

Deploying as Physical Appliance

- Prerequisites and guidelines for deploying Nexus Dashboard as a physical appliance, on page 1
- Physical Node Cabling, on page 5
- Deploying Nexus Dashboard as Physical Appliance, on page 7

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 general and service-specific prerequisites described in Prerequisites: Nexus Dashboard.
- Review and complete any additional prerequisites that is described in the *Release Notes* for the services you plan to deploy.

  You can find the service-specific documents at the following links:

  - Nexus Dashboard Fabric Controller Release Notes
  - Nexus Dashboard Insights Release Notes
  - Nexus Dashboard Orchestrator Release Notes

- 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 on the UCS-C220-M5 (`SE-NODE-G2`) and UCS-C225-M6 (`ND-NODE-L4`) original Cisco Nexus Dashboard platform hardware only. The following table lists the PIDs and specifications of the physical appliance servers:

Table 1: Supported UCS-C220-M5 Hardware

Product ID	Hardware
SE-NODE-G2	• Cisco UCS C220 M5 Chassis • 2x 10-core 2.2-GHz Intel Xeon Silver CPU • 256 GB of RAM • 4x 2.4-TB HDDs 400-GB SSD 1.2-TB NVME drive • Cisco UCS Virtual Interface Card 1455 (4x25G Ports) • 1050-W power supply
SE-CL-L3	A cluster of 3x `SE-NODE-G2=` appliances.

Table 2: Supported UCS-C225-M6 Hardware

Product ID	Hardware

ND-NODE-L4	• Cisco UCS C225 M6 Chassis
	• 2.8-GHz AMD CPU
	• 256 GB of RAM
	• 4x 2.4-TB HDDs
	960-GB SSD
	1.6-TB NVME drive
	• Intel X710T2LG 2x10 GbE (Copper)
	• One of the following:
	  • Intel E810XXVDA2 2x25/10 GbE (Fiber Optic)
	  • Cisco UCS Virtual Interface Card 1455 (4x25G Ports)
	• 1050-W power supply

ND-CLUSTER-L4	A cluster of 3x `ND-NODE-L4=` appliances.

Note The above hardware supports Cisco Nexus Dashboard software only. 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.

Note

- For Nexus Dashboard release 3.2(1), do not upgrade the CIMC version to 4.3(5.x) or higher.
- For Nexus Dashboard release 3.2(2), if you are going to upgrade your UCS server firmware to 4.3(5x), you must first upgrade or install your Nexus Dashboard release to 3.2(2) before you can upgrade your UCS server firmware to 4.3(5x). See the "Supported UCS server firmware" section in the Nexus Dashboard and Services Release Notes, Release 3.2.2 for the UCS server firmware versions that are supported.

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

  See Configure a Cisco Integrated Management Controller IP address, on page 4.

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

  See Enable Serial over LAN in the Cisco Integrated Management Controller, on page 4.

  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 release 2.3.2 image pre-installed, but you want to deploy release 3.2.1 instead, we recommend:

  1. First, bring up the release 2.3.2 cluster, as described in the deployment guide for that release.
  2. Then upgrade to release 3.2.1, as described in Upgrading Existing ND 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 3-node cluster. Extra secondary nodes can be added for horizontal scaling if required by the number of services you deploy. 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
```

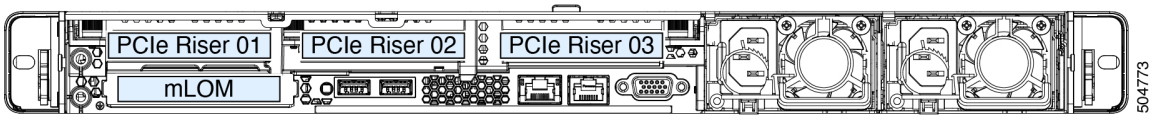
Step 3 In the command output, verify that `com0` is the com port for SoL.

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

Physical Node Cabling

Physical nodes can be deployed in UCS-C220-M5 (`SE-NODE-G2`) and UCS-C225-M6 (`ND-NODE-L4`) physical servers with the following guidelines:

Figure 1: mLOM and PCIe Riser 01 Card Used for Node Connectivity



- Both servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The UCS-C220-M5 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
- The UCS-C225-M6 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.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode.
- For management network:

- You must use the `mgmt0` and `mgmt1` on the mLOM card.
- All ports must have the same speed, either 1G or 10G.

- For data network:
  - On the UCS-C220-M5 server, you must use the VIC1455 card.
  - On the UCS-C225-M6 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 interfaces must be connected to individual host-facing switch ports; Fabric Extenders (FEX), PortChannel (PC), and Virtual PortChannel (vPC) are not supported.
  - Configure Cisco VIC-connected data uplinks as trunk ports on the connected switches. Add a native VLAN on the upstream trunk port if VLAN tagging is not used in Nexus Dashboard for the data network.



Note When connecting Nexus Dashboard nodes to Arista switches on the data interfaces (fabric interfaces), configure the switch interfaces as access ports rather than trunk ports, with the access VLAN set to the data subnet VLAN.

  - Verify that the port-channel feature is not enabled 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 and verify that the option is not enabled (check box is unchecked).

       If the **Port Channel** option is enabled, remove the check from the check box before beginning the Nexus Dashboard deployment process.

  - All ports must have the same speed, either 10G or 25G.
  - Port-1 corresponds to `fabric0` in Nexus Dashboard and Port-2 corresponding to `fabric1`.

    You can use both `fabric0` and `fabric1` for data network connectivity.

Note When using a 4-port card, the order of ports depends on the model of the server you are using:

- On the UCS-C220-M5 server, the order from left to right is Port-1, Port-2, Port-3, Port-4.
- On the UCS-C225-M6 server, the order from left to right is Port-4, Port-3, Port-2, Port-1.
- Port-1 and Port-2 OR Port-3 and Port-4 are used in pairs. We recommend that you use the Port-1/Port-2 pair.

- If you connect the nodes to Cisco Catalyst switches, you must add `switchport voice vlan dot1p` command to the switch interfaces.

  On the Catalyst switches, packets are tagged 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.

Note This command is not supported on Cisco Catalyst 9500X or 9600X models. For those models, you must configure the interface as a trunk interface and specify the VLAN in the Nexus Dashboard data network configuration.

Deploying Nexus Dashboard as Physical Appliance

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. This section describes how to configure and bring up the initial Nexus Dashboard cluster.

Before you begin

- Ensure that you 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 3.2.1
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 pure IPv6 mode, provide 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 initial bootstrap process to complete.

After you provide 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.
```

