

Deploying as a Physical Appliance

- Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance, on page 1
- Understanding and cabling physical appliances, on page 3
- Deploy Nexus Dashboard as a physical appliance, on page 48

Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance

Before you proceed with deploying the Nexus Dashboard cluster, you must:

- Review and complete the prerequisites described in Prerequisites and Guidelines.
- Review the *Cisco Nexus Dashboard Release Notes* for any information that can affect your deployment. See the Cisco Nexus Dashboard documentation landing page.
- Ensure you are using the following hardware and the servers are racked and connected as described in *Cisco Nexus Dashboard Hardware Setup Guide* specific to the model of server you have.

  The physical appliance form factor is supported only on these versions of the original Cisco Nexus Dashboard platform hardware:

  - `SE-NODE-G2` (UCS-C220-M5). The product ID of the 3-node cluster chassis is `SE-CL-L3`.
  - `ND-NODE-L4` (UCS-C225-M6). The product ID of the 3-node cluster chassis is `ND-CLUSTER-L4`.
  - `ND-NODE-G5S` (UCS-C225-M8). The product ID of the 3-node cluster chassis is `ND-CLUSTERG5S`.
  - `ND-NODE-G5L` (UCS-C225-M8). The product ID of the 3-node cluster chassis is `ND-CLUSTERG5L`.

Note This hardware only supports Cisco Nexus Dashboard software. If any other operating system is installed, the node can no longer be used as a Cisco Nexus Dashboard node.

- Ensure that you are running a supported version of Cisco Integrated Management Controller (CIMC).

  The minimum that is supported and recommended versions of CIMC are listed in the "Compatibility" section of the *Release Notes* for your Cisco Nexus Dashboard release.

- Ensure that you have configured an IP address for the server's CIMC.

  See Configure a Cisco Integrated Management Controller IP address, on page 2.

- Ensure that Serial over LAN (SoL) is enabled in CIMC.

  See Enable Serial over LAN in the Cisco Integrated Management Controller, on page 3.

  You might have a misconfiguration of SoL if the bootstrap fails at the `bootstrap peer nodes` point with this error:

  ```
  Waiting for firstboot prompt on NodeX
  ```

- Ensure that all nodes are running the same release version image.

- If your Cisco Nexus Dashboard hardware came with a different release image than the one you want to deploy, we recommend deploying the cluster with the existing image first and then upgrading it to the needed release.

  For example, if the hardware you received came with the release 3.2.1 image pre-installed, but you want to deploy release 4.3.1 instead, we recommend:

  1. First, bring up the release 3.2.1 cluster, as described in the deployment guide for that release.
  2. Then upgrade to release 4.3.1, as described in Upgrading a Nexus Dashboard 3.2.2m Cluster to This Release.



Note For brand new deployments, you can also choose to simply re-image the nodes with the latest version of the Cisco Nexus Dashboard (for example, if the hardware came with an image which does not support a direct upgrade to this release through the GUI workflow) before returning to this document for deploying the cluster. This process is described in the "Re-Imaging Nodes" section of the *Troubleshooting* article for this release.

- You must have at least a 1-node cluster.
- Extra secondary nodes can be added for horizontal scaling if required. For the maximum number of `secondary` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

Configure a Cisco Integrated Management Controller IP address

Follow these steps to configure a Cisco Integrated Management Controller (CIMC) IP address.

Procedure

Step 1 Power on the server.

After the hardware diagnostic is complete, you will be prompted with different options controlled by the function (Fn) keys.

Step 2 Press the **F8** key to enter the **Cisco IMC configuration Utility**.

Step 3 Follow these substeps.

a) Set **NIC mode** to `Dedicated`.

b) Choose between the **IPv4** and **IPv6** IP modes.

You can choose to enable or disable DHCP. If you disable DHCP, provide the static IP address, subnet, and gateway information.

c) Ensure that **NIC Redundancy** is set to `None`..

d) Press **F1** for more options such as hostname, DNS, default user passwords, port properties, and reset port profiles.

Step 4 Press **F10** to save the configuration and then restart the server.

Enable Serial over LAN in the Cisco Integrated Management Controller

Serial over LAN (SoL) is required for the `connect host` command, which you use to connect to a physical appliance node to provide basic configuration information. To use the SoL, you must first enable it on your Cisco Integrated Management Controller (CIMC).

Follow these steps to enable Serial over LAN in the Cisco Integrated Management Controller.

Procedure

Step 1 SSH into the node using the CIMC IP address and enter the sign-in credentials.

Step 2 Run these commands:

```
Server# scope sol
Server /sol # set enabled yes
Server /sol *# set baud-rate 115200
Server /sol *# commit
Server /sol *#
Server /sol # show

C220-WZP23150D4C# scope sol
C220-WZP23150D4C /sol # show

Enabled Baud Rate(bps)  Com Port SOL SSH Port
------- --------------- -------- -------------
yes     115200          com0     2400
```

Step 3 In the command output, verify that `com0` is the com port for SoL.

This enables the system to monitor the console using the `connect host` command from the CIMC CLI, which is necessary for the cluster bringup.

Understanding and cabling physical appliances

These sections provide overview and cabling information for these supported Nexus Dashboard physical appliances.

ND-NODE-G5L (UCS-C225-M8)

These sections provide overview and cabling information for the ND-NODE-G5L (UCS-C225-M8) physical appliance.

Understanding the ND-NODE-G5L (UCS-C225-M8)

This section provides overview information for the ND-NODE-G5L (UCS-C225-M8) physical appliance.

The ND-NODE-G5L physical appliance is a Nexus Dashboard appliance that is based on the UCS C225 M8 server. Even though the Nexus Dashboard ND-NODE-G5L is based on the UCS C225 M8 server, you cannot replace components within the Nexus Dashboard ND-NODE-G5L as you would with the base UCS C225 M8 server; instead, you will perform a Return Material Authorization (RMA) operation for the appliance if a component within the Nexus Dashboard ND-NODE-G5L is faulty.

Because the Nexus Dashboard ND-NODE-G5L is based on the UCS C225 M8 server, you can refer to the *Cisco UCS C225 M8 Server Installation and Service Guide* for important UCS C225 M8 server information that is also applicable for the Nexus Dashboard ND-NODE-G5L, such as the information provided in these chapters in that document:

- Overview
- Installing the Server
- Server Specifications
- Storage Controller Considerations

However, because you cannot replace components within the Nexus Dashboard ND-NODE-G5L as you would with the base UCS C225 M8 server, these chapters from that document are not applicable for the Nexus Dashboard ND-NODE-G5L and should be ignored:

- Maintaining the Server
- Recycling Server Components
- Installation For Cisco UCS Manager Integration
- GPU Installation

The following sections provide information that is applicable for the Nexus Dashboard ND-NODE-G5L.

Overview

Cisco Nexus Dashboard provides a common platform for deploying Cisco Data Center applications. These applications provide real time analytics, visibility and assurance for policy and infrastructure.

The Cisco Nexus Dashboard server is required for installing and hosting the Cisco Nexus Dashboard application.

The appliance is orderable in the following versions:

- **ND-NODE-G5L**: Single-node appliance
- **ND-CLUSTERG5L**: Three-node version that leverages the same configuration as ND-NODE-G5L but with three appliances included

Components

The ND-NODE-G5L appliance is configured with the following components:

- **ND-M8-MLB**: ND M8 RACK MLB
- **ND-C225-M8S**: UCS C225 M8 Rack w/oCPU mem drv 1U wSFF HDD/SSD backplane
- **CIMC-LATEST-D**: IMC SW (Recommended) latest release for C-Series Servers
- **CON-L1NCO-UCSC5M8S**: CX LEVEL 1 8X7XNCDOS UCS C225 M8 Rack w/o CPU mem drv 1U w
- **ND-RAIL-D**: Ball Bearing Rail Kit for C220 & C240 M7/M8 rack servers
- **ND-CPU-A9655P**: ND AMD 9655P 2.6GHz 400W 96C/384MB Cache DDR5 6000MT/s
- **ND-M2-I240GB-D**: ND 240GB M.2 Boot SATA Intel SSD
- **ND-M2-HWRAID-D**: ND Cisco Boot optimized M.2 Raid controller
- **ND-TPM2-002D-D**: ND TPM 2.0 FIPS 140-2 MSW2022 compliant AMD M8 servers
- **ND-RIS1A-225M8**: ND C225 M8 1U Riser 1A PCIe Gen4 x16 HH
- **ND-HD24TB10KJ4-D**: ND 2.4TB 12G SAS 10K RPM SFF HDD (4Kn)
- **ND-SD960GBM3XEPD**: ND 960GB 2.5in Enter Perf 6G SATA Micron G2 SSD (3X)
- **ND-SD32TKA3XEP-D**: 3.2TB 2.5in Enter Perf 24G SAS Kioxia PM7 SSD (3X)
- **ND-M-V5Q50GV2-D**: Cisco VIC 15427 4x 10/25/50G mLOM C-Series w/Secure Boot
- **ND-FBRS2-C225M8**: C225 M8 Riser2 HH Filler Blank
- **ND-FBRS-C220-D**: C220M7 HH Riser3 blank
- Power supplies:
  - 1200W AC Titanium Power Supply for C-series Rack Servers
  - 1050W -48V DC Power Supply for UCS Rack Server
  - 1050W -48V DC Power Supply for APIC servers (India)
- **ND-MRX64G2RE5**: ND 64GB DDR5-6400 RDIMM 2Rx4 (16Gb)
- **ND-RAID-M1L16**: ND 24G Tri-Mode M1 RAID Controller w/4GB FBWC 16Drv
- **ND-P-ID10GC-D**: Cisco-Intel X710T2LG 2x10 GbE RJ45 PCIe NIC
- **ND-HSLP-C225M8**: ND C225 M8 Heatsink
- **ND-BBLKD-M8**: ND C-Series M6 & M8 SFF drive blanking panel
- **ND-DDR5-BLK**: ND DDR5 DIMM Blanks
- **ND-PSU-BLK-D**: Power Supply Blanking Panel for M7 / M8 servers
- **ND-HPBKT-225M8**: C225 M8 Tri-Mode 24G SAS RAID Controller Bracket
- **CBL-SAS-C225M8**: C225M8 SAS Cable Mainboard to RAID card

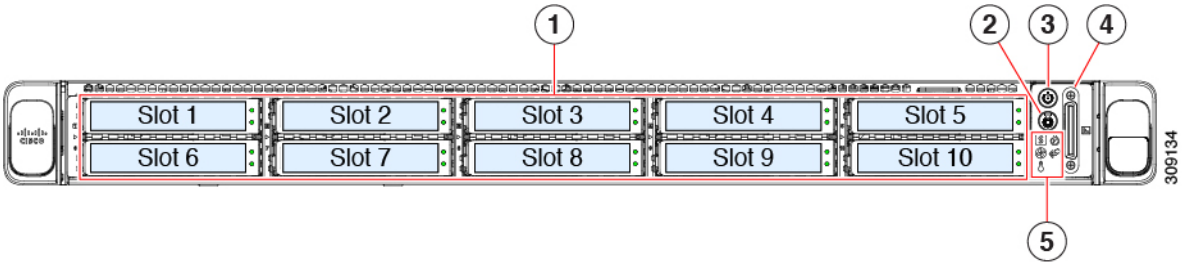
- **NO-POWER-CORD**: ECO friendly green option no power cable will be shipped
- **ND-DLOM-01-D**: Dedicated Mode BIOS setting for C-Series Servers
- **CNDL-DESELECT-D**: Conditional Deselect
- **OPTOUT-ENTL-SWAP**: License not needed: Entitlements updated in Smart Account

External Features

ND-NODE-G5L Front Panel Features

The following figure shows the front panel features of the small form-factor drive versions of the server.

Figure 1: ND-NODE-G5L Front Panel



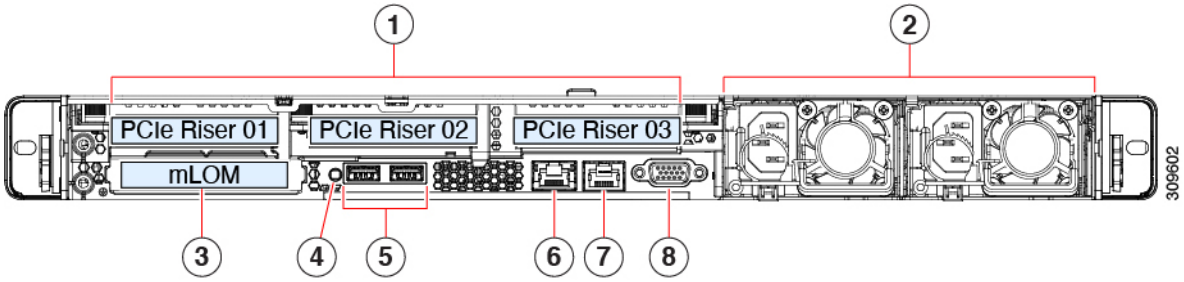
1	Drive bays 1 – 10 support SAS/SATA hard disk drives (HDDs) and solid state drives (SSDs).	**2**	Unit identification button/LED
3	Power button/power status LED	**4**	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)
5	System LED cluster: • Fan status LED • System Status LED • Power supply status LED • Network link activity LED • Temperature status LED		-

ND-NODE-G5L Rear Panel Features

The rear panel features can be different depending on the number and type of PCIe cards in the server.

