take a few minutes for the dynamic pinning to happen and the vNICs are ready for failover. Check and make sure all the vNICs are in a pinned state before moving to the next step.
10. Check and verify that the virtual NIC in the hypervisor is up and showing a negotiated 25Gbps link speed.

All	Properties	CDP	LLDP	RDMA
Adapter	Cisco Systems Inc Cisco VIC Ethernet NIC			
Name	vnic0			
Location	PCI 0000:3f:00:0			
Driver	nenic			
Status				
Status	Connected			
Actual speed, Duplex	25 Gbit/s, Full Duplex			
Configured speed, Duplex	25 Gbit/s, Full Duplex			
Networks	No networks			
SR-IOV				
Status	Not supported			
Cisco Discovery Protocol				
Version	2			
Timeout	60			
Time to live	154			

Figure 23.
Hypervisor uplink speed

11. Disconnect the link from the same HyperFlex node to the primary Fabric Interconnect and replace with a new 25G-capable cable.
12. Repeat steps 8 through 10 for the second link to confirm the speed increase.
13. Repeat steps 4 through 12 on each node in the HyperFlex cluster.

Note: All the nodes in a HyperFlex cluster must be running with the same link speed. It is therefore imperative to ensure that all server links in the cluster are upgraded during the same maintenance window.

Optional: Adding more links for higher network bandwidth

Cisco VIC 1457 is a 10/25-Gbps quad-port Small Form-Factor Pluggable (SFP28) mLOM card designed for the M5 generation of Cisco HyperFlex servers. If your applications require higher network bandwidth and your HyperFlex servers have the VIC 1457 adapter installed, you have an option to double the network bandwidth to each Fabric Interconnect. This is accomplished by adding one additional network link from each HyperFlex node to each Fabric Interconnect.

Note: This is a complete online upgrade without any downtime. However, it is recommended to perform this upgrade in a maintenance window.

Procedure:

Before you start adding additional links, ensure that the new cables are of the same speed as the existing cables. For example, if your current connectivity is 25Gbps, add an identical 25Gbps cable as the second link to each fabric interconnect. Mixing the speeds is not a recommended configuration.

1. Check and confirm that the HyperFlex cluster state is online and healthy by running the command **hxcli cluster health**.
2. Connect the new cables as shown in Figure 24 below.




Figure 24.
Cisco VIC 1457 quad-port connections

- Log in to UCS Manager and configure the newly connected ports as server ports on both Fabric Interconnects.

