



Preparing for Cisco NFVI Installation

Before you can install and configure Cisco NFVI, you must complete the following hardware and application preparation procedures provided in the following topics.

- [Installing the Cisco NFVI Hardware, on page 1](#)
- [Configuring ToR Switches for C-Series Pods, on page 4](#)
- [Configuring ToR Switches for UCS B-Series Pods , on page 8](#)
- [Preparing Cisco IMC and Cisco UCS Manager, on page 10](#)
- [Installing the Management Node, on page 10](#)
- [Setting Up the UCS C-Series Pod, on page 13](#)
- [Setting Up the UCS B-Series Pod, on page 16](#)
- [Configuring the Out-of-Band Management Switch, on page 17](#)

Installing the Cisco NFVI Hardware

Before installing Cisco Virtualization Infrastructure Manager (VIM), switch on the Cisco UCS C-Series or B-Series hardware. Depending upon the pod type, the CIMC connection or UCSM IP has to be set up ahead of time. Following table lists the UCS hardware options and network connectivity protocol that can be used with either virtual extensible LAN (VXLAN) over a Linux bridge, VLAN over OVS or VLAN over VPP. If Cisco Virtual Topology Services, an optional Cisco NFVI application, is installed, Virtual Topology Forwarder (VTF) can be used with VXLAN for tenants, and VLANs for providers on C-Series pods.

Table 1: Cisco NFVI Hardware and Network Connectivity Protocol

UCS Pod Type	Compute and Controller Node	Storage Node	Network Connectivity Protocol
C-Series	UCS C220/240 M4.	UCS C240 M4 (SFF) with two internal SSDs.	VXLAN/Linux Bridge or OVS/VLAN or VPP/VLAN, or ACI/VLAN.
C-Series	Controller: UCS C220/240 Compute: HP DL360 Gen9	UCS C240 M4 (SFF) with two internal SSDs.	OVS/VLAN

UCS Pod Type	Compute and Controller Node	Storage Node	Network Connectivity Protocol
C-Series with Cisco VTS	UCS C220/240 M4.	UCS C240 M4 (SFF) with two internal SSDs.	For tenants: VTF with VXLAN. For providers: VLAN
C-Series Micro Pod	<p>UCS 240 M4 /M5 with 12 HDD and 2 external SSDs. Pod can be expanded to 18 computes. Each compute has 2x1.2 TB HDD</p> <p>Or</p> <p>UCS 220 M4/M5 with 7 HDD and 1 external SSD. Pod can be expanded to 18 computes. Each compute has 2x1.2 TB HDD</p> <p>Note Refer the BOM for SSD-based install for M5; Also the M5 BOM is based on Intel X710 for control and data plane and XL710 for SRIOV</p>	Not Applicable as C-Series Micro Pod is integrated with Compute, and Controller.	<p>UCS M4 Support: OVS/VLAN or VPP/VLAN or ACI/VLAN.</p> <p>UCS M5 Support: OVS/VLAN or VPP/VLAN</p>
C-Series Hyperconverged	UCS 240 M4.	UCS C240 M4 (SFF) with 12 HDD and two external SSDs, also acts a compute node.	OVS/VLAN
B-Series	UCS B200 M4.	UCS C240 M4 (SFF) with two internal SSDs.	VXLAN/Linux Bridge or OVS/VLAN.
B-Series with UCS Manager Plugin	UCS B200 M4s	UCS C240 M4 (SFF) with two internal SSDs.	OVS/VLAN



Note The storage nodes boot off two internal SSDs. It also has four external solid state drives (SSDs) for journaling, which gives a 1:5 SSD-to-disk ratio (assuming a chassis that is filled with 20 spinning disks). Each C-Series pod has either a 2 port 10-GE Cisco vNIC 1227 card or 2 of 4 port Intel 710 X card. UCS M4 blades only support Cisco 1340 and 1380 NICs. For more information about the Cisco vNICs, see [LAN and SAN Connectivity for a Cisco UCS Blade](#). Cisco VIM has a micro pod (based on UCS M4 hardware) which works on Cisco VIC 1227 or Intel NIC 710, with OVS/VLAN or VPP/VLAN as the virtual network protocol. The micro pod supports users with a small, functional, but redundant cloud with capability of adding standalone computes to an existing pod. The manifestation of the UCS-M4 micro-pod works on Cisco VIC (1227) or Intel NIC 710, with OVS/VLAN of VPP/VLAN as the virtual network protocol.

Cisco VIM supports M4-based micropod on a VIC/NIC system with OVS, to extend the SRIOV support on a 2x2-port Intel 520 NIC card. Also, the M5 based Micro pod is based on Intel NIC 710, and supports SRIOV over XL710, with OVS/VLAN or VPP/VLAN as the virtual network protocol. In Cisco VIM 2.4.1, we have added the support of M5-based micropod on a VIC/NIC system, to extend the SRIOV support on a 2x2 port Intel XL710 NIC card.