The following figure shows the rear panel features of the server with three riser configuration.

Figure 2: ND-NODE-G5L Rear Panel Three-Riser Configuration



1	PCIe slots Following PCIe Riser combinations are available for 3 HH Riser cage configuration: • **Riser 1:** • Riser 1A (PCIe Gen4): Half-height, 3/4 length, x16.		
2	Power supply units (PSUs), two, which can be redundant when configured in 1+1 power mode.	**3**	Modular LAN-on-motherboard (mLOM) card bay (x16 PCIe lane)
4	System identification button/LED	**5**	USB 3.0 ports (two)
6	Dedicated 1 GB Ethernet management port	**7**	COM port (RJ-45 connector)
8	VGA video port (DB-15 connector)		

Status LEDs and Buttons

This section contains information for interpreting front, rear, and internal LED states.

Front-Panel LEDs

Figure 3: Front Panel LEDs

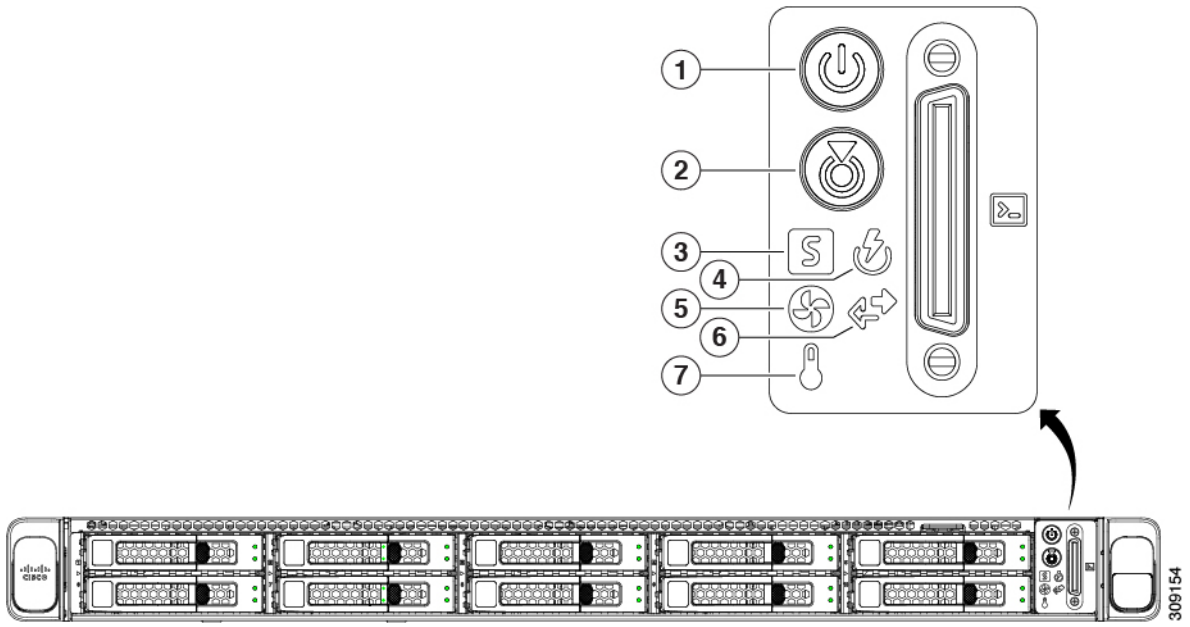


Table 1: Front Panel LEDs, Definition of States

	LED Name	States
1	Power button/LED ( )	• Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.
2	Unit identification ( )	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.

3	System health ()	• Green—The server is running in normal operating condition.
		• Green, blinking—The server is performing system initialization and memory check.
		• Amber, steady—The server is in a degraded operational state (minor fault). For example:
		  • Power supply redundancy is lost.
		  • CPUs are mismatched.
		  • At least one CPU is faulty.
		  • At least one DIMM is faulty.
		  • At least one drive in a RAID configuration failed.
		• Amber, 2 blinks—There is a major fault with the system board.
		• Amber, 3 blinks—There is a major fault with the memory DIMMs.
		• Amber, 4 blinks—There is a major fault with the CPUs.

4	Power supply status ()	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
5	Fan status ()	• Green—All fan modules are operating properly. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
6	Network link activity ()	• Off—The Ethernet LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.
7	Temperature status ()	• Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.

Rear-Panel LEDs

Figure 4: Rear Panel LEDs

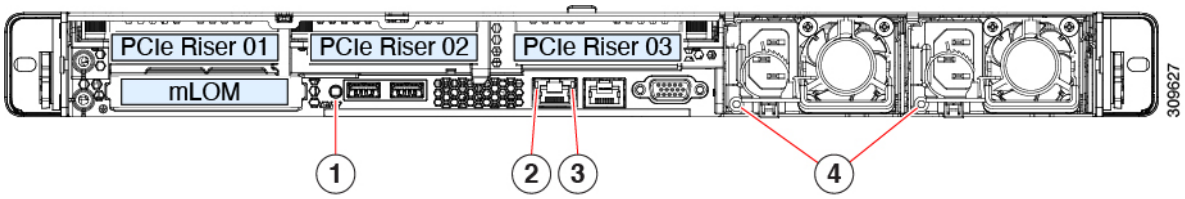


Table 2: Rear Panel LEDs, Definition of States

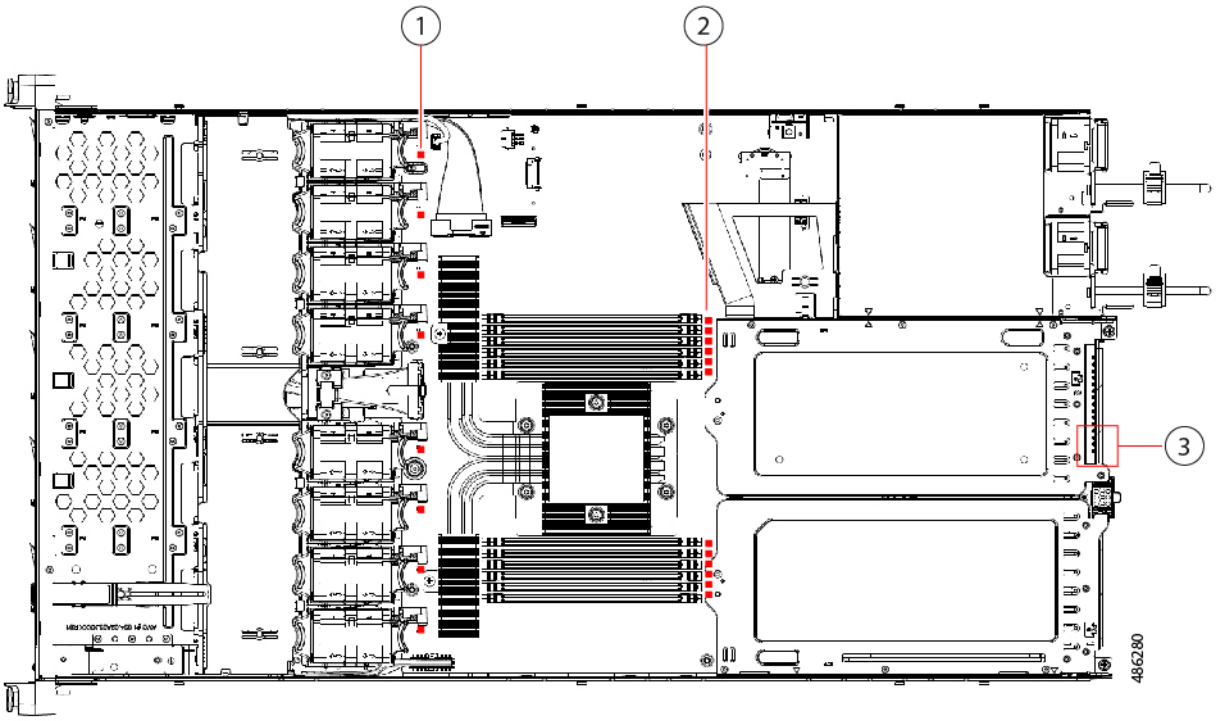
	LED Name	States
1	Rear unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
2	1-Gb Ethernet dedicated management link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.
3	1-Gb Ethernet dedicated management link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.

4	Power supply status (one LED each power supply unit)	**AC power supplies:** • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). **DC power supplies:** • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).

Internal Diagnostic LEDs

The server has internal fault LEDs for CPUs, DIMMs, and fan modules.

Figure 5: Internal Diagnostic LED Locations

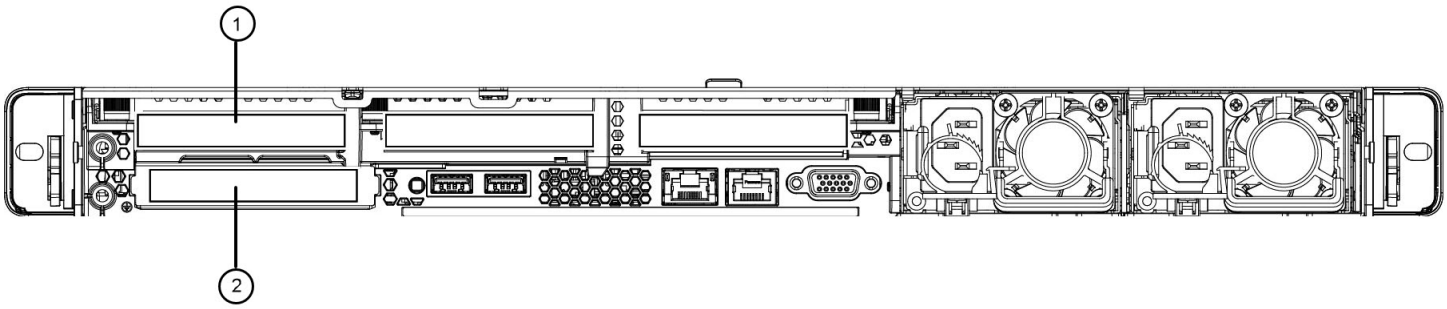


1	Fan module fault LEDs (one behind each fan connector on the motherboard) • Amber—Fan has a fault or is not fully seated. • Green—Fan is OK.	2	DIMM fault LEDs (one behind each DIMM socket on the motherboard) These LEDs operate only when the server is in standby power mode. • Amber—DIMM has a fault. • Off—DIMM is OK.
3	CPU fault LEDs (beside rear USB2 connector). These LEDs operate only when the server is in standby power mode. • Amber—CPU has a fault. • Off—CPU is OK.	-	

Physical node cabling

This figure shows the `ND-NODE-G5L` (UCS-C225-M8) physical server, where you will make these connections.

Figure 6: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-G5L (UCS-C225-M8)



1	**Management connections**: Through the two MGMT ports in the UCSC-O-ID10GC PCIe card installed in the PCIe Riser 01 location.
2	**Data connections**: Ports are numbered 1, 2, 3, and 4, from left-to-right in the UCSC-M-V5Q50GV2-D (Cisco UCS VIC 15427 Quad Port 10/25/50G CNA) Modular LAN-on-motherboard (mLOM). See the "Data network connections" information below for the supported port channel configurations.

Note The OCP card included on the ND-NODE-G5L servers supports a 1Gb copper connection for management only. All other network connections for Nexus Dashboard need to leverage the four port VIC card (callout 2 in the figure above). This VIC card supports 10/25/50Gbps, and the recommended SFP+ cable is the SFP-10G-AOC3M, but Cisco also offers 5-meter and 7-meter options as well. The VIC card requires a minimum of two connections per server for data network connectivity. These VIC connections can leverage any supported SFP, but Cisco recommends this connection for seamless deployments of Nexus Dashboard.

The physical nodes can be deployed with these guidelines:

- The ND-NODE-G5L comes with a UCSC-O-ID10GC PCIe card (shown in the above diagram), which you use for Nexus Dashboard management network connectivity.
- The ND-NODE-G5L also comes with a UCSC-M-V5Q50GV2-D (Cisco UCS VIC 15427 Quad Port 10/25/50G CNA) Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard data network.

Note Verify the Cisco UCS VIC card port mapping for the data interfaces using the information provided in "Data network connections" before cabling the node. The data interfaces must connect to individual host-facing switch ports; do not configure switch port channels or vPC for these connections.

- Verify that the port-channel feature is enabled on the VIC.
  - For the ND-NODE-L4 (UCS-C225-M6) and ND-NODE-G5S/ND-NODE-G5L (UCS-C225-M8), the port-channel feature is enabled by default at the factory. No additional configurations are necessary for those platforms, but you can use these procedures to verify that the port-channel feature is enabled correctly and enable it on those platforms, if necessary.

- For the `SE-NODE-G2` (UCS-C220-M5), you must manually enable the port-channel feature on the VIC.

1. Log into the CIMC.
2. Click the hamburger icon, then navigate to:

   Networking > **Adapter Card 1** > **General**
3. In the **Adapter Card Properties** area, locate the **Port Channel** field.

   If the **Port Channel** option is disabled (has no check in the check box), enable the feature (put a check in the check box) before beginning the Nexus Dashboard deployment process.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds, one for the data interfaces (bond0) and one for the management interfaces (bond1), running in Active-Standby mode.
- **Management network connections**:
  - You must use the `mgmt0` and `mgmt1` on the PCIe card.
  - All ports must have the same speed, either 1G or 10G.
- **Data network connections**:
  - On the `ND-NODE-G5L` server, you must use optical connections through the necessary port channel combinations in the UCSC-M-V5Q50GV2-D (Cisco UCS VIC 15427 Quad Port 10/25/50G CNA) PCIe card.