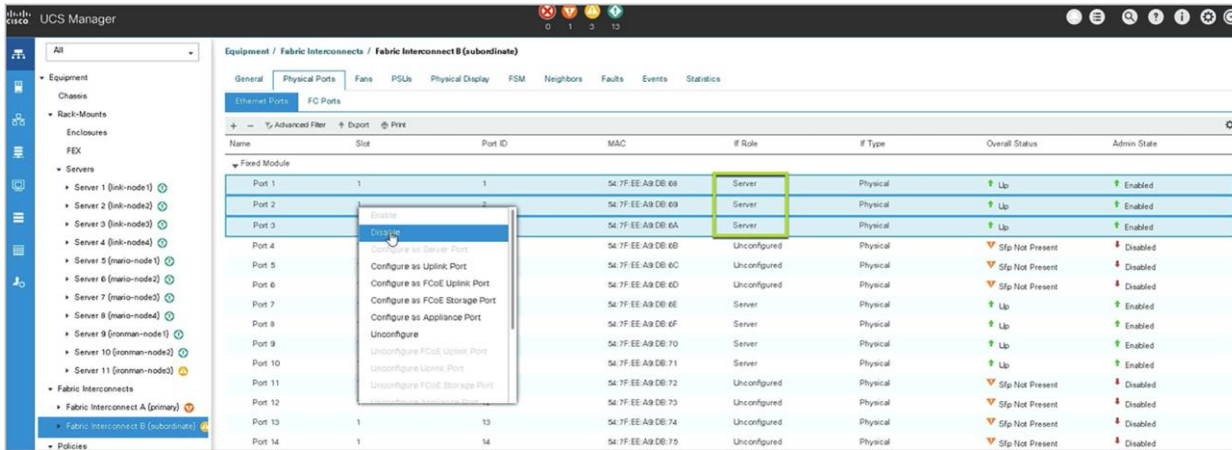


Figure 25.
Configure the server ports

- In UCS Manager, navigate to **Equipment -> Rack-Mounts -> Servers -> your HyperFlex Node -> Adapters -> Adapter 1 -> DCE Interfaces**, confirm that it shows up all four links now and that they are members of the port channel.
- Optionally log in to each of the UCS Fabric Interconnects and confirm that the port channel bandwidth is now doubled.

```
ironman-B(nx-os)# show interface port-channel 1285
port-channel1285 is up
admin state is up,
Hardware: Port-Channel, address: 003a.9cb7.5e4f (bia 003a.9cb7.5e4f)
Description: Server
MTU 9216 bytes, BW 50000000 Kbit, DLY 10 usec
reliability 255/255, txload 0/255, rxload 1/255
Encapsulation ARPA, medium is broadcast
Port mode is vntag
full-duplex, 25 Gb/s
Input flow-control is off, output flow-control is off
Auto-mdix is turned off
Switchport monitor is off
EtherType is 0x8100
Members in this channel: Eth1/8, Eth1/12
Last clearing of "show interface" counters never
3 interface resets
Load-Interval #1: 30 seconds
 30 seconds input rate 1152504 bits/sec, 391 packets/sec
 30 seconds output rate 1182568 bits/sec, 406 packets/sec
input rate 1.15 Mbps, 391 pps; output rate 1.18 Mbps, 406 pps
Load-Interval #2: 5 minute (300 seconds)
 300 seconds input rate 965512 bits/sec, 174 packets/sec
 300 seconds output rate 1094584 bits/sec, 160 packets/sec
input rate 965.51 Kbps, 174 pps; output rate 1.09 Mbps, 160 pps
RX
1329982538 unicast packets 253763 multicast packets 17172 broadcast packets
1330253473 input packets 2939880023121 bytes
318709809 jumbo packets 0 storm suppression bytes
0 runs 0 giants 0 CRC 0 no buffer
0 input error 0 short frame 0 overrun 0 underrun 0 ignored
0 watchdog 0 bad etype drop 0 bad proto drop 0 if down drop
0 input with dribble 7978 input discard
8 Rx pause
TX
1338625539 unicast packets 3700393 multicast packets 163645 broadcast packets
1342489577 output packets 3000859568296 bytes
333929236 jumbo packets
0 output error 0 collision 0 deferred 0 late collision
0 lost carrier 0 no carrier 0 babble 0 output discard
0 Tx pause
ironman-B(nx-os)#
```

Figure 26.
Server bandwidth doubled

Important Note: Due to a known cosmetic defect, both ESXi CLI and vCenter do not update their status to show the new link speed. However, the updated bandwidth is available for the workloads to consume when needed. This display issue can be resolved by a reboot of the server, which can be scheduled during the next planned patch upgrade or any other maintenance window. The steps below for server reboot are optional, if you wish to update the status right away.

6. Reboot each HyperFlex node one by one.
 - Log in to HyperFlex Connect and confirm that the cluster is online and healthy and that one or more node failures can be tolerated.