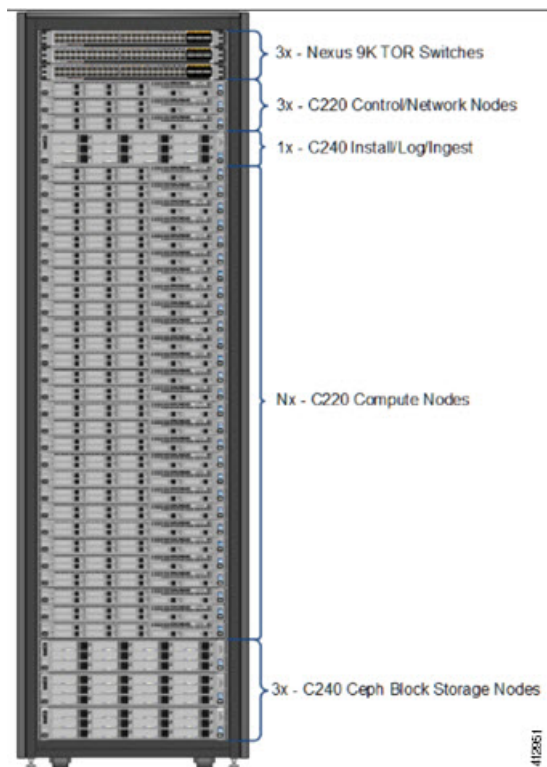
In addition, the Cisco Nexus 9372 or 93180YC, or 9396PX must be available to serve the Cisco NFVI ToR function.

After verifying that you have the required Cisco UCS servers and blades and the Nexus 93xx, install the hardware using the following links:

- [Cisco UCS C220 M4 Server Installation and Service Guide](#)
- [Cisco UCS C240 M4 Server Installation and Service Guide](#)
- [Cisco UCS B200 Blade Server and Installation Note](#)
- [Cisco Nexus 93180YC, 9396PX, 9372PS, and 9372PX-E NX-OS Mode Switches Hardware Installation Guide](#)

The following figure shows a C-Series Cisco NFVI pod. Although the figure shows that a full complement of UCS C220 computes nodes, the number of compute nodes can vary, depending on the implementation requirements. The UCS C220 control and compute nodes can be replaced with UCS 240 series. However in that case the number of computes fitting in one chassis system is reduced by half.

Figure 1: Cisco NFVI C-Series Pod



Note The combination of UCS-220 and 240 within the compute and control nodes is not supported.

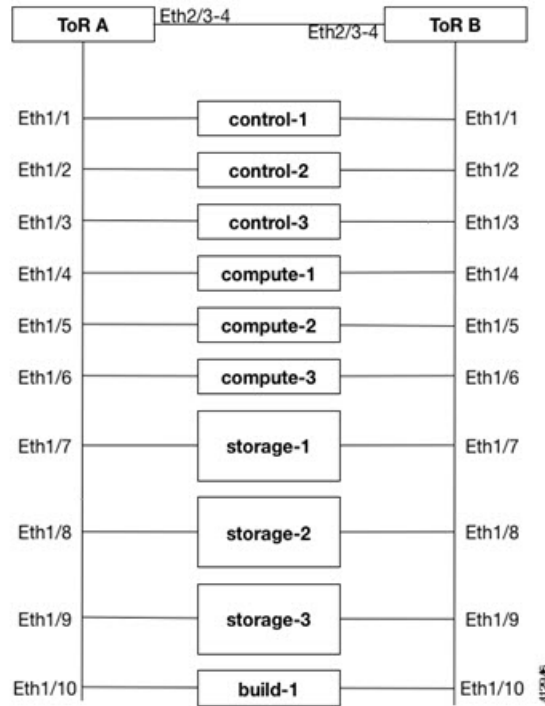
For more information on wiring schematic of various pod configurations, refer to Appendix section.

Configuring ToR Switches for C-Series Pods

During installation, the Cisco VIM installer creates vNIC's on each of the two physical interfaces and creates a bond for the UCS C-Series pod. Before this you have to manually configure the ToR switches to create a vPC with the two interfaces connected to each server. Use identical Cisco Nexus 9372, or 93180YC, or 9396PX switches for the ToRs. We recommend you to use the N9K TOR software versions for setup: 7.0(3)I4(6) 7.0(3)I6(1). Also, you can refer to, the Appendix section for the wiring details for each pod type on a C-series-based install.

Complete the following steps to create a vPC on a pair of Cisco Nexus ToR switches. The steps use the following topology as an example. Modify the configuration as it applies to your environment. In Cisco VIM, we have introduced a feature which is called auto-configuration of ToR (for N9K series only). This is an optional feature, and if you decide to take this route, the following steps can be skipped.