Note For 25/50 GB speed connections, you will need one of the following pairs of Forward Error Correction (FEC) configurations:

On the N9000	On the Nexus Dashboard applicance (via CIMC)
FEC AUTO	cl74
FC-FEC	cl74
FEC OFF	FEC OFF

  - All ports must have the same speed, either 10G, 25G, or 50G.
  - Attach the cables for the data network connections, where one cable connects to `fabric0` and the second cable connects to `fabric1`.

    `fabric0` and `fabric1` in the `ND-NODE-G5L` server corresponds to these ports:
    - Port-1 and Port-2 correspond to `fabric0`
    - Port-3 and Port-4 correspond to `fabric1`

    You can therefore have any one of these port channel combinations:

First data cable (`fabric0`)	Second data cable (`fabric1`)
Port-1	Port-3
Port-1	Port-4
Port-2	Port-3
Port-2	Port-4

See Figure 6: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-G5L (UCS-C225-M8), on page 13 for the port locations.

You can use both `fabric0` and `fabric1` for data network connectivity as Active-Standby.

Caution

- Do not configure a switch port channel or vPC for these connections. The supported pairings listed above identify the valid physical port mappings on the VIC card only.
- If you connect the two cables for the data network connections using different port channel combinations from those listed above, there will be MAC move notifications on the upstream switch and the ports will flap.

- If you connect the nodes to Cisco Catalyst switches, packets are tagged on those Catalyst switches with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

  If you are using Nexus Silicon One switch models, then you must configure the connected switchport interface as a trunk interface, and configure the VLAN that is specified in the Nexus Dashboard data network configuration.

What's next

Go to Deploy Nexus Dashboard as a physical appliance, on page 48 to deploy the Nexus Dashboard.

ND-NODE-G5S (UCS-C225-M8)

These sections provide overview and cabling information for the `ND-NODE-G5S` (UCS-C225-M8) physical appliance.

Understanding the ND-NODE-G5S (UCS-C225-M8)

This section provides overview information for the `ND-NODE-G5S` (UCS-C225-M8) physical appliance.

The `ND-NODE-G5S` physical appliance is a Nexus Dashboard appliance that is based on the UCS C225 M8 server. Even though the Nexus Dashboard `ND-NODE-G5S` is based on the UCS C225 M8 server, you cannot replace components within the Nexus Dashboard `ND-NODE-G5S` as you would with the base UCS C225 M8 server; instead, you will perform a Return Material Authorization (RMA) operation for the appliance if a component within the Nexus Dashboard `ND-NODE-G5S` is faulty.

Because the Nexus Dashboard ND-NODE-G5S is based on the UCS C225 M8 server, you can refer to the *Cisco UCS C225 M8 Server Installation and Service Guide* for important UCS C225 M8 server information that is also applicable for the Nexus Dashboard ND-NODE-G5S, such as the information provided in these chapters in that document:

- Overview
- Installing the Server
- Server Specifications
- Storage Controller Considerations

However, because you cannot replace components within the Nexus Dashboard ND-NODE-G5S as you would with the base UCS C225 M8 server, these chapters from that document are not applicable for the Nexus Dashboard ND-NODE-G5S and should be ignored:

- Maintaining the Server
- Recycling Server Components
- Installation For Cisco UCS Manager Integration
- GPU Installation

The following sections provide information that is applicable for the Nexus Dashboard ND-NODE-G5S.

Overview

Cisco Nexus Dashboard provides a common platform for deploying Cisco Data Center applications. These applications provide real time analytics, visibility and assurance for policy and infrastructure.

The Cisco Nexus Dashboard server is required for installing and hosting the Cisco Nexus Dashboard application.

The appliance is orderable in the following versions:

- **ND-NODE-G5S**: Single-node appliance
- **ND-CLUSTERG5S**: Three-node version that leverages the same configuration as ND-NODE-G5S but with three appliances included

Components

The ND-NODE-G5S appliance is configured with the following components:

- **CIMC-LATEST-D**: IMC SW (Recommended) latest release for C-Series Servers
- **ND-CPU-A9454P**: ND AMD 9454P 2.75GHz 290W 48C/256MB Cache DDR5 4800MT/s
- **ND-M2-240G-D**: ND 240GB M.2 SATA Micron G2 SSD
- **ND-M2-HWRAID-D**: ND Cisco Boot optimized M.2 Raid controller
- **ND-TPM2-002D-D**: ND TPM 2.0 FIPS 140-2 MSW2022 compliant AMD M8 servers
- **ND-RIS1A-225M8**: ND C225 M8 1U Riser 1A PCIe Gen4 x16 HH
- **ND-HD24TB10KJ4-D**: ND 2.4TB 12G SAS 10K RPM SFF HDD (4Kn)
- **ND-SD960GBM3XEPD**: ND 960GB 2.5in Enter Perf 6G SATA Micron G2 SSD (3X)

- **ND-RAIL-D:** Ball Bearing Rail Kit for C220 & C240 M7/M8 rack servers
- Power supplies:
  - 1200W AC Titanium Power Supply for C-series Rack Servers
  - 1050W -48V DC Power Supply for UCS Rack Server
  - 1050W -48V DC Power Supply for APIC servers (India)
- **ND-MRX32G1RE3**: ND 32GB DDR5-5600 RDIMM 1Rx4 (16Gb)
- **ND-RAID-M1L16**: ND 24G Tri-Mode M1 RAID Controller w/4GB FBWC 16Drv
- **ND-O-ID10GC-D**: Intel X710T2LOCPV3G1L 2x10GbE RJ45 OCP3.0 NIC
- **ND-OCP3-KIT-D**: C2XX OCP 3.0 Interposer W/Mech Assy
- **ND-P-V5Q50G-D**: Cisco VIC 15425 4x 10/25/50G PCIe C-Series w/Secure Boot

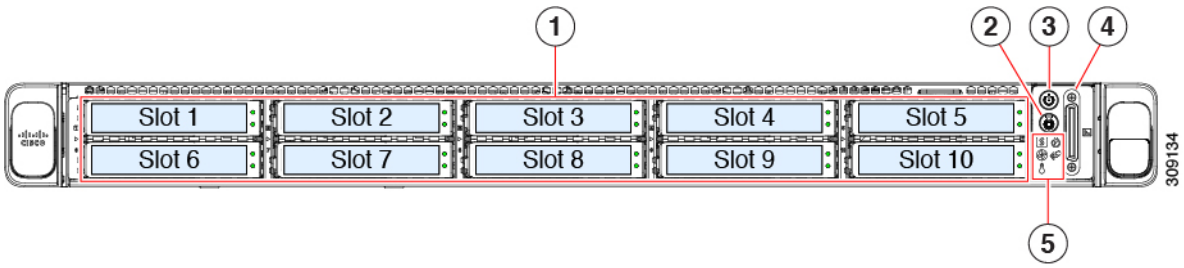
External Features

ND-NODE-G5S Front Panel Features

The following figure shows the front panel features of the small form-factor drive versions of the server.

For definitions of LED states, see Front-Panel LEDs, on page 8.

Figure 7: ND-NODE-G5S Front Panel



1	Drive bays 1 – 10 support SAS/SATA hard disk drives (HDDs) and solid state drives (SSDs).	**2**	Unit identification button/LED
3	Power button/power status LED	**4**	KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)
5	System LED cluster: • Fan status LED • System Status LED • Power supply status LED • Network link activity LED • Temperature status LED		-

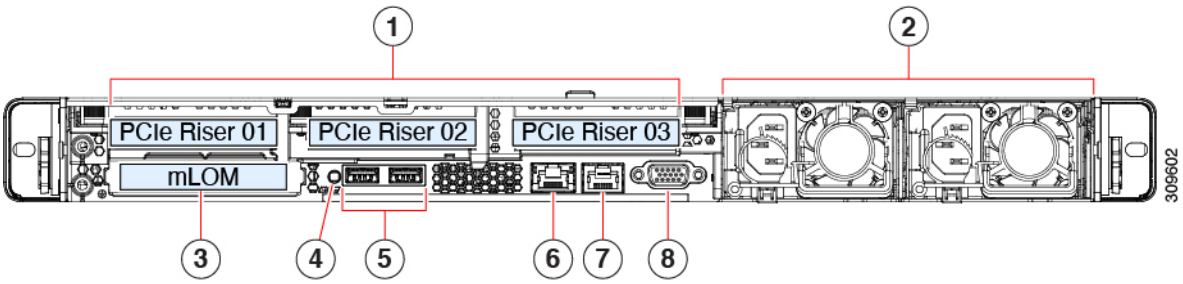
ND-NODE-G5S Rear Panel Features

The rear panel features can be different depending on the number and type of PCIe cards in the server.

The following figure shows the rear panel features of the server with three riser configuration.

For definitions of LED states, see Rear-Panel LEDs, on page 10.

Figure 8: ND-NODE-G5S Rear Panel Three-Riser Configuration



1	PCIe slots Following PCIe Riser combinations are available for 3 HH Riser cage configuration: • **Riser 1:** • Riser 1A (PCIe Gen4): Half-height, 3/4 length, x16, NCSI, Single Wide GPU.		
2	Power supply units (PSUs), two, which can be redundant when configured in 1+1 power mode.	**3**	Modular LAN-on-motherboard (mLOM) card bay (x16 PCIe lane)
4	System identification button/LED	**5**	USB 3.0 ports (two)
6	Dedicated 1 GB Ethernet management port	**7**	COM port (RJ-45 connector)
8	VGA video port (DB-15 connector)		

PCIe Slot Specifications

The following tables describe the specifications for the slots in three riser combination.

Table 3: PCIe Riser 1

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support
1A	Gen4 x16	x24 connector	¾ length	Half-height	Yes
1B	Gen5 x16	x24 connector	¾ length	Half-height	Yes

Status LEDs and Buttons

This section contains information for interpreting front, rear, and internal LED states.

Front-Panel LEDs

Figure 9: Front Panel LEDs

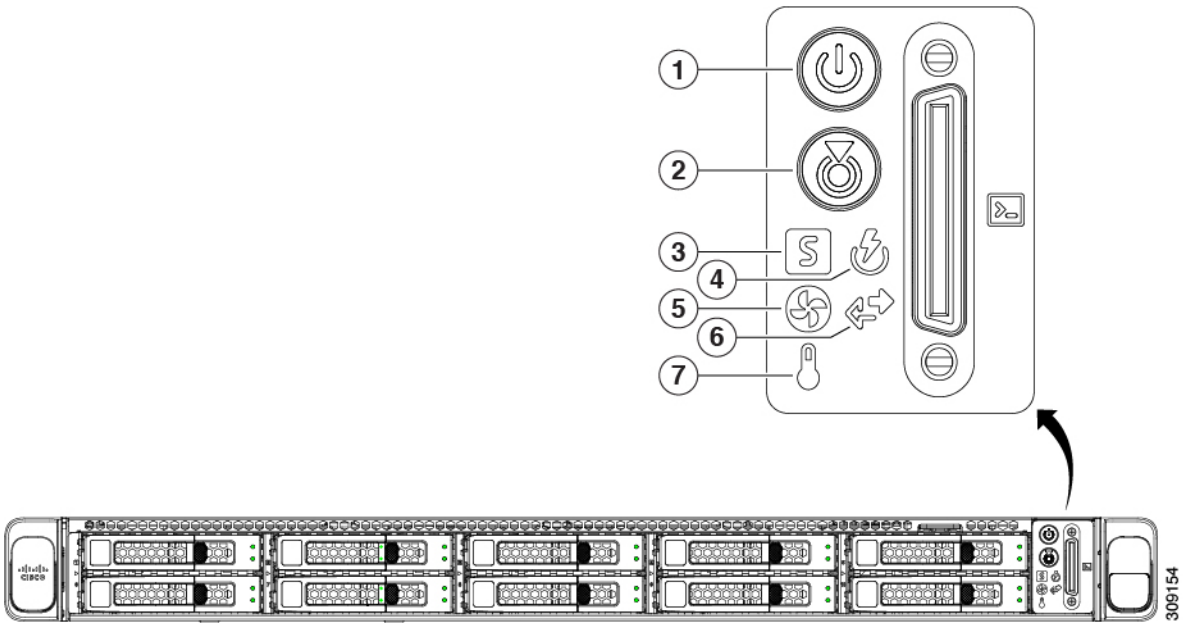


Table 4: Front Panel LEDs, Definition of States

	LED Name	States
1	Power button/LED ( )	• Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.
2	Unit identification ( )	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.

3	System health ()	• Green—The server is running in normal operating condition.
		• Green, blinking—The server is performing system initialization and memory check.
		• Amber, steady—The server is in a degraded operational state (minor fault). For example:
		  • Power supply redundancy is lost.
		  • CPUs are mismatched.
		  • At least one CPU is faulty.
		  • At least one DIMM is faulty.
		  • At least one drive in a RAID configuration failed.
		• Amber, 2 blinks—There is a major fault with the system board.
		• Amber, 3 blinks—There is a major fault with the memory DIMMs.
		• Amber, 4 blinks—There is a major fault with the CPUs.