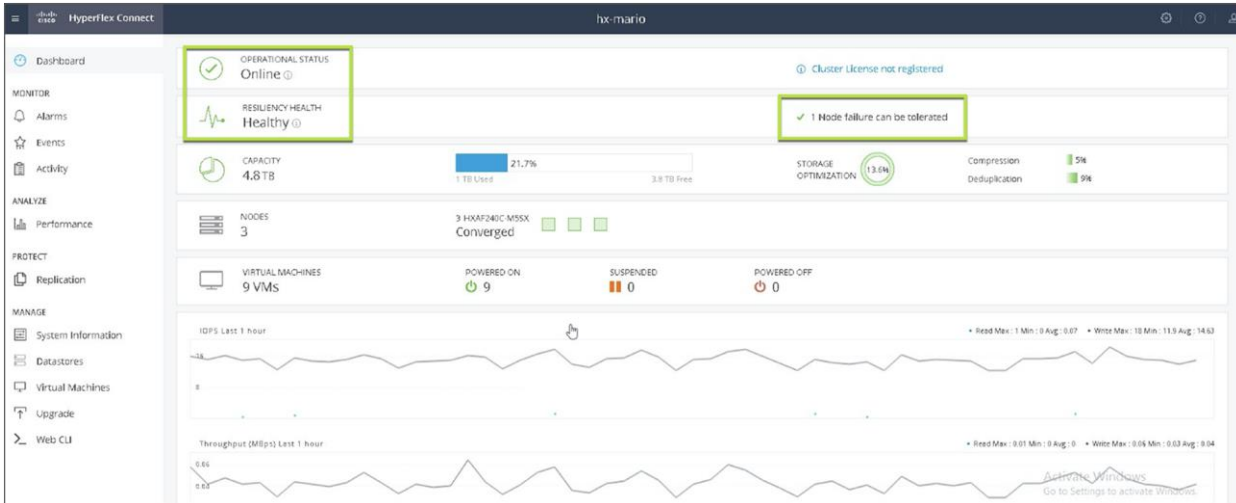


Figure 27.
HyperFlex cluster health

- Place a HyperFlex node in Maintenance Mode.

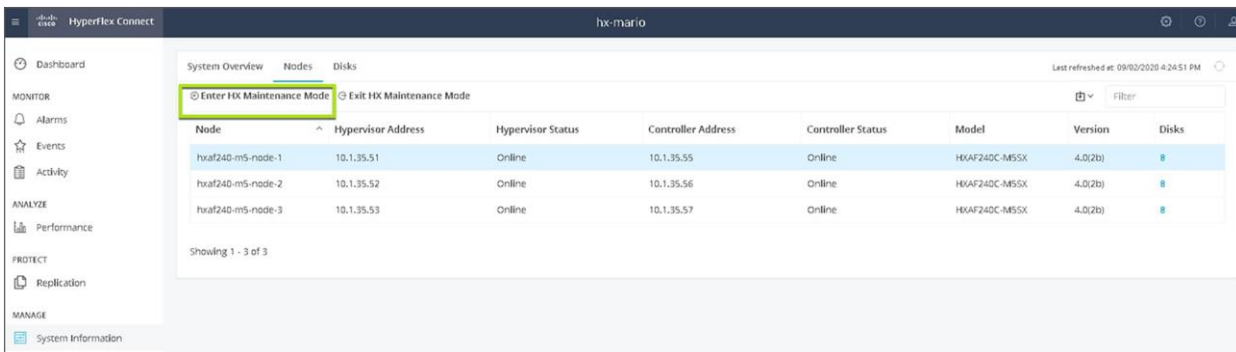


Figure 28.
Enter HyperFlex Maintenance Mode

- Once the node is successfully in Maintenance Mode, reboot the node from the vCenter.
- Wait for the node to reboot, and reestablish connectivity back to vCenter.
- Check and confirm that the new updated link speed is displayed in vCenter or the ESXi shell.