Figure 2: ToR Configuration Sample



Step 1

Change the vPC domain ID for your configuration. The vPC domain ID can be any number as long as it is unique. The IP address on the other switch mgmt0 port is used for the keepalive IP. Change it to the IP used for your network.

For the preceding example, the following is the configuration:

```
ToR-A (mgmt0 is 172.18.116.185)
feature vpc
vpc domain 116
peer-keepalive destination 172.18.116.186
ToR-B (mgmt0 is 172.18.116.186)
feature vpc
vpc domain 116
peer-keepalive destination 172.18.116.185
```

Because both switches are cabled identically, the remaining configuration is identical on both switches. In this example, topology Eth2/3 and Eth2/4 are connected to each other and combined into a port channel that functions as the vPC peer link.

```
feature lACP
interface Ethernet2/3-4
channel-group 116 mode active
interface port-channel116
switchport mode trunk
vpc peer-link
```

Step 2

For each VLAN type, (mgmt_vlan, tenant_vlan_range, storage, api, external, provider), execute the following on each ToR:

```
vlan <vlan_type>
  no shut
```

Step 3 Configure all the interfaces are connected to the servers to be members of the port channels. In the example, only ten interfaces are shown. But you must configure all interfaces that are connected to the server.

Note *If interfaces have configurations from previous deployments, you can remove them by entering **default interface Eth1/1-10**, then **no interface Po1-10**.*

1. For deployment with any mechanism driver on Cisco VIC

There is no configuration differences among different roles (controllers/computes/storages). The same configuration applies to all interfaces.

```
interface Ethernet 1/1
channel-group 1 mode active
interface Ethernet 1/2
channel-group 2 mode active
interface Ethernet 1/3
channel-group 3 mode active
interface Ethernet 1/4
channel-group 4 mode active
interface Ethernet 1/5
channel-group 5 mode active
interface Ethernet 1/6
channel-group 6 mode active
interface Ethernet 1/7
channel-group 7 mode active
interface Ethernet 1/8
channel-group 8 mode active
interface Ethernet 1/9
channel-group 9 mode active
interface Ethernet 1/10
channel-group 10 mode active
```

2. For deployment with OVS/VLAN or LinuxBridge on Intel VIC

The interface configuration is the same as Cisco VIC case as shown in the preceding section. However, number of switch interfaces that are configured is more in the case of Intel NIC as we have dedicated control, and data physical ports participating in the case of Intel NIC. Also for SRIOV switchport, no port channel is configured, and the participating VLAN can be in trunk mode.

3. For deployment with VPP/VLAN or VTS on Intel VIC

In this case VPP/VLAN or VTS is used as the mechanism driver. The interface configuration varies based on the server roles. Assume Ethernet1/1 to Ethernet1/3 are controller interfaces, Ethernet1/4 to Ethernet1/6 are storage interfaces, and Ethernet1/7 to Ethernet1/10 are the compute interfaces. The sample configurations looks like:

```
interface Ethernet 1/1
channel-group 1 mode active
interface Ethernet 1/2
channel-group 2 mode active
interface Ethernet 1/3
channel-group 3 mode active
interface Ethernet 1/4
channel-group 4 mode active
interface Ethernet 1/5
channel-group 5 mode active
interface Ethernet 1/6
channel-group 6 mode active
interface Ethernet 1/7
channel-group 7
```

```

interface Ethernet 1/8
channel-group 8
interface Ethernet 1/9
channel-group 9
interface Ethernet 1/10
channel-group 10

```

Note When using VPP/VLAN with Intel NIC, ensure that LACP is turned off for those port channels that are connected to the compute nodes. In the sample configurations that is given, preceding codes corresponds to Ethernet 1/7 to 1/10"

Step 4 Configure the port channel interface to be vPC, and trunk all VLANs. Native vlan has to be configured and set to the mgmt vlan on the control ports for Intel NIC. Without it PXE boot fails. Skip to listen or learn in spanning tree transitions, and do not suspend the ports if they do not receive LACP packets. Also, configure it with large MTU of 9216 (this is important or else Ceph install hangs).The last configuration allows you to start the servers before the bonding is set up.

```

interface port-channel1-9
shutdown
spanning-tree port type edge trunk
switchport mode trunk
switchport trunk native vlan mgmt_vlan for the control ports when Intel NIC is used
switchport trunk allowed vlan <mgmt_vlan, tenant_vlan_range, storage, api, external, provider>
no lACP suspend-individual
mtu 9216
vpc <1-9>
no shutdown

```

Step 5 Identify the port channel interface that connects to the management node on the ToR:

```

interface port-channel10
shutdown
spanning-tree port type edge trunk
switchport mode trunk
switchport trunk allowed vlan <mgmt_vlan>
no lACP suspend-individual
vpc 10
no shutdown