4	Power supply status ()	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
5	Fan status ()	• Green—All fan modules are operating properly. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
6	Network link activity ()	• Off—The Ethernet LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.
7	Temperature status ()	• Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.

Rear-Panel LEDs

Figure 10: Rear Panel LEDs

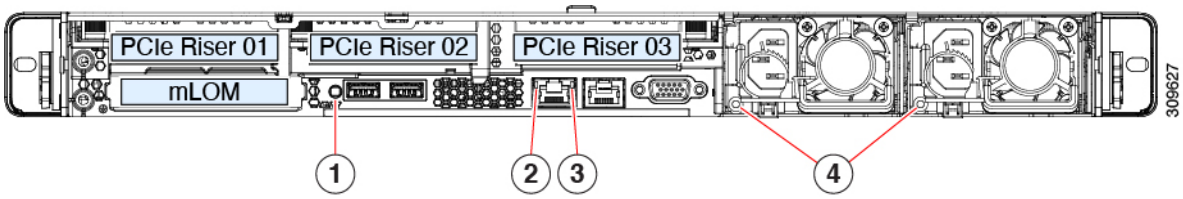


Table 5: Rear Panel LEDs, Definition of States

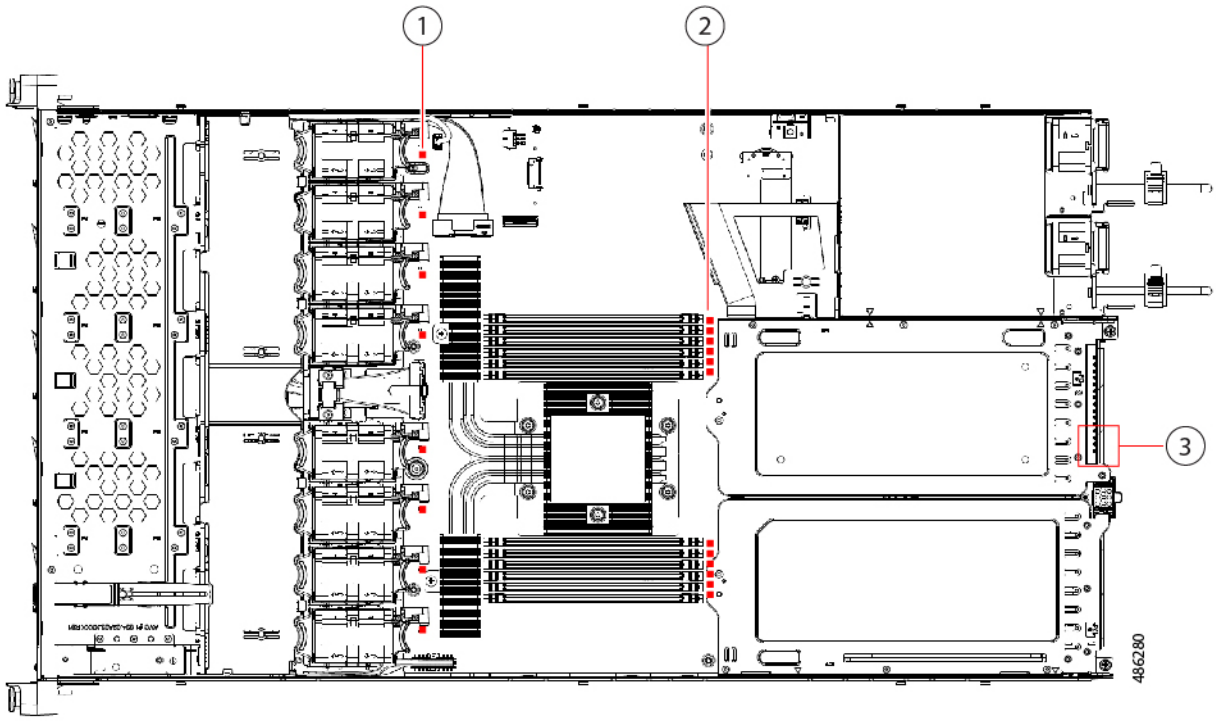
	LED Name	States
1	Rear unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
2	1-Gb Ethernet dedicated management link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.
3	1-Gb Ethernet dedicated management link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.

4	Power supply status (one LED each power supply unit)	**AC power supplies:** • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). **DC power supplies:** • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).

Internal Diagnostic LEDs

The server has internal fault LEDs for CPUs, DIMMs, and fan modules.

Figure 11: Internal Diagnostic LED Locations

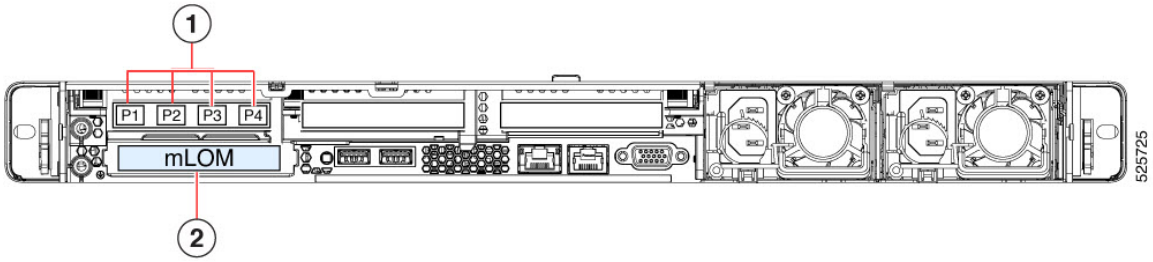


1	Fan module fault LEDs (one behind each fan connector on the motherboard) • Amber—Fan has a fault or is not fully seated. • Green—Fan is OK.	2	DIMM fault LEDs (one behind each DIMM socket on the motherboard) These LEDs operate only when the server is in standby power mode. • Amber—DIMM has a fault. • Off—DIMM is OK.
3	CPU fault LEDs (beside rear USB2 connector). These LEDs operate only when the server is in standby power mode. • Amber—CPU has a fault. • Off—CPU is OK.	-	

Physical node cabling

This figure shows the `ND-NODE-G5S` (UCS-C225-M8) physical server, where you will make these connections.

Figure 12: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-G5S (UCS-C225-M8)



1	**Data connections**: Ports are numbered 1, 2, 3, and 4, from left-to-right in the UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card installed in the PCIe Riser 01 location. See the "Data network connections" information below for the supported port channel configurations.
2	**Management connections**: Through the two MGMT ports in the Modular LAN-on-motherboard (mLOM).

Note The OCP card included on the ND-NODE-G5S servers supports a 1Gb copper connection for management only. All other network connections for Nexus Dashboard need to leverage the four port VIC card (callout 1 in the figure above). This VIC card supports 10/25/50Gbps, and the recommended SFP+ cable is the SFP-10G-AOC3M, but Cisco also offers 5-meter and 7-meter options as well. The VIC card requires a minimum of two connections per server for data network connectivity. These VIC connections can leverage any supported SFP, but Cisco recommends this connection for seamless deployments of Nexus Dashboard.

The physical nodes can be deployed with these guidelines:

- All servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The ND-NODE-G5S includes a UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card in the "PCIe-Riser-01" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity.

Note Verify the Cisco UCS VIC card port mapping for the data interfaces using the information provided in "Data network connections" before cabling the node. The data interfaces must connect to individual host-facing switch ports; do not configure switch port channels or vPC for these connections.

- Verify that the port-channel feature is enabled on the VIC.
  - For the ND-NODE-L4 (UCS-C225-M6) and ND-NODE-G5S/ND-NODE-G5L (UCS-C225-M8), the port-channel feature is enabled by default at the factory. No additional configurations are necessary for those platforms, but you can use these procedures to verify that the port-channel feature is enabled correctly and enable it on those platforms, if necessary.

- For the `SE-NODE-G2` (UCS-C220-M5), you must manually enable the port-channel feature on the VIC.

1. Log into the CIMC.
2. Click the hamburger icon, then navigate to:

   Networking > **Adapter Card 1** > **General**
3. In the **Adapter Card Properties** area, locate the **Port Channel** field.

   If the **Port Channel** option is disabled (has no check in the check box), enable the feature (put a check in the check box) before beginning the Nexus Dashboard deployment process.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds, one for the data interfaces (bond0) and one for the management interfaces (bond1), running in Active-Standby mode.
- **Management network connections**:
  - You must use the `mgmt0` and `mgmt1` on the mLOM card.
  - All ports must have the same speed, either 1G or 10G.
- **Data network connections**:
  - On the `ND-NODE-G5S` server, you must use optical connections through the necessary port channel combinations in the UCSC-P-V5Q50G-D (Cisco UCS VIC 15425 Quad Port 10/25/50G CNA) PCIE card.



Note For 25/50 GB speed connections, you will need one of the following pairs of Forward Error Correction (FEC) configurations:

On the N9000	On the Nexus Dashboard applicance (via CIMC)
FEC AUTO	cl74
FC-FEC	cl74
FEC OFF	FEC OFF

  - All ports must have the same speed, either 10G, 25G, or 50G.
  - Attach the cables for the data network connections, where one cable connects to `fabric0` and the second cable connects to `fabric1`.

    `fabric0` and `fabric1` in the `ND-NODE-G5S` server corresponds to these ports:
    - Port-1 and Port-2 correspond to `fabric0`
    - Port-3 and Port-4 correspond to `fabric1`

    You can therefore have any one of these port channel combinations:

First data cable (`fabric0`)	Second data cable (`fabric1`)
Port-1	Port-3
Port-1	Port-4
Port-2	Port-3
Port-2	Port-4

See Figure 12: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-G5S (UCS-C225-M8), on page 24 for the port locations.

You can use both `fabric0` and `fabric1` for data network connectivity as Active-Standby.

Caution

- Do not configure a switch port channel or vPC for these connections. The supported pairings listed above identify the valid physical port mappings on the VIC card only.
- If you connect the two cables for the data network connections using different port channel combinations from those listed above, there will be MAC move notifications on the upstream switch and the ports will flap.

- If you connect the nodes to Cisco Catalyst switches, packets are tagged on those Catalyst switches with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

  If you are using Nexus Silicon One switch models, then you must configure the connected switchport interface as a trunk interface, and configure the VLAN that is specified in the Nexus Dashboard data network configuration.

What's next

Go to Deploy Nexus Dashboard as a physical appliance, on page 48 to deploy the Nexus Dashboard.

ND-NODE-L4 (UCS-C225-M6)

These sections provide overview and cabling information for the `ND-NODE-L4` (UCS-C225-M6) physical appliance.

Understanding the ND-NODE-L4 (UCS-C225-M6)

This section provides overview information for the `ND-NODE-L4` (UCS-C225-M6) physical appliance.

The `ND-NODE-L4` physical appliance is a Nexus Dashboard appliance that is based on the UCS-C225-M6 server. Even though the Nexus Dashboard `ND-NODE-L4` is based on the UCS-C225-M6 server, you cannot replace components within the Nexus Dashboard `ND-NODE-L4` as you would with the base UCS-C225-M6 server; instead, you will perform a Return Material Authorization (RMA) operation for the appliance if a component within the Nexus Dashboard `ND-NODE-L4` is faulty.

Because the Nexus Dashboard ND-NODE-L4 is based on the UCS-C225-M6 server, you can refer to the *Cisco UCS C225 M6 Server Installation and Service Guide* for important UCS-C225-M6 server information that is also applicable for the Nexus Dashboard ND-NODE-L4, such as the information provided in these chapters in that document:

- Overview
- Installing the Server
- Server Specifications
- Storage Controller Considerations

However, because you cannot replace components within the Nexus Dashboard ND-NODE-L4 as you would with the base UCS-C225-M6 server, these chapters from that document are not applicable for the Nexus Dashboard ND-NODE-L4 and should be ignored:

- Maintaining the Server
- Installation For Cisco UCS Manager Integration
- GPU Installation

The following sections provide information that is applicable for the Nexus Dashboard ND-NODE-L4.

Overview

Cisco Nexus Dashboard provides a common platform for deploying Cisco Data Center applications. These applications provide real time analytics, visibility and assurance for policy and infrastructure.

The Cisco Nexus Dashboard server is required for installing and hosting the Cisco Nexus Dashboard application.

The server is orderable in the following version:

- ND-NODE-L4 — Small form-factor (SFF) drives, with 10-drive backplane. Supports up to 10 2.5-inch SAS/SATA drives. Drive bays 1 and 2 support NVMe SSDs.

Following PCIe Riser combinations are available:

- One half-height riser card in PCIe Riser 1
- Three half-height riser cards in PCIe Riser 1, 2, 3
- Two full-height riser cards Riser 1 and 3

- Riser 1—Supports Riser1. Supports single x16 PCIe supporting full height 3/4 length cards in 2 riser configuration (or) Half-height 3/4-length cards in 3 riser configuration and NC-SI from Pilot4.
- Riser 2—Supports Riser 1. Supports single x16 PCIe supporting only Half-height 3/4-length cards in 3-riser configuration.
- Riser 3—Supports Riser 3A, 3B. PCIe slot 3 with the following options:
  - Riser3A Supports single x16 PCIe supporting half height 3/4 length cards in 3 riser configuration and NC-SI.
  - Riser3B Supports single x16 PCIe supporting full height 3/4-length cards in 2 riser configuration and NC-SI.