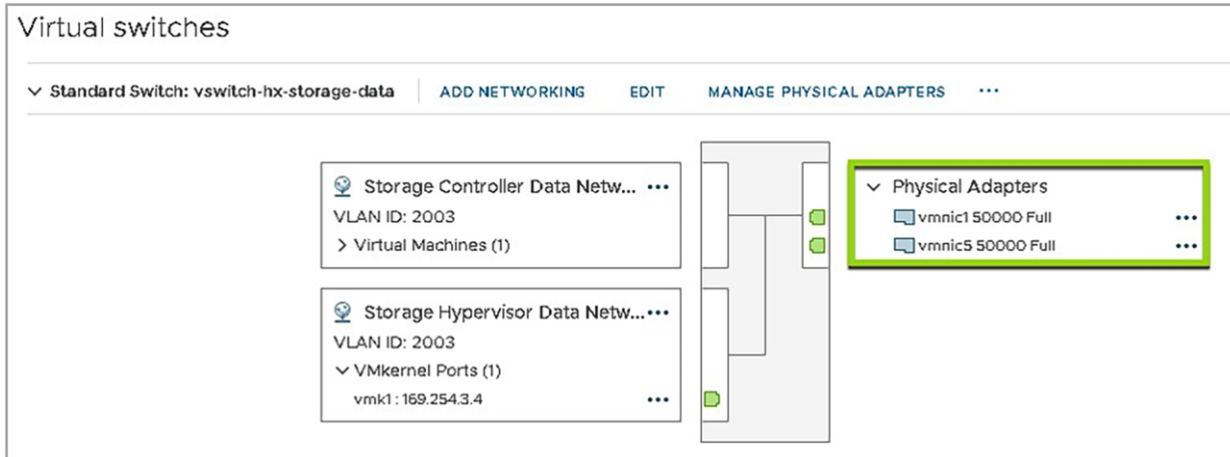


Figure 29.
Updated physical NIC bandwidth

- Exit the node from Maintenance Mode from HyperFlex Connect.

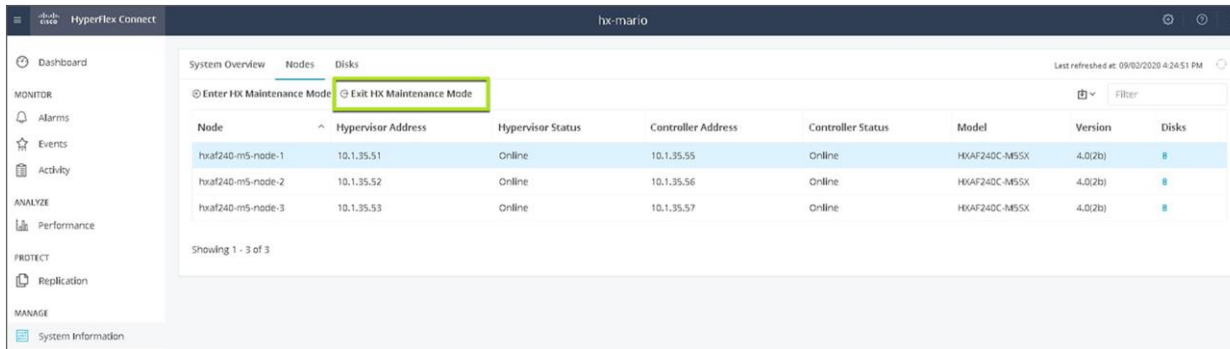


Figure 30.
Exit HyperFlex Maintenance Mode

- Wait for the cluster to be online and healthy and for node failures that can tolerate as 1 or higher.
- Repeat the steps above for the remaining nodes in the cluster.

For more information

Read more about Cisco HyperFlex:

<https://www.cisco.com/site/us/en/products/computing/hyperconverged-infrastructure/index.html>

Read more about the Cisco UCS 6454 Fabric Interconnect:

<https://www.cisco.com/c/en/us/products/servers-unified-computing/fabric-interconnects.html>

Read more about Cisco VIC 1457 adapters:

<https://www.cisco.com/c/en/us/products/collateral/interfaces-modules/unified-computing-system-adapters/datasheet-c78-741130.html#Productoverview>

Read more about UCS Fabric Interconnect hardware migration:

http://cs.co/ucs_fi_6200_to_6400_migration

Americas Headquarters

Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)