```

Step 6 Check the port channel summary status. The ports connected to the neighbor switch have to be in (P) state. Before the server installation, the server facing interfaces has to be in (I) state. After installation, they have to be in (P) state, which means they are up and in port channel mode.

```

gen-leaf-1# show port-channel summary
Flags: D - Down P - Up in port-channel (members)
I - Individual H - Hot-standby (LACP only)
s - Suspended r - Module-removed
S - Switched R - Routed
U - Up (port-channel)
M - Not in use. Min-links not met

```

```

-----
Group Port- Type Protocol Member Ports
Channel
-----

```

```

1 Po1(SD) Eth LACP Eth1/1(I)
2 Po2(SD) Eth LACP Eth1/2(I)
3 Po3(SD) Eth LACP Eth1/3(I)
4 Po4(SD) Eth LACP Eth1/4(I)
5 Po5(SD) Eth LACP Eth1/5(I)
6 Po6(SD) Eth LACP Eth1/6(I)
7 Po7(SD) Eth LACP Eth1/7(I)
8 Po8(SD) Eth LACP Eth1/8(I)

```

```

9 Po9(SD) Eth LACP Eth1/9(I)
10 Po10(SD) Eth LACP Eth1/10(I)
116 Po116(SU) Eth LACP Eth2/3(P) Eth2/4(P)

```

Step 7 Enable automatic Cisco NX-OS errdisable state recovery:

```

errdisable recovery cause link-flap
errdisable recovery interval 30

```

Cisco NX-OS places links that flap repeatedly into errdisable state to prevent spanning tree convergence problems caused by non-functioning hardware. During Cisco VIM installation, the server occasionally triggers the link flap threshold, so enabling automatic recovery from this error is recommended.

```

errdisable recovery cause link-flap
errdisable recovery interval 30

```

Step 8 If you are installing Cisco Virtual Topology Systems, an optional Cisco NFVI application, enable jumbo packets and configure 9216 MTU on the port channel or Ethernet interfaces. For example:

Port channel:

```

interface port-channel10
  switchport mode trunk
  switchport trunk allowed vlan 80,323,680,860,2680,3122-3250
  mtu 9216
  vpc 10

```

Ethernet:

```

interface Ethernet1/25
  switchport mode trunk
  switchport trunk allowed vlan 80,323,680,860,2680,3122-3250
  mtu 9216

```

Configuring ToR Switches for UCS B-Series Pods

Complete the following steps to create a vPC on a pair of Cisco Nexus ToR switches for a UCS B-Series pod. The steps are similar to configuring ToR switches for C-Series pods, with some differences. In the steps, the two ToR switches are Storm-tor-1 (mgmt0 is 172.18.116.185), and Storm-tor-2 (mgmt0 is 172.18.116.186). Modify the configuration as it applies to your environment. If no multicast or QOS configuration is required, and Auto-configuration of TOR is chosen as an option, the steps listed below can be skipped.

Before you begin

Step 1 Change the vPC domain ID for your configuration. The vPC domain ID can be any number as long as it is unique. The IP address on the other switch mgmt0 port is used for the keepalive IP. Please change it to the IP used for your network. Storm-tor-1 (mgmt0 is 172.18.116.185)

a)

```

feature vpc
vpc domain 116
  peer-keepalive destination 172.18.116.186

```



```

for each vlan_type (mgmt_vlan, tenant_vlan_range, storage, api, external, provider); # execute the
following for each vlan
  vlan <vlan_type>
  no shut
vrf context management
  ip route 0.0.0.0/0 172.18.116.1

interface mgmt0
  vrf member management
  ip address 172.18.116.185/24

```

Storm-tor-2 (mgmt0 is 172.18.116.186)

```

feature vpc
vpc domain 116
  peer-keepalive destination 172.18.116.185
for each vlan_type (mgmt_vlan, tenant_vlan_range, storage, api, external, provider); # execute the
following for each vlan
  vlan <vlan_type>
  no shut
vrf context management
  ip route 0.0.0.0/0 172.18.116.1

interface mgmt0
  vrf member management
  ip address 172.18.116.186/24

```

Step 2 Since both switches are cabled identically, the rest of the configuration is identical on both switches. Configure all the interfaces connected to the fabric interconnects to be in the VPC as well.

```

feature lacp
interface port-channel1
  description "to fabric interconnect 1"
  switchport mode trunk
  vpc 1
interface port-channel2
  description "to fabric interconnect 2"
  switchport mode trunk
  vpc 2
interface Ethernet1/43
  description "to fabric interconnect 1"
  switchport mode trunk
  channel-group 1 mode active
interface Ethernet1/44
  description "to fabric interconnect 2"
  switchport mode trunk
  channel-group 2 mode active