- 2 10GBase-T Ethernet LAN over Motherboard (LOM) ports for network connectivity, plus one 1 Gigabit Ethernet dedicated management port
- One mLOM/VIC card provides 10G/25G/40G/50G/100G connectivity. Supported cards are:
  - Cisco VIC 1455 VIC PCIE – Quad Port 10/25G SFP28 (UCSC-PCIE-C25Q-04)

External Features

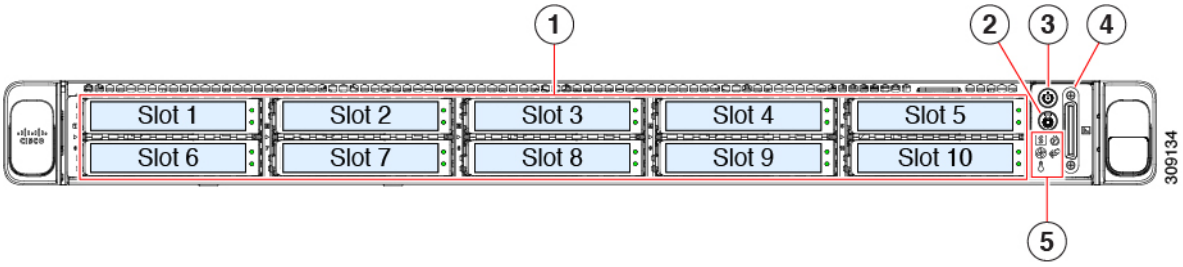
This topic shows the external features of the server versions.

Cisco ND-NODE-L4 (SFF Drives) Front Panel Features

The following figure shows the front panel features of the small form-factor drive versions of the server.

For definitions of LED states, see Front-Panel LEDs, on page 30.

Figure 13: ND-NODE-L4 (SFF Drives) Front Panel



1 UCSC-C225-M6S Version—Drive bays 1 – 10 support SAS/SATA hard disk drives (HDDs) and solid state drives (SSDs). As an option, drive bays 1-4 can contain up to 4 NVMe drives in any number up to 4. Drive bays 5 through 10 support only SAS/SATA HDDs or SSDs. UCSC-C225-M6N Version—Drive bays 1—10 supports 2.5-inch NVMe-only SSDs.	2 Unit identification button/LED
3 Power button/power status LED	4 KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB 2.0 connectors)
5 System LED cluster: • Fan status LED • System Status LED • Power supply status LED • Network link activity LED • Temperature status LED	-

Cisco ND-NODE-L4 Rear Panel Features

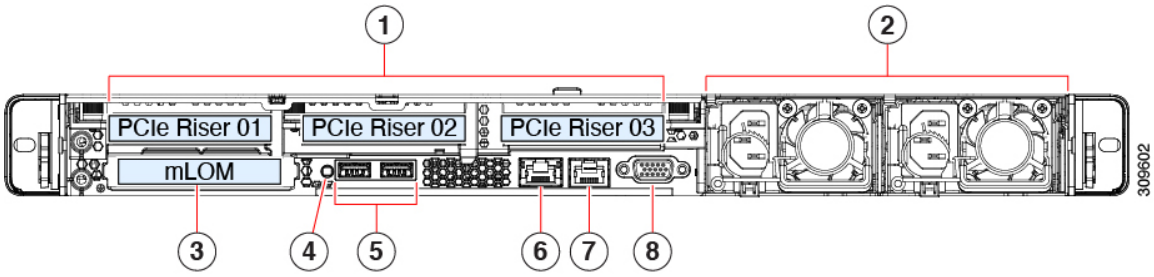
The rear panel features can be different depending on the number and type of PCIe cards in the server.

By default, single CPU servers come with only one half-height riser 1 installed, and dual CPU servers support all three half-height risers.

The following figure shows the rear panel features of the server with three riser configuration.

For definitions of LED states, see Rear-Panel LEDs, on page 32.

Figure 14: Cisco ND-NODE-L4 Rear Panel Three Riser Configuration



The following figure shows the rear panel features of the server with three riser configuration.

1 PCIe slots Following PCIe Riser combinations are available: • One half-height riser card in PCIe Riser 1	
2 Power supply units (PSUs), two which can be redundant when configured in 1+1 power mode.	3 Modular LAN-on-motherboard (mLOM) card bay (x16 PCIe lane)
4 System identification button/LED	5 USB 3.0 ports (two)
6 Dedicated 1 GB Ethernet management port	7 COM port (RJ-45 connector)
8 VGA video port (DB-15 connector)	

PCIe Slot Specifications

The following tables describe the specifications for the slots in three riser combination.

Table 6: PCIe Riser 1

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support
1	Gen-3 and 4 x16	x24 connector	¾ length	Half-height	Yes

Status LEDs and Buttons

This section contains information for interpreting front, rear, and internal LED states.

Front-Panel LEDs

Figure 15: Front Panel LEDs

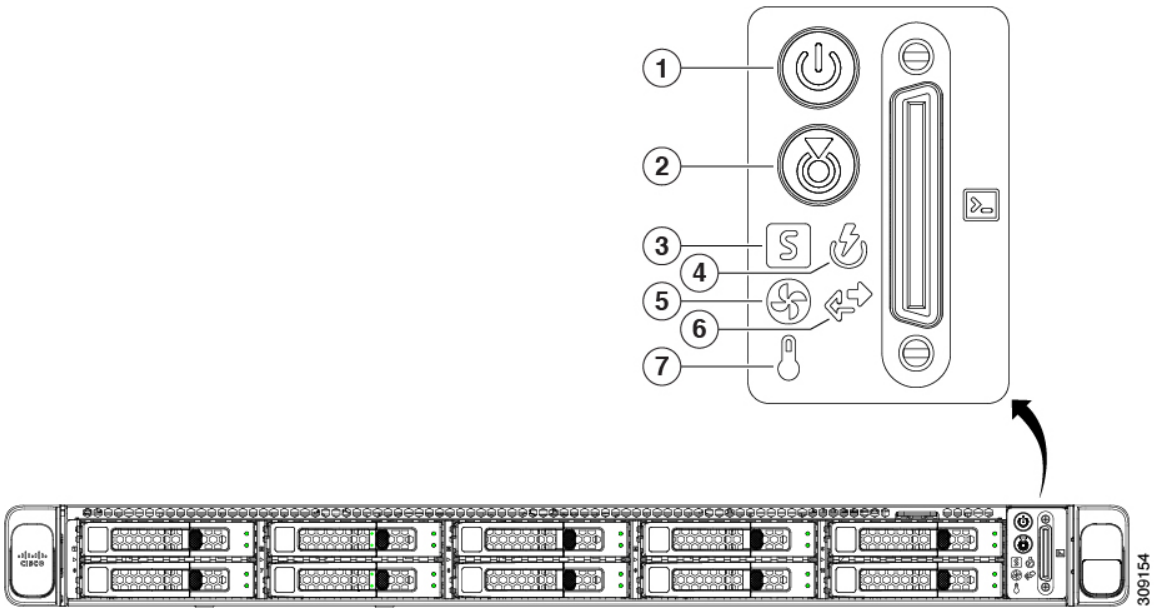


Table 7: Front Panel LEDs, Definition of States

LED Name	States
1 Power button/LED ( )	• Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.
2 Unit identification ( )	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.

3
System health ()
	• Green—The server is running in normal operating condition.
	• Green, blinking—The server is performing system initialization and memory check.
	• Amber, steady—The server is in a degraded operational state (minor fault). For example:
	  • Power supply redundancy is lost.
	  • CPUs are mismatched.
	  • At least one CPU is faulty.
	  • At least one DIMM is faulty.
	  • At least one drive in a RAID configuration failed.
	• Amber, 2 blinks—There is a major fault with the system board.
	• Amber, 3 blinks—There is a major fault with the memory DIMMs.
	• Amber, 4 blinks—There is a major fault with the CPUs.

4 Power supply status ()	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
5 Fan status ()	• Green—All fan modules are operating properly. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
6 Network link activity ()	• Off—The Ethernet LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.
7 Temperature status ()	• Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.

Rear-Panel LEDs

Figure 16: Rear Panel LEDs

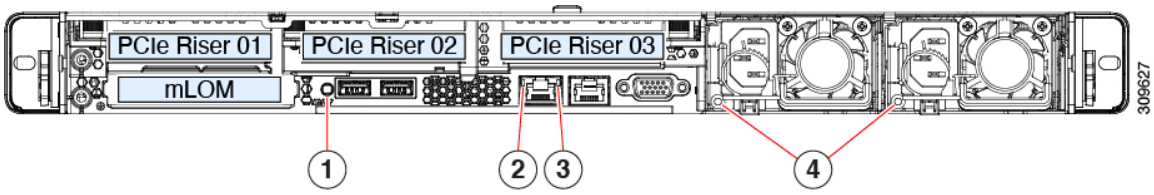


Table 8: Rear Panel LEDs, Definition of States

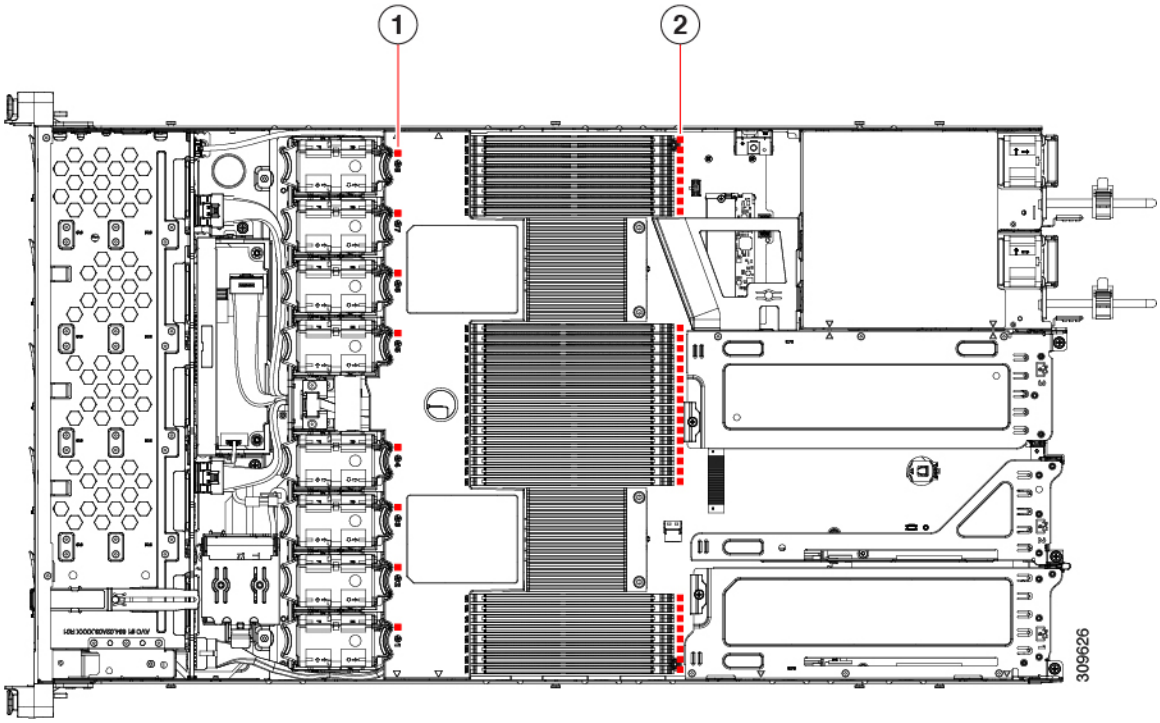
LED Name	States
1 Rear unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.
2 1-Gb Ethernet dedicated management link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.
3 1-Gb Ethernet dedicated management link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.

4 Power supply status (one LED each power supply unit)	**AC power supplies:** • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). **DC power supplies:** • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).

Internal Diagnostic LEDs

The server has internal fault LEDs for CPUs, DIMMs, and fan modules.

Figure 17: Internal Diagnostic LED Locations

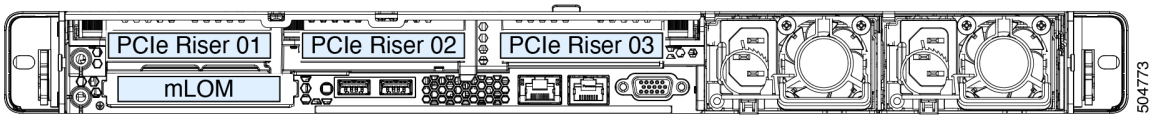


1 Fan module fault LEDs (one behind each fan connector on the motherboard) • Amber—Fan has a fault or is not fully seated. • Green—Fan is OK.	2 DIMM fault LEDs (one behind each DIMM socket on the motherboard) These LEDs operate only when the server is in standby power mode. • Amber—DIMM has a fault. • Off—DIMM is OK.
3 CPU fault LEDs (one behind each CPU socket on the motherboard). These LEDs operate only when the server is in standby power mode. • Amber—CPU has a fault. • Off—CPU is OK.	-

Physical node cabling

Physical nodes can be deployed in the `ND-NODE-L4` (UCS-C225-M6) physical server.

Figure 18: mLOM and PCIe riser 01 card used for node connectivity: ND-NODE-L4 (UCS-C225-M6)



The physical nodes can be deployed with these guidelines:

- All servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The ND-NODE-L4 server includes either a 2x10GbE NIC (APIC-P-ID10GC), or 2x25/10GbE SFP28 NIC (APIC-P-I8D25GF), or the VIC1455 card in the "PCIe-Riser-01" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity.