```

Step 3 Create the port-channel interface on the ToR that is connecting to the management node:

```

interface port-channel3
  description "to management node"
  spanning-tree port type edge trunk
  switchport mode trunk
  switchport trunk allowed vlan <mgmt_vlan>
  no lacp suspend-individual
  vpc 3
interface Ethernet1/2
  description "to management node"
  switchport mode trunk
  channel-group 3 mode active

```

Step 4 Enable jumbo frames for each ToR port-channel that connects to the Fabric Interconnects:

```
interface port-channel<number>
  mtu 9216
```

Note You must also enable jumbo frames in the `setup_data.yaml` file. See the UCS Manager Common Access Information for B-Series Pods topic in [Setting Up the Cisco VIM Data Configurations](#)

Preparing Cisco IMC and Cisco UCS Manager

Cisco NFVI requires specific Cisco Integrated Management Controller (IMC) and Cisco UCS Manager firmware versions and parameters. The Cisco VIM bare metal installation uses the Cisco IMC credentials to access the server Cisco IMC interface, which is used to delete and create vNICs and to create bonds.

Complete the following steps to verify Cisco IMC and UCS Manager are ready for Cisco NFVI installation:

-
- Step 1** Verify that each Cisco UCS server has one of the following Cisco IMC firmware versions is running 2.0(13i) or greater (preferably 2.0(13n)). Cisco IMC version can be 3.0 series (use 3.0.3(f) or later).. The latest Cisco IMC ISO image can be downloaded from the Cisco Software Download site. For upgrade procedures, see the [Cisco UCS C-Series Rack-Mount Server BIOS Upgrade Guide](#).
 - Step 2** For UCS B-Series pods, verify that the Cisco UCS Manager version is one of the following: 2.2(5a), 2.2(5b), 2.2(6c), 2.2(6e), 3.1(c).
 - Step 3** For UCS C-Series pods, verify the following Cisco IMC information is added: IP address, username, and password.
 - Step 4** For UCS B-Series pods, verify the following UCS Manager information is added: username, password, IP address, and resource prefix. The resource prefix maximum length is 6. The provisioning network and the UCS Manager IP address must be connected.
 - Step 5** Verify that no legacy DHCP/Cobbler/PXE servers are connected to your UCS servers. If so, disconnect or disable the interface connected to legacy DHCP, Cobbler, or PXE server. Also, delete the system from the legacy cobbler server.
 - Step 6** Verify Cisco IMC has NTP enabled and is set to the same NTP server and time zone as the operating system.
-

Installing the Management Node

This procedure installs RHEL 7.4 with the following modifications:

- Hard disk drives are setup in RAID 6 configuration with one spare HDD for eight HDDs deployment, two spare HDDs for 9 to 16 HDDs deployment, or four spare HDDs for 17 to 24 HDDs deployment
- Networking—Two bridge interfaces are created, one for the installer API and one for provisioning. Each bridge interface has underlying interfaces bonded together with 802.3ad. Provision interfaces are 10 GE Cisco VICs. API interfaces are 1 GE LOMs. If the NFVIBENCH, is planned to be used, another 2xIntel 520 or 4xIntel710 X is needed.
- The installer code is placed in `/root/`.
- SELinux is enabled on the management node for security.

Before you begin

Verify that the Cisco NFVI management node where you plan to install the Red Hat for Enterprise Linux (RHEL) operating system is a Cisco UCS C240 M4 Small Form Factor (SFF) with 8, 16 or 24 hard disk drives (HDDs). In addition, the management node must be connected to your enterprise NTP and DNS servers. If your management node server does not meet these requirements, do not continue until you install a qualified UCS C240 server. Also, verify that the pod has MRAID card.

Step 1 Log into the Cisco NFVI management node.

Step 2 Follow steps in [Configuring the Server Boot Order](#) to set the boot order to boot from Local HDD.

Step 3 Follow steps in Cisco UCS [Configure BIOS Parameters](#) to set the following advanced BIOS settings:

For Management node based on UCS M4 boxes set the following for BIOS Parameters:

- PCI ROM CLP—Disabled
- PCH SATA Mode—AHCI
- All Onboard LOM Ports—Enabled
- LOM Port 1 OptionROM—Disabled
- LOM Port 2 OptionROM—Disabled
- All PCIe Slots OptionROM—Enabled
- PCIe Slot:1 OptionROM—Enabled
- PCIe Slot:2 OptionROM—Enabled
- PCIe Slot: MLOM OptionROM—Disabled
- PCIe Slot:HBA OptionROM—Enabled
- PCIe Slot:FrontPcie1 OptionROM—Enabled
- PCIe Slot:MLOM Link Speed—GEN3
- PCIe Slot:Riser1 Link Speed—GEN3
- PCIe Slot:Riser2 Link Speed—GEN3
- • MLOM OptionROM—Enabled

Step 4 Click **Save Changes**.