Note Verify the port mapping for the data interfaces using the information provided in "Data network connections" before cabling the node. The data interfaces must connect to individual host-facing switch ports; do not configure switch port channels or vPC for these connections.

- Verify that the port-channel feature is enabled on the VIC.
  - For the ND-NODE-L4 (UCS-C225-M6) and ND-NODE-G5S/ND-NODE-G5L (UCS-C225-M8), the port-channel feature is enabled by default at the factory. No additional configurations are necessary for those platforms, but you can use these procedures to verify that the port-channel feature is enabled correctly and enable it on those platforms, if necessary.
  - For the SE-NODE-G2 (UCS-C220-M5), you must manually enable the port-channel feature on the VIC.

  1. Log into the CIMC.
  2. Click the hamburger icon, then navigate to:

     Networking > **Adapter Card 1** > **General**
  3. In the **Adapter Card Properties** area, locate the **Port Channel** field.

     If the **Port Channel** option is disabled (has no check in the check box), enable the feature (put a check in the check box) before beginning the Nexus Dashboard deployment process.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds, one for the data interfaces (bond0) and one for the management interfaces (bond1), running in Active-Standby mode.
- **Management network connections**:
  - You must use the mgmt0 and mgmt1 on the mLOM card.
  - All ports must have the same speed, either 1G or 10G.
- **Data network connections**:

- On the ND-NODE-L4 server, you can use the 2x10GbE NIC (APIC-P-ID10GC), or 2x25/10GbE SFP28 NIC (APIC-P-I8D25GF), or the VIC1455 card.

Note If you connect using the 25G Intel NIC, you must disable the FEC setting on the switch port to match the setting on the NIC:

```
(config-if)# fec off
# show interface ethernet 1/34
Ethernet1/34 is up
admin state is up, Dedicated Interface
  [...]
  FEC mode is off
```

- All ports must have the same speed, either 10G, 25G, or 50G.
- Attach the cables for the data network connections, where one cable connects to fabric0 and the second cable connects to fabric1.

  fabric0 and fabric1 in the ND-NODE-L4 server corresponds to certain ports, depending on the type of card:

  - For the 2-port Intel cards:
    - Port-1 corresponds to fabric0
    - Port-2 corresponds to fabric1
  - For the 4-port cards:
    - Port-1 and Port-2 correspond to fabric0
    - Port-3 and Port-4 correspond to fabric1

  You can therefore have any one of these port channel combinations:

First data cable (fabric0)	Second data cable (fabric1)
Port-1	Port-3
Port-1	Port-4
Port-2	Port-3
Port-2	Port-4

  When using a 4-port card on the ND-NODE-L4 server, the order from left to right is Port-4, Port-3, Port-2, Port-1.

You can use both fabric0 and fabric1 for data network connectivity as Active-Standby.

Caution

- Do not configure a switch port channel or vPC for these connections. The supported pairings listed above identify the valid physical port mappings on the VIC card only.
- If you connect the two cables for the data network connections using different port channel combinations from those listed above, there will be MAC move notifications on the upstream switch and the ports will flap.

- If you connect the nodes to Cisco Catalyst switches, packets are tagged on those Catalyst switches with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

  If you are using Nexus Silicon One switch models, then you must configure the connected switchport interface as a trunk interface, and configure the VLAN that is specified in the Nexus Dashboard data network configuration.

What's next

Go to Deploy Nexus Dashboard as a physical appliance, on page 48 to deploy the Nexus Dashboard.

SE-NODE-G2 (UCS-C220-M5)

These sections provide overview and cabling information for the `SE-NODE-G2` (UCS-C220-M5) physical appliance.

Understanding the SE-NODE-G2 (UCS-C220-M5)

This section provides overview information for the `SE-NODE-G2` (UCS-C220-M5) physical appliance.

The `SE-NODE-G2` physical appliance is a Nexus Dashboard appliance that is based on the UCS C220 M5 server. Even though the Nexus Dashboard `SE-NODE-G2` is based on the UCS C220 M5 server, you cannot replace components within the Nexus Dashboard `SE-NODE-G2` as you would with the base UCS C220 M5 server; instead, you will perform a Return Material Authorization (RMA) operation for the appliance if a component within the Nexus Dashboard `SE-NODE-G2` is faulty.

Because the Nexus Dashboard `SE-NODE-G2` is based on the UCS C220 M5 server, you can refer to the *Cisco UCS C220 M5 Server Installation and Service Guide* for important UCS C220 M5 server information that is also applicable for the Nexus Dashboard `SE-NODE-G2`, such as the information provided in these chapters in that document:

- Overview
- Installing the Server
- Server Specifications
- Storage Controller Considerations

However, because you cannot replace components within the Nexus Dashboard `SE-NODE-G2` as you would with the base UCS C220 M5 server, these chapters from that document are not applicable for the Nexus Dashboard `SE-NODE-G2` and should be ignored:

- Maintaining the Server
- GPU Installation
- Installation For Cisco UCS Manager Integration

The following sections provide information that is applicable for the Nexus Dashboard `SE-NODE-G2`.

Overview

Cisco Nexus Dashboard provides a common platform for deploying Cisco Data Center applications. These applications provide real time analytics, visibility and assurance for policy and infrastructure.

The Cisco Nexus Dashboard server is required for installing and hosting the Cisco Nexus Dashboard application.

The server is orderable in the following version:

- SE-CL-L3 — Small form-factor (SFF) drives, with 10-drive backplane. Supports up to 10 2.5-inch SAS/SATA drives. Drive bays 1 and 2 support NVMe SSDs.

External Features

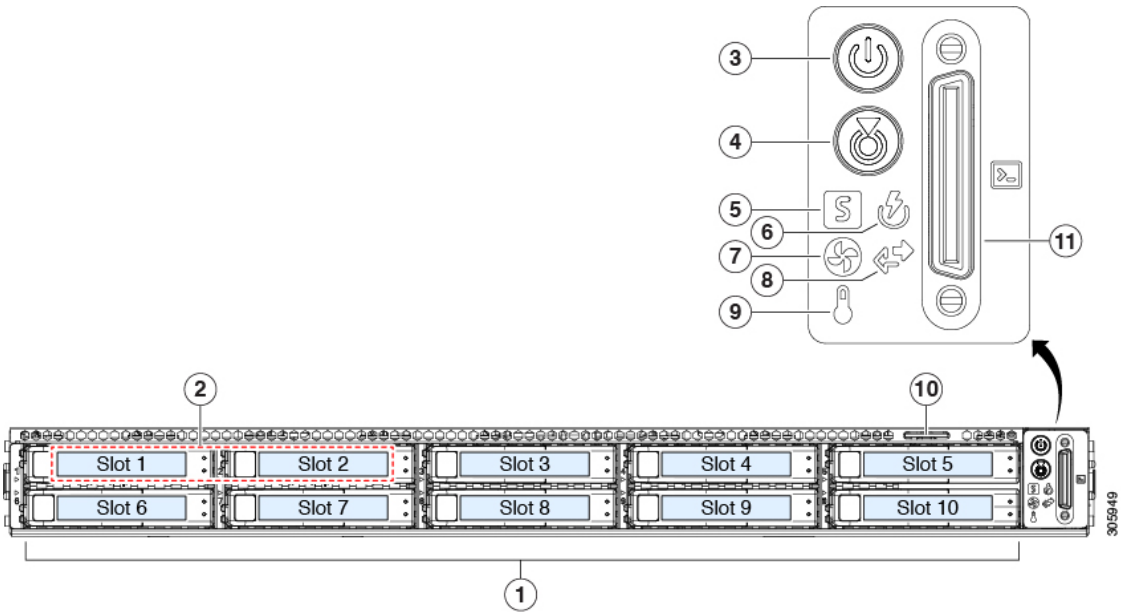
This topic shows the external features of the server versions.

Cisco SE-CL-L3 (SFF Drives) Front Panel Features

The following figure shows the front panel features of the small form-factor drive versions of the server.

For definitions of LED states, see Front-Panel LEDs, on page 41.

Figure 19: Cisco SE-CL-L3 (SFF Drives) Front Panel



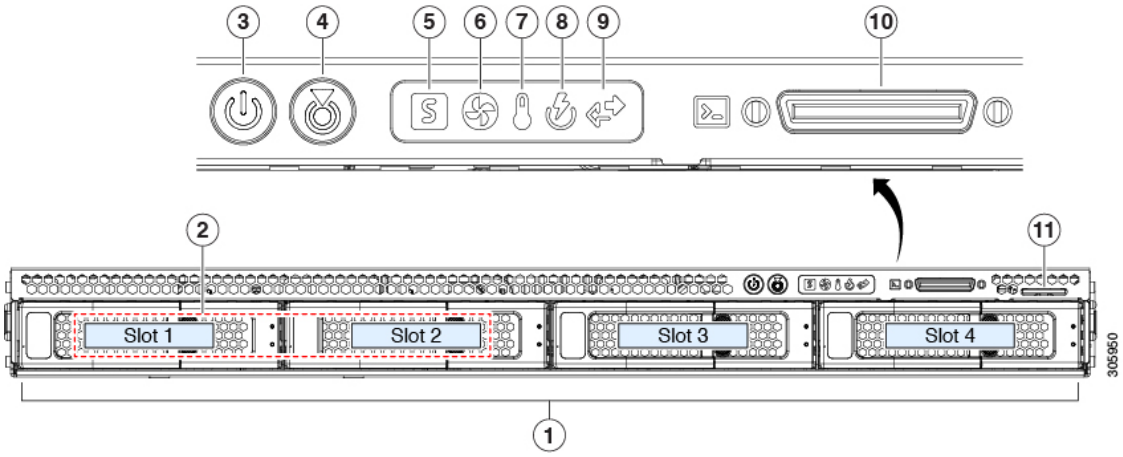
1 Drive bays 1 – 10 support SAS/SATA hard disk drives (HDDs) and solid state drives (SSDs)	7 Fan status LED
2 • SE-CL-L3 : Drive bays 1 and 2 support NVMe PCIe SSDs.	8 Network link activity LED
3 Power button/power status LED	9 Temperature status LED
4 Unit identification button/LED	10 Pull-out asset tag
5 System status LED	11 KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB connectors)
6 Power supply status LED	-

SE-CL-L3 (LFF Drives) Front Panel Features

The following figure shows the front panel features of the large form-factor drive version of the server.

For definitions of LED states, see Front-Panel LEDs, on page 41.

Figure 20: SE-CL-L3 (LFF Drives) Front Panel



1 Drive bays 1 – 4 support SAS/SATA HDDs and SSDs	7 Temperature status LED
2 Drive bays 1 and 2 support NVMe PCIe SSDs. A size-converter drive sled is required if 2.5-inch SSDs are used.	8 Power supply status LED
3 Power button/power status LED	9 Network link activity LED
4 Unit identification button/LED	10 KVM connector (used with KVM cable that provides one DB-15 VGA, one DB-9 serial, and two USB connectors)
5 System health LED	11 Pull-out asset tag

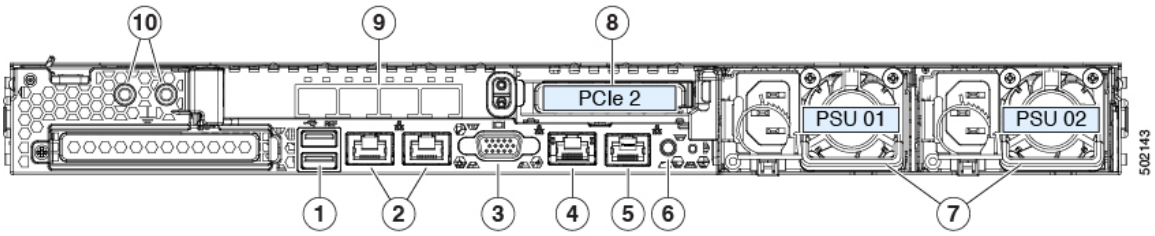
6 Fan status LED	-

SE-CL-L3 Rear Panel Features

The rear panel features are the same for all versions of the server.

For definitions of LED states, see Rear-Panel LEDs, on page 44.

Figure 21: SE-CL-L3 Rear Panel

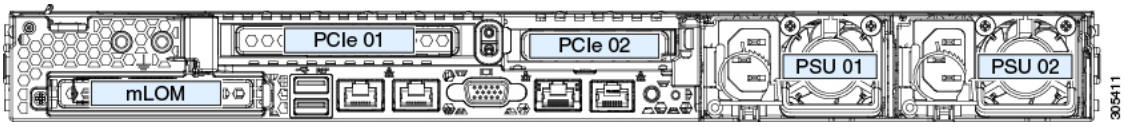


1 USB 3.0 ports (two)	6 Rear unit identification button/LED
2 Dual 1-Gb/10-Gb Ethernet ports (LAN1 and LAN2) The dual LAN ports can support 1 Gbps and 10 Gbps, depending on the link partner capability. These correspond to `eth1-1` (`eth0`) and `eth1-2` (`eth1`) respectively.	7 Power supplies (two, redundant as 1+1)
3 VGA video port (DB-15 connector)	8 PCIe riser 1/slot 1 (x16 lane) Includes PCIe cable connectors for front-loading NVMe SSDs (x8 lane)
4 1-Gb Ethernet dedicated management port	9 Quad 10-Gb/25-Gb ports. These correspond to `eth 2-1` to `eth 2-4`. Only 2 interfaces out of the 4 are active at a time (`eth2-1/2-2` or `eth2-3/2-4`) in active/standby mode.
5 Serial port (RJ-45 connector)	10 Threaded holes for dual-hole grounding lug

PCIe Slot Specifications

The server contains two PCIe slots on one riser assembly for horizontal installation of PCIe cards. Both slots support the NCSI protocol and 12V standby power.

Figure 22: Rear Panel, Showing PCIe Slot Numbering



The following tables describe the specifications for the slots.

Table 9: PCIe Riser 1/Slot 1

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support
1	Gen-3 x16	x24 connector	¾ length	Full-height	Yes
Micro SD card slot	One socket for Micro SD card				

Table 10: PCIe Riser 2/Slot 2

Slot Number	Electrical Lane Width	Connector Length	Maximum Card Length	Card Height (Rear Panel Opening)	NCSI Support
2	Gen-3 x16	x24 connector	½ length	½ height	Yes
PCIe cable connector for front-panel NVMe SSDs	Gen-3 x8	Other end of cable connects to front drive backplane to support front-panel NVMe SSDs.			

Note Riser 2/Slot 2 is not available in single-CPU configurations.

Status LEDs and Buttons

This section contains information for interpreting front, rear, and internal LED states.

Front-Panel LEDs

Figure 23: Front Panel LEDs

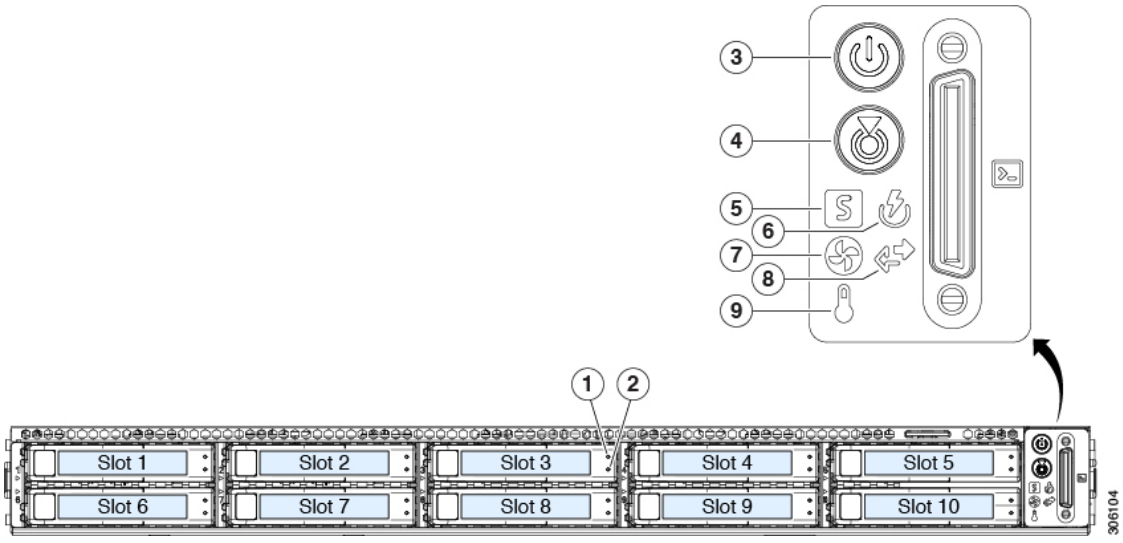


Table 11: Front Panel LEDs, Definition of States

LED Name	States
1 SAS/SATA drive fault Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	• Off—The hard drive is operating properly. • Amber—Drive fault detected. • Amber, blinking—The device is rebuilding. • Amber, blinking with one-second interval—Drive locate function activated in the software.
2 SAS/SATA drive activity LED S	• Off—There is no hard drive in the hard drive tray (no access, no fault). • Green—The hard drive is ready. • Green, blinking—The hard drive is reading or writing data.
1 NVMe SSD drive fault Note NVMe solid state drive (SSD) drive tray LEDs have different behavior than SAS/SATA drive trays.	• Off—The drive is not in use and can be safely removed. • Green—The drive is in use and functioning properly. • Green, blinking—the driver is initializing following insertion or the driver is unloading following an eject command. • Amber—The drive has failed. • Amber, blinking—A drive Locate command has been issued in the software.
2 NVMe SSD activity M	• Off—No drive activity. • Green, blinking—There is drive activity.
3 Power button/LED	• Off—There is no AC power to the server. • Amber—The server is in standby power mode. Power is supplied only to the Cisco IMC and some motherboard functions. • Green—The server is in main power mode. Power is supplied to all server components.
4 Unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.

5 System health	• Green—The server is running in normal operating condition.
	• Green, blinking—The server is performing system initialization and memory check.
	• Amber, steady—The server is in a degraded operational state (minor fault). For example:
	  • Power supply redundancy is lost.
	  • CPUs are mismatched.
	  • At least one CPU is faulty.
	  • At least one DIMM is faulty.
	  • At least one drive in a RAID configuration failed.
	• Amber, 2 blinks—There is a major fault with the system board.
	• Amber, 3 blinks—There is a major fault with the memory DIMMs.
	• Amber, 4 blinks—There is a major fault with the CPUs.

6 Power supply status	• Green—All power supplies are operating normally. • Amber, steady—One or more power supplies are in a degraded operational state. • Amber, blinking—One or more power supplies are in a critical fault state.
7 Fan status	• Green—All fan modules are operating properly. • Amber, blinking—One or more fan modules breached the non-recoverable threshold.
8 Network link activity	• Off—The Ethernet LOM port link is idle. • Green—One or more Ethernet LOM ports are link-active, but there is no activity. • Green, blinking—One or more Ethernet LOM ports are link-active, with activity.
9 Temperature status	• Green—The server is operating at normal temperature. • Amber, steady—One or more temperature sensors breached the critical threshold. • Amber, blinking—One or more temperature sensors breached the non-recoverable threshold.

Rear-Panel LEDs

Figure 24: Rear Panel LEDs

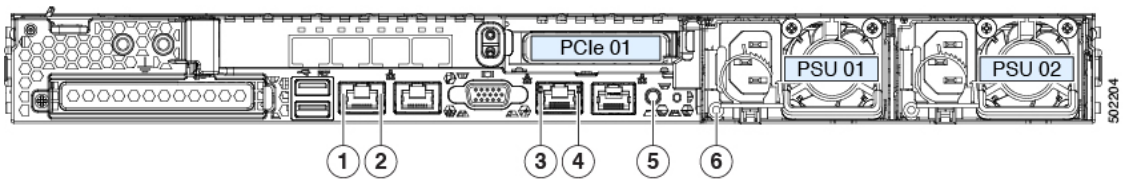


Table 12: Rear Panel LEDs, Definition of States

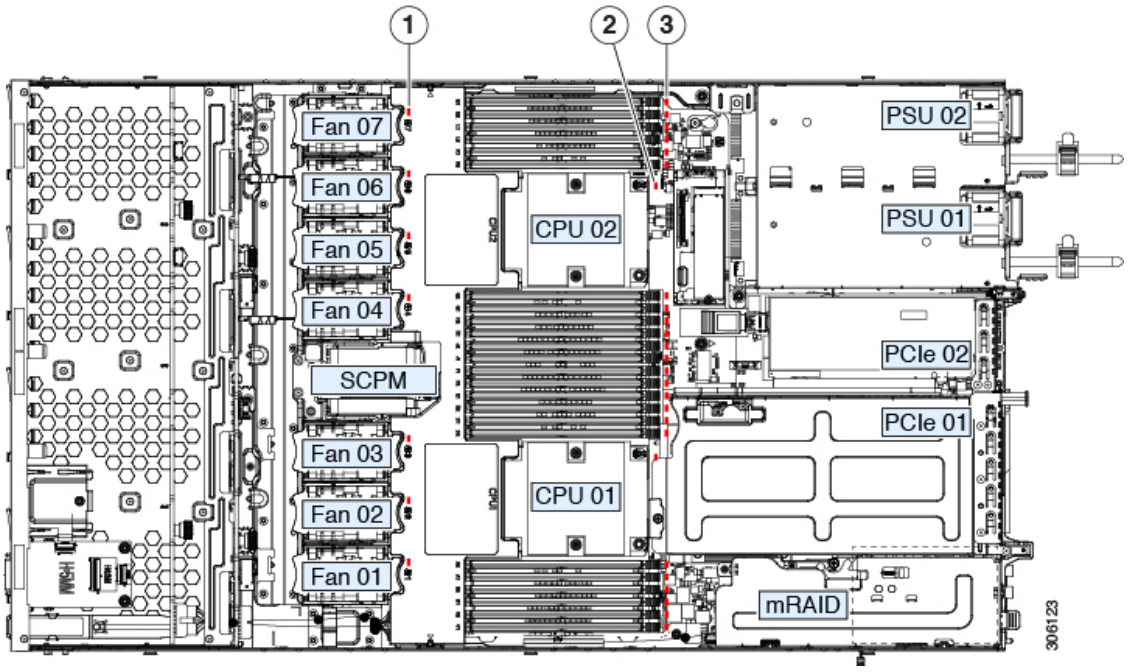
LED Name	States
1-Gb/10-Gb Ethernet link speed (on both LAN1 and LAN2)	• Amber—Link speed is 100 Mbps. • Amber—Link speed is 1 Gbps. • Green—Link speed is 10 Gbps.
2-Gb/10-Gb Ethernet link status (on both LAN1 and LAN2)	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
3-Gb Ethernet dedicated management link speed	• Off—Link speed is 10 Mbps. • Amber—Link speed is 100 Mbps. • Green—Link speed is 1 Gbps.
4-Gb Ethernet dedicated management link status	• Off—No link is present. • Green—Link is active. • Green, blinking—Traffic is present on the active link.
5Rear unit identification	• Off—The unit identification function is not in use. • Blue, blinking—The unit identification function is activated.

6 Power supply status (one LED each power supply unit)	**AC power supplies:** • Off—No AC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure). **DC power supplies:** • Off—No DC input (12 V main power off, 12 V standby power off). • Green, blinking—12 V main power off; 12 V standby power on. • Green, solid—12 V main power on; 12 V standby power on. • Amber, blinking—Warning threshold detected but 12 V main power on. • Amber, solid—Critical error detected; 12 V main power off (for example, over-current, over-voltage, or over-temperature failure).

Internal Diagnostic LEDs

The server has internal fault LEDs for CPUs, DIMMs, and fan modules.

Figure 25: Internal Diagnostic LED Locations

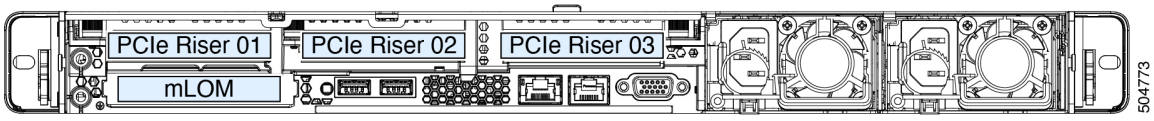


1 Fan module fault LEDs (one behind each fan connector on the motherboard) • Amber—Fan has a fault or is not fully seated. • Green—Fan is OK.	3 DIMM fault LEDs (one behind each DIMM socket on the motherboard) These LEDs operate only when the server is in standby power mode. • Amber—DIMM has a fault. • Off—DIMM is OK.
2 CPU fault LEDs (one behind each CPU socket on the motherboard). These LEDs operate only when the server is in standby power mode. • Amber—CPU has a fault. • Off—CPU is OK.	-

Physical node cabling

Physical nodes can be deployed in the `SE-NODE-G2` (UCS-C220-M5) physical server.

Figure 26: mLOM and PCIe riser 01 card used for node connectivity: `SE-NODE-G2` (UCS-C220-M5)



The physical nodes can be deployed with these guidelines:

- All servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The `SE-NODE-G2` server includes a 4-port VIC1455 card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity

Note Verify the Cisco UCS VIC card port mapping for the data interfaces using the information provided in "Data network connections" before cabling the node. The data interfaces must connect to individual host-facing switch ports; do not configure switch port channels or vPC for these connections.

- Verify that the port-channel feature is enabled on the VIC.
  - For the `ND-NODE-L4` (UCS-C225-M6) and `ND-NODE-G5S`/`ND-NODE-G5L` (UCS-C225-M8), the port-channel feature is enabled by default at the factory. No additional configurations are necessary for those platforms, but you can use these procedures to verify that the port-channel feature is enabled correctly and enable it on those platforms, if necessary.
  - For the `SE-NODE-G2` (UCS-C220-M5), you must manually enable the port-channel feature on the VIC.

1. Log into the CIMC.
2. Click the hamburger icon, then navigate to:

   Networking > **Adapter Card 1** > **General**
3. In the **Adapter Card Properties** area, locate the **Port Channel** field.

   If the **Port Channel** option is disabled (has no check in the check box), enable the feature (put a check in the check box) before beginning the Nexus Dashboard deployment process.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds, one for the data interfaces (bond0) and one for the management interfaces (bond1), running in Active-Standby mode.
- **Management network connections**:
  - You must use the `mgmt0` and `mgmt1` on the mLOM card.
  - All ports must have the same speed, either 1G or 10G.
- **Data network connections**:
  - On the `SE-NODE-G2` server, you must use the VIC1455 card.
  - All ports must have the same speed, either 10G, 25G, or 50G.
  - Attach the cables for the data network connections, where one cable connects to `fabric0` and the second cable connects to `fabric1`.