Step 5 Add the management node vNICs to the provisioning VLAN to provide the management node with access to the provisioning network:

- a) In the CIMC navigation area, click the **Server** tab and select **Inventory**.
- b) In the main window, click the **Cisco VIC Adapters** tab.
- c) Under Adapter Card, click the **vNICs** tab.
- d) Click the first vNIC and choose **Properties**.
- e) In the vNIC Properties dialog box, enter the provisioning VLAN in the Default VLAN field and click **Save Changes**.
- f) Repeat Steps **a** through **e** for the second vNIC.

Note Delete any additional vNICs configured on the UCS server beyond the two default ones.

- Step 6** Download the Cisco VIM ISO image to your computer from the location that is provided to you by the account team.
- Step 7** In CIMC, launch the KVM console.
- Step 8** Mount the Cisco VIM ISO image as a virtual DVD.
- Step 9** Reboot the UCS server, then press **F6** to enter the boot menu.
- Step 10** Select the KVM-mapped DVD to boot the Cisco VIM ISO image supplied with the install artifacts.
- Step 11** When the boot menu appears, select **Install Cisco VIM Management Node**. This is the default selection, and is automatically be chosen after the timeout.
- Step 12** At the prompts, answer the following questions to install the Management node as unified management node only or not:

- Hostname—Enter the management node hostname (The hostname length must be 32 or less characters).
- Select Yes to Install as Unified Management only when the node is only planned to be used as unified management node; Migration from one to another is not supported)
- API IPv4 address—Enter the management node API IPv4 address in CIDR (Classless Inter-Domain Routing) format. For example, 172.29.86.62/26
- API Gateway IPv4 address—Enter the API network default gateway IPv4 address.
- MGMT IPv4 address—Enter the management node MGMT IPv4 address in CIDR format. For example, 10.30.118.69/26

Note This field is not be prompted if the management node is installed as “unified management node only”

- Prompt to enable static IPv6 address configuration—Enter 'yes' to continue input similar IPv6 address configuration for API and MGMT network or 'no' to skip if IPv6 is not needed.
- API IPv6 address—Enter the management node API IPv6 address in CIDR (Classless Inter-Domain Routing) format. For example, 2001:c5c0:1234:5678:1001::5/8.
- Gateway IPv6 address—Enter the API network default gateway IPv6 address.
- MGMT IPv6 address—Enter the management node MGMT IPv6 address in CIDR format. For example, 2001:c5c0:1234:5678:1002::5/80
- DNS server—Enter the DNS server IPv4 address or IPv6 address if static IPv6 address is enabled.
- Option for Teaming Drier for Link Aggregation (answer ‘yes’ when Nexus Switch is the TOR, and answer ‘no’ when NCS-5500 is TOR): <yes|no> "

After you enter the management node IP addresses, the Installation options menu appears. In the installation menu, there are more options, fill in the options that are listed below (option8 and 2) and leave everything else as it is. If there is problem to start the installation, enter "r" to refresh the Installation menu.

- Step 13** In the Installation menu, select option **8** to enter the root password.
- Step 14** At the password prompts, enter the root password, then enter it again to confirm.
- Step 15** At the Installation Menu, select option **2** to enter the time zone.
- Step 16** At the Timezone settings prompt, select the option 1 as option 2 is not supported.
- Step 17** Enter the number corresponding to your time zone.
- Step 18** At the next prompt, enter the number for your region.
- Step 19** At the next prompt, choose the city, then confirm the time zone settings.

Step 20 After confirming your time zone settings, enter **b** to start the installation.

Step 21 After the installation is complete, press **Return** to reboot the server.

Step 22 After the reboot, check the management node clock using the Linux **date** command to ensure that the TLS certificates are valid, for example:

```
#date
Mon Aug 22 05:36:39 PDT 2016

To set date:
#date -s '2016-08-21 22:40:00'
Sun Aug 21 22:40:00 PDT 2016

To check for date:
#date
Sun Aug 21 22:40:02 PDT 2016
```

Setting Up the UCS C-Series Pod

After you install the RHEL OS on the management node, perform the following steps to set up the Cisco UCS C-Series servers:

Step 1 Follow steps in Configuring the Server Boot Order to set the boot order to boot from Local HDD.

Step 2 Follow steps in Configure BIOS Parameters to set the LOM, HBA, and PCIe slots to the following settings:

For servers based on UCS M4 boxes set the following for BIOS Parameters:

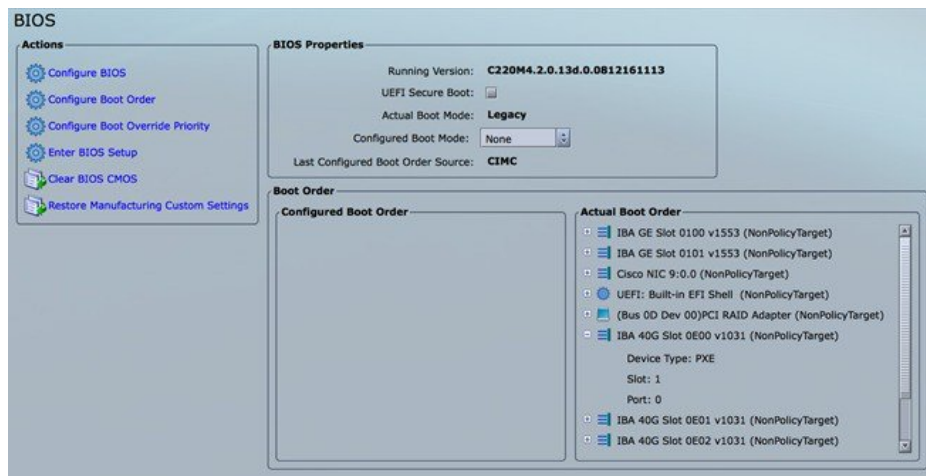
- CDN Support for VIC—Disabled
- PCI ROM CLP—Disabled
- PCH SATA Mode—AHCI
- All Onboard LOM Ports—Enabled
- LOM Port 1 OptionROM—Disabled
- LOM Port 2 OptionROM—Disabled
- All PCIe Slots OptionROM—Enabled
- PCIe Slot:1 OptionROM—Enabled
- PCIe Slot:2 OptionROM—Enabled
- PCIe Slot: MLOM OptionROM—Enabled
- PCIe Slot:HBA OptionROM—Enabled
- PCIe Slot:N1 OptionROM—Enabled
- PCIe Slot:N2 OptionROM—Enabled
- PCIe Slot:HBA Link Speed—GEN3

• For servers based on UCS M5 boxes set the following for BIOS Parameters:

- • All Onboard LOM Ports—Enabled
- • LOM Port 1 OptionROM—Disabled
- • LOM Port 2 OptionROM—Disabled
- • PCIe Slot:1 OptionROM—Enabled
- • PCIe Slot:2 OptionROM—Enabled
- • MLOM OptionROM—Enabled
- • MRAID OptionROM—Enabled

For other parameters, leave it at their default settings.

- More steps have to be taken to setup C-series pod with Intel NIC. In the Intel NIC testbed, each C-series server has 2, 4-port Intel 710 NIC cards. Ports A, B, and C for each Intel NIC card has to be connected to the respective TOR. Also, ensure that the PCI slot in which the Intel NIC cards are inserted are enabled in the BIOS setting (BIOS > Configure BIOS > Advanced > LOM and PCI Slot Configuration -> All PCIe Slots OptionROM-Enabled and enable respective slots). To identify the slots, check the slot-id information under the Network-Adapter tab listed under the Inventory link on the CIMC pane. All the Intel NIC ports should be displayed in the BIOS summary page under the Actual Boot Order pane, as IBA 40G Slot xyza with Device Type is set to PXE.



For UCS M5 look for “IBA 40G Slot ...” under the BIOS Properties



In case, the boot-order for the Intel NICs are not listed as above, the following one-time manual step needs to be executed to flash the Intel NIC x710 to enable PXE.

1. Boot each server with a CentOS image.
2. Download Intel Ethernet Flash Firmware Utility (Preboot.tar.gz) for X710 from the above link for Linux platform: <https://www.intel.com/content/www/us/en/support/software/manageability-products/000005790.html>.
3. Copy downloaded PREBOOT.tar to UCS server having X710 card.

```
mkdir -p /tmp/Intel/
tar xvf PREBOOT.tar -C /tmp/Intel/
cd /tmp/Intel/
cd APPS/BootUtil/Linux_x64
chmod a+x bootutili64e
./bootutili64e -h # help
./bootutili64e # list out the current settings for NIC
./bootutili64e -bootenable=pxe -all
shutdown -r now
# now go with PXE
# Check result of the flash utility (watch out for PXE Enabled on 40GbE interface)

#./bootutili64e
```

```
Intel(R) Ethernet Flash Firmware Utility
BootUtil version 1.6.20.1
Copyright (C) 2003-2016 Intel Corporation
```

Type BootUtil -? for help

Port	Network Address	Location	Series	WOL	Flash Firmware	Version
1	006BF10829A8	18:00.0	Gigabit	YES	UEFI,CLP,PXE Enabled,iSCSI	1.5.53
2	006BF10829A8	18:00.1	Gigabit	YES	UEFI,CLP,PXE Enabled,iSCSI	1.5.53
3	3CFDFEA471F0	10:00.0	40GbE	N/A	UEFI,CLP,PXE Enabled,iSCSI	1.0.31
4	3CFDFEA471F1	10:00.1	40GbE	N/A	UEFI,CLP,PXE Enabled,iSCSI	1.0.31
5	3CFDFEA471F2	10:00.2	40GbE	N/A	UEFI,CLP,PXE,iSCSI	-----
6	3CFDFEA471F3	10:00.3	40GbE	N/A	UEFI,CLP,PXE,iSCSI	-----
7	3CFDFEA47130	14:00.0	40GbE	N/A	UEFI,CLP,PXE Enabled,iSCSI	1.0.31
8	3CFDFEA47131	14:00.1	40GbE	N/A	UEFI,CLP,PXE Enabled,iSCSI	1.0.31
9	3CFDFEA47132	14:00.2	40GbE	N/A	UEFI,CLP,PXE,iSCSI	-----
10	3CFDFEA47133	14:00.3	40GbE	N/A	UEFI,CLP,PXE,iSCSI	-----