    `fabric0` and `fabric1` in `SE-NODE-G2` corresponds to these ports:

- Port-1 and Port-2 correspond to `fabric0`
- Port-3 and Port-4 correspond to `fabric1`

You can therefore have any one of these port channel combinations:

First data cable (`fabric0`)	Second data cable (`fabric1`)
Port-1	Port-3
Port-1	Port-4
Port-2	Port-3
Port-2	Port-4

When using a 4-port card on the `SE-NODE-G2` server, the order from left to right is Port-1, Port-2, Port-3, Port-4.

You can use both `fabric0` and `fabric1` for data network connectivity as Active-Standby.

Caution

- Do not configure a switch port channel or vPC for these connections. The supported pairings listed above identify the valid physical port mappings on the VIC card only.
- If you connect the two cables for the data network connections using different port channel combinations from those listed above, there will be MAC move notifications on the upstream switch and the ports will flap.

- If you connect the nodes to Cisco Catalyst switches, packets are tagged on those Catalyst switches with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

  If you are using Nexus Silicon One switch models, then you must configure the connected switchport interface as a trunk interface, and configure the VLAN that is specified in the Nexus Dashboard data network configuration.

What's next

Go to Deploy Nexus Dashboard as a physical appliance, on page 48 to deploy the Nexus Dashboard.

Deploy Nexus Dashboard as a physical appliance

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. Follow these steps to deploy Nexus Dashboard as a physical appliance.

Before you begin

Complete the requirements and guidelines described in Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance, on page 1.

Procedure

Step 1 Configure the first node's basic information.

You must configure only a single ("first") node as described in this step. Other nodes will be configured during the GUI-based cluster deployment process described in the following steps and will accept settings from the first `primary` node. The other two `primary` nodes do not require any additional configuration besides ensuring that their CIMC IP addresses are reachable from the first `primary` node and login credentials are set, as well as network connectivity between the nodes is established on the data network.

a) SSH into the node using CIMC management IP and use the `connect host` command to connect to the node's console.

```
C220-WZP23150D4C# connect host
CISCO Serial Over LAN:
Press Ctrl+x to Exit the session
```

After connecting to the host, press **Enter** to continue.

b) After you see the Nexus Dashboard setup utility prompt, press **Enter**.

```
Starting Nexus Dashboard setup utility
Welcome to Nexus Dashboard 4.2.1
Press Enter to manually bootstrap your first master node...
```

c) Enter and confirm the `admin` password

This password will be used for the `rescue-user` CLI login as well as the initial GUI password.

```
Admin Password:
Reenter Admin Password:
```

d) Enter the management network information.

```
Management Network:
  IP Address/Mask: 192.168.9.172/24
  Gateway: 192.168.9.1
```

Note If you want to configure IPv6-only mode, enter the IPv6 in the above example instead.

e) Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, enter the capital letter `N` to proceed. If you want to change any of the entered information, enter `y` to re-start the basic configuration script.

```
Please review the config
Management network:
  Gateway: 192.168.9.1
  IP Address/Mask: 192.168.9.172/24

Re-enter config? (y/N): N
```

Step 2 Wait for the process to complete.

After you enter and confirm management network information of the first node, the initial setup configures the networking and brings up the UI, which you will use to add two and configure other nodes and complete the cluster deployment.

```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

System UI online, please login to https://192.168.9.172 to continue.
```

Step 3 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

Enter the password you entered in a previous step and click **Login**

Step 4 After logging in, review the **Meet Nexus Dashboard** welcome screens, then proceed to the **Journey** screen. Click on **Go** under the **Cluster bringup** step to begin the bootstrap process.

Step 5 Enter the requested information in the **Basic Information** page of the **Cluster Bringup** wizard.

a) For **Cluster Name**, enter a name for this Nexus Dashboard cluster.

The cluster name must follow the RFC-1123 requirements.

b) For **Select the Nexus Dashboard Implementation type**, choose either **LAN** or **SAN** then click **Next**.

Step 6 Enter the requested information in the **Configuration** page of the **Cluster Bringup** wizard.

a) (Optional) If you want to enable IPv6 functionality for the cluster, put a check in the **Enable IPv6** checkbox.

b) Click **+Add DNS provider** to add one or more DNS servers, enter the DNS provider IP address, then click the checkmark icon.

c) (Optional) Click **+Add DNS search domain** to add a search domain, enter the DNS search domain IP address, then click the checkmark icon.

d) (Optional) If you want to enable NTP server authentication, put a check in the **NTP Authentication** checkbox.

e) If you enabled NTP authentication, click **+ Add Key**, enter the required information, and click the checkmark icon to save the information.

- **Key**–Enter the NTP authentication key, which is a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP servers. You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP authentication key.
- **ID**–Enter a key ID for the NTP host. Each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Authentication Type**–Choose authentication type for the NTP key.
- Put a check in the **Trusted** checkbox if you want this key to be trusted. Untrusted keys cannot be used for NTP authentication.

For the complete list of NTP authentication requirements and guidelines, see General prerequisites and guidelines.

If you want to enter additional NTP keys, click **+ Add Key** again and enter the information.

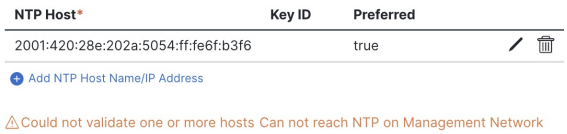
f) If you enabled NTP authentication, click **+Add NTP Host Name/IP Address**, enter the required information, and click the checkmark icon to save the information.

- **NTP Host**–Enter an IP address; fully qualified domain names (FQDN) are not supported.
- **Key ID**–Enter the key ID of the NTP key you defined in the previous substep.

If NTP authentication is disabled, this field is grayed out.

- Put a check in the **Preferred** checkbox if you want this host to be preferred.

Note If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and entered an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet and is unable to connect to an IPv6 address of the NTP server. You will enter IPv6 address in the next step. In this case, enter the other required information as described in the following steps and click **Next** to proceed to the next page where you will enter IPv6 addresses for the nodes.

If you want to enter additional NTP servers, click **+Add NTP Host Name/IP Address** again and enter the information.

g) For **Proxy Server**, enter the URL or IP address of a proxy server.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

You can click **+Add Ignore Host** to enter one or more destination IP addresses for which traffic will skip using the proxy.

The proxy server must permit these URLs:

```
svc.intersight.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

If you do not want to configure a proxy, click **Skip Proxy** then click **Confirm**.

h) (Optional) If your proxy server requires authentication, put a check in the **Authentication required for Proxy** checkbox and enter the login credentials.

i) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure these settings:

- **App Network**–The address space used by the application's services running in the Nexus Dashboard. Enter the IP address and netmask.
- **Service Network**–An internal network used by Nexus Dashboard and its processes. Enter the IP address and netmask.
- **App Network IPv6**–If you put a check in the **Enable IPv6** checkbox earlier, enter the IPv6 subnet for the app network.
- **Service Network IPv6**–If you put a check in the **Enable IPv6** checkbox earlier, enter the IPv6 subnet for the service network.

For more information about the application and service networks, see General prerequisites and guidelines.

j) Click **Next**.

Step 7 In the **Node Details** page, update the first node's information.

You have defined the Management network and IP address for the node into which you are currently logged in during the initial node configuration in earlier steps, but you must also enter the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.

a) For **Cluster Connectivity**, if your cluster is deployed in L3 mode, choose **BGP**. Otherwise, choose **L2**.

BGP configuration is required for the persistent IP addresses feature used by telemetry. This feature is described in more detail in the BGP configuration and persistent IP addresses and Nexus Dashboard persistent IP addresses sections.

Note You can enable BGP at this time or in the Nexus Dashboard GUI after the cluster is deployed. All remaining nodes need to configure BGP if it is configured. You must enable BGP now if the data network of nodes have different subnets.

b) Click the **Edit** button next to the first node.

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated, but you must enter the other information.

c) For **Name**, enter a name for the node.

The node's **Name** will be set as its hostname, so it must follow the RFC-1123 requirements.

Note If you need to change the name but the **Name** field is not editable, run the CIMC validation again to fix this issue.

d) For **Type**, choose **Primary**.

The first nodes of the cluster must be set to **Primary**. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address configuration later, you would need to redeploy the cluster.

All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4/IPv6.

f) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

g) Click **Save** to save the changes.

Step 8 If you are deploying a multi-node cluster, in the **Node Details** page, click **Add Node** to add the second node to the cluster.

a) In the **Deployment Details** area, enter the **CIMC IP Address**, **Username**, and **Password** for the second node.

Note For **Username** for the second node, enter the admin user ID.

b) Click **Validate** to verify connectivity to the node.

The node's serial number is automatically populated after CIMC connectivity is validated.

c) For **Name**, enter the name for the node.

The node's name will be set as its hostname, so it must follow the RFC-1123 requirements.

d) For **Type**, choose `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Management Network** area, enter the node's management network information.

You must enter the management network IP address, netmask, and gateway.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4/IPv6.

f) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address configuration later, you would need to redeploy the cluster.

All nodes in the cluster must be configured with either only IPv4, only IPv6, or dual stack IPv4 and IPv6.

g) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

h) Click **Save** to save the changes.

i) Repeat this step for the final (third) primary node of the cluster.

Step 9 (Optional) Repeat the previous step to enter information about any additional secondary or standby nodes.



Note To support higher scale, you must provide a sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.

You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 10 In the **Node Details** page, verify the information that you entered, then click **Next**.

Step 11 In the **Persistent IPs** page, if you want to add more persistent IP addresses, click **+ Add Data Service IP Address**, enter the IP address, and click the checkmark icon. Repeat this step as many times as desired, then click **Next**.

You must configure the minimum number of required persistent IP addresses during the bootstrap process. This step enables you to add more persistent IP addresses if desired.

Step 12 In the **Summary** page, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

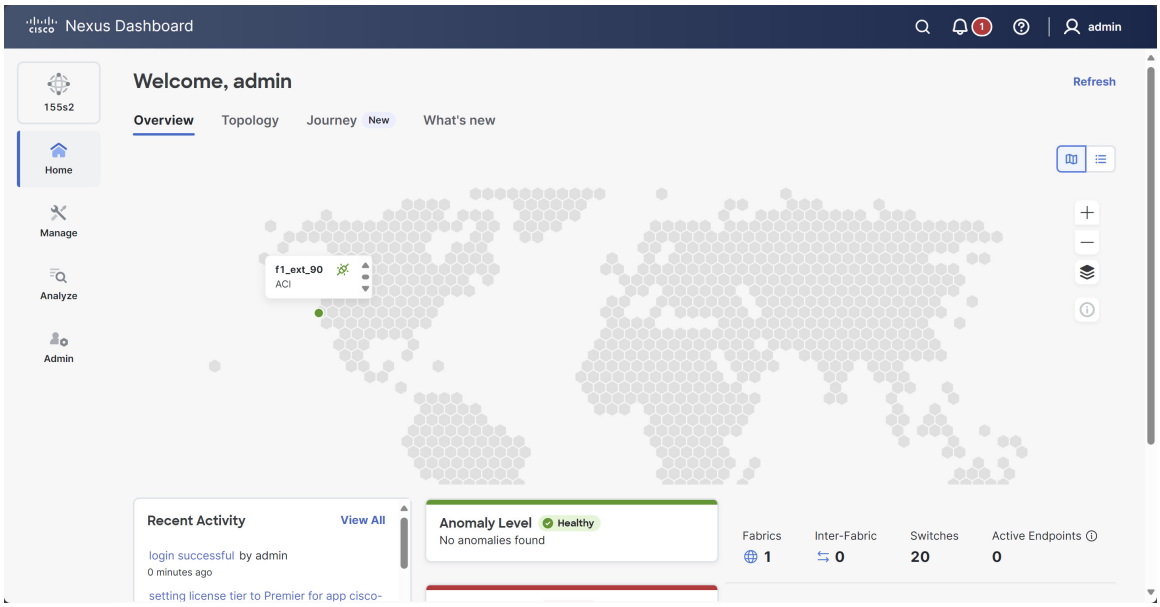
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

It may take up to 60 minutes or more for the cluster to form, depending on the number of nodes in the cluster, and all the features to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 13 Verify that the cluster is healthy.

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled".

After all the cluster is deployed and all services are started, you can look at the **Anomaly Level** on the **Home** > **Overview** page to ensure the cluster is healthy:



Alternatively, you can log in to any one node using SSH as the `rescue-user` using the password you entered during node deployment and using the `acs health` command to see the status:

- While the cluster is converging, you may see the following output:

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

```
$ acs health
All components are healthy
```

Note In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

```
deploy base system services
```

This is due to an issue with `etcd` on the node after a reboot of the physical Nexus Dashboard cluster.

To resolve the issue, enter the `acs reboot clean` command on the affected node.

Step 14 (Optional) Connect your Cisco Nexus Dashboard cluster to Cisco Intersight for added visibility and benefits. Refer to *Working with Cisco Intersight* for detailed steps.

Step 15 After you have deployed Nexus Dashboard, see the collections page for this release for configuration information.

What to do next

The next task is to create the fabrics and fabric groups. See the *Creating Fabrics and Fabric Groups* article for this release on the Cisco Nexus Dashboard collections page.