Note If the message 'FLASH Not Present' is displayed on execution of “./bootutil64e”, on some UCS-M5 servers you need to perform a server reboot. To perform the reboot run the shutdown command- **shutdown –r now**. After the server reboot, execution of “./bootutil64e”, reveals that the PXE is enabled on the targeted Intel NIC port.

Setting Up the UCS B-Series Pod

After you install the RHEL OS on the management node, complete the following steps to configure a Cisco NFVI B-Series pod:

Step 1 Log in to Cisco UCS Manager, connect to the console of both fabrics and execute the following commands:

```
# connect local-mgmt
# erase config
All UCS configurations are erased and system starts to reboot. Are you sure? (yes/no): yes
Removing all the configuration. Please wait...
```

Step 2 Go through the management connection and clustering wizards to configure Fabric A and Fabric B:

Fabric Interconnect A

```
# connect local-mgmt
# erase config
Enter the configuration method. (console/gui) console
Enter the setup mode; setup newly or restore from backup. (setup/restore) ? setup
You have chosen to setup a new Fabric interconnect. Continue? (y/n): y
Enforce strong password? (y/n) [y]: n
Enter the password for "admin":
Confirm the password for "admin":
Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes
Enter the switch fabric (A/B) []: A
Enter the system name: skull-fabric
Physical Switch Mgmt0 IPv4 address : 10.30.119.58
Physical Switch Mgmt0 IPv4 netmask : 255.255.255.0
IPv4 address of the default gateway : 10.30.119.1
Cluster IPv4 address : 10.30.119.60
Configure the DNS Server IPv4 address? (yes/no) [n]: y
DNS IPv4 address : 172.29.74.154
Configure the default domain name? (yes/no) [n]: y
Default domain name : ctocllab.cisco.com
```

Join centralized management environment (UCS Central)? (yes/no) [n]: **n**

Following configurations are applied:

```
Switch Fabric=A
System Name=skull-fabric
Enforced Strong Password=no
Physical Switch Mgmt0 IP Address=10.30.119.58
Physical Switch Mgmt0 IP Netmask=255.255.255.0
Default Gateway=10.30.119.1
DNS Server=172.29.74.154
Domain Name=ctocllab.cisco.com
Cluster Enabled=yes
Cluster IP Address=10.30.119.60
NOTE: Cluster IP is configured only after both Fabric Interconnects are initialized
```



```
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait..
```

Fabric Interconnect B

```
Enter the configuration method. (console/gui) ? console
```

```
Installer has detected the presence of a peer Fabric interconnect. This Fabric interconnect is 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 IP Address: 10.30.119.58
Peer Fabric interconnect Mgmt0 IP Netmask: 255.255.255.0
Cluster IP address : 10.30.119.60
Physical Switch Mgmt0 IPv4 address : 10.30.119.59
Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes
Applying configuration. Please wait.
```

Step 3 Configure the NTP:

- a) In UCS Manager navigation area, click the **Admin** tab.
- b) In the Filter drop-down list, choose **Time Zone Management**.
- c) In the main window under Actions, click **Add NTP Server**.
- d) In the Add NTP Server dialog box, enter the NTP hostname or IP address, then click **OK**.

Step 4 Following instructions in [Cisco UCS Manager GUI Configuration Guide, Release 2.2](#), "Configuring Server Ports with the Internal Fabric Manager" section, configure the Fabric Interconnect A and Fabric Interconnect B uplinks to the Cisco NFVI top of rack (ToR) switches as **Uplink Ports**, **Server Ports**, and **Port Channels**.

Step 5 Configure the downlinks to the B-Series server chassis as **Server Ports**.

Step 6 Acknowledge all chassis.

Configuring the Out-of-Band Management Switch

Cisco VIM installer API and SSH bonded interface occurs on 1-GB Intel NICs that connect the Cisco NFVI management node and the Cisco Catalyst switch. Following is a sample configuration for creating a port channel on a Catalyst switch. Modify the configuration for your environment:

```
interface GigabitEthernet0/39
  channel-group 2 mode active
  speed 1000

interface GigabitEthernet0/40
  channel-group 2 mode active
  speed 1000

interface Port-channel2
  switchport access vlan 165
  switchport mode access
```