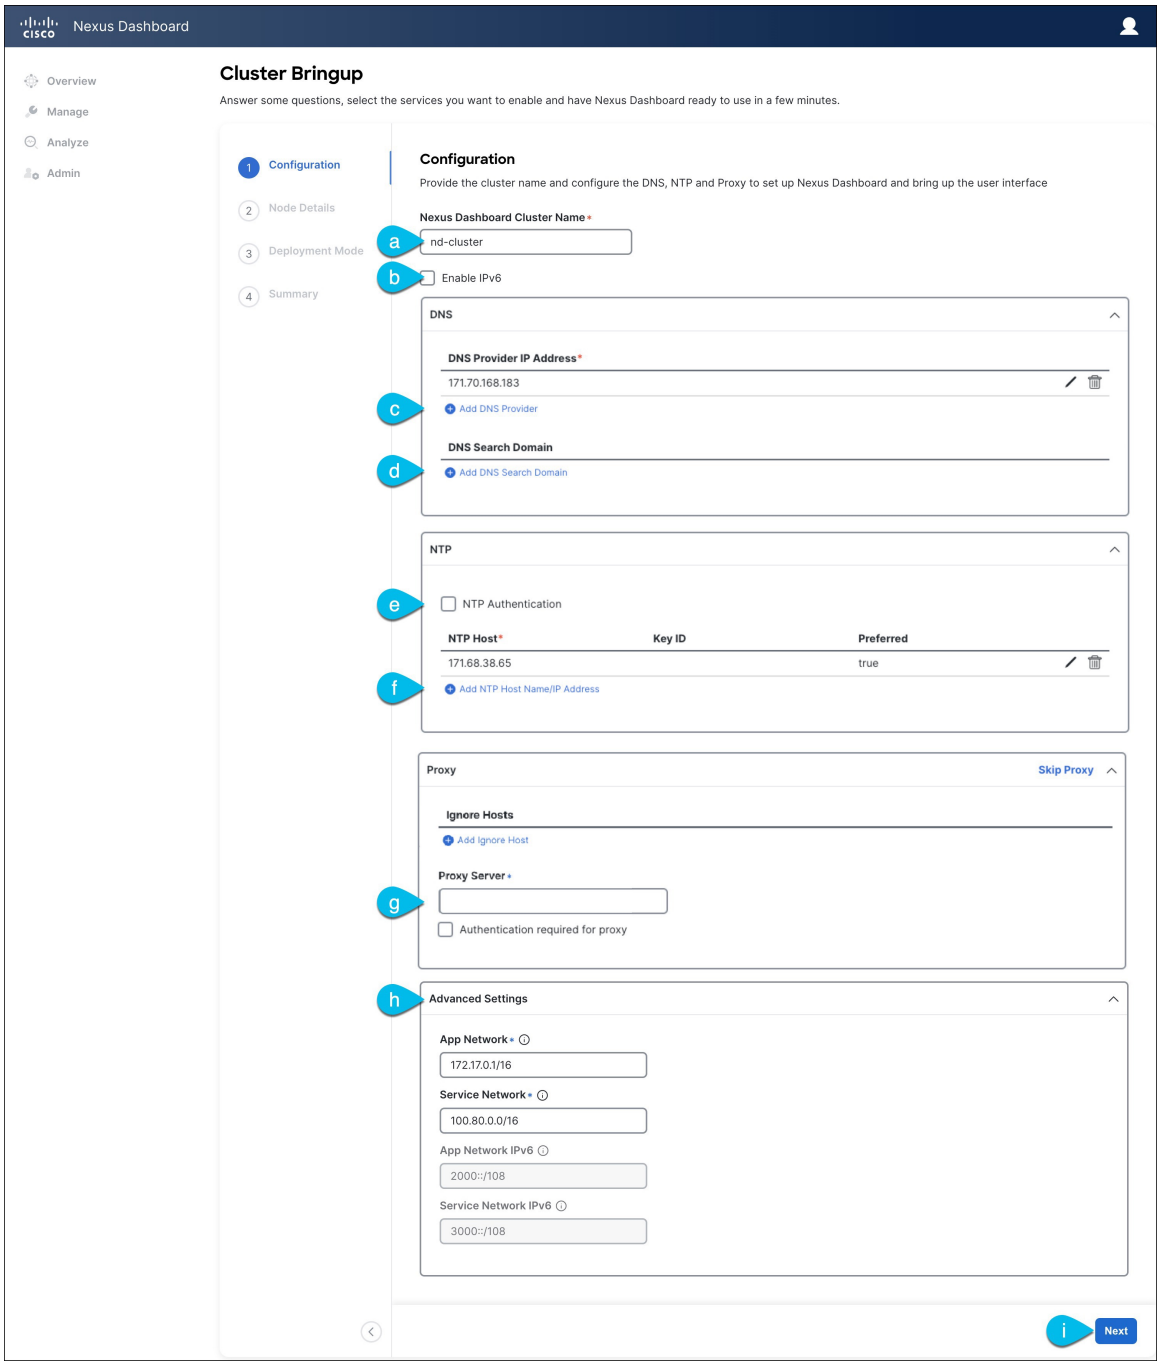
Step 3 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

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 provided in a previous step and click **Login**

Step 4 Provide the **Cluster Details**.

In the **Cluster Details** screen of the **Cluster Bringup** wizard, provide the following information:



a) Provide the **Cluster Name** for this Nexus Dashboard cluster.

- The cluster name must follow the RFC-1123 requirements.
- If you use multi-cluster connectivity to establish connectivity between multiple Nexus Dashboard clusters, the names of the clusters that you plan to connect together must be unique.

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

c) Click **+Add DNS Provider** to add one or more DNS servers.

After you've entered the information, click the checkmark icon to save it.

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

After you've entered the information, click the checkmark icon to save it.

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – 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.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
After you've entered the information, click the checkmark icon to save it.

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and guidelines for all enabled services.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

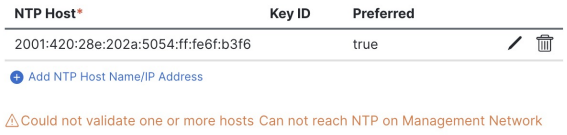
In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

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 provided 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 (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

g) Provide a **Proxy Server**, then click **Validate** it.

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 also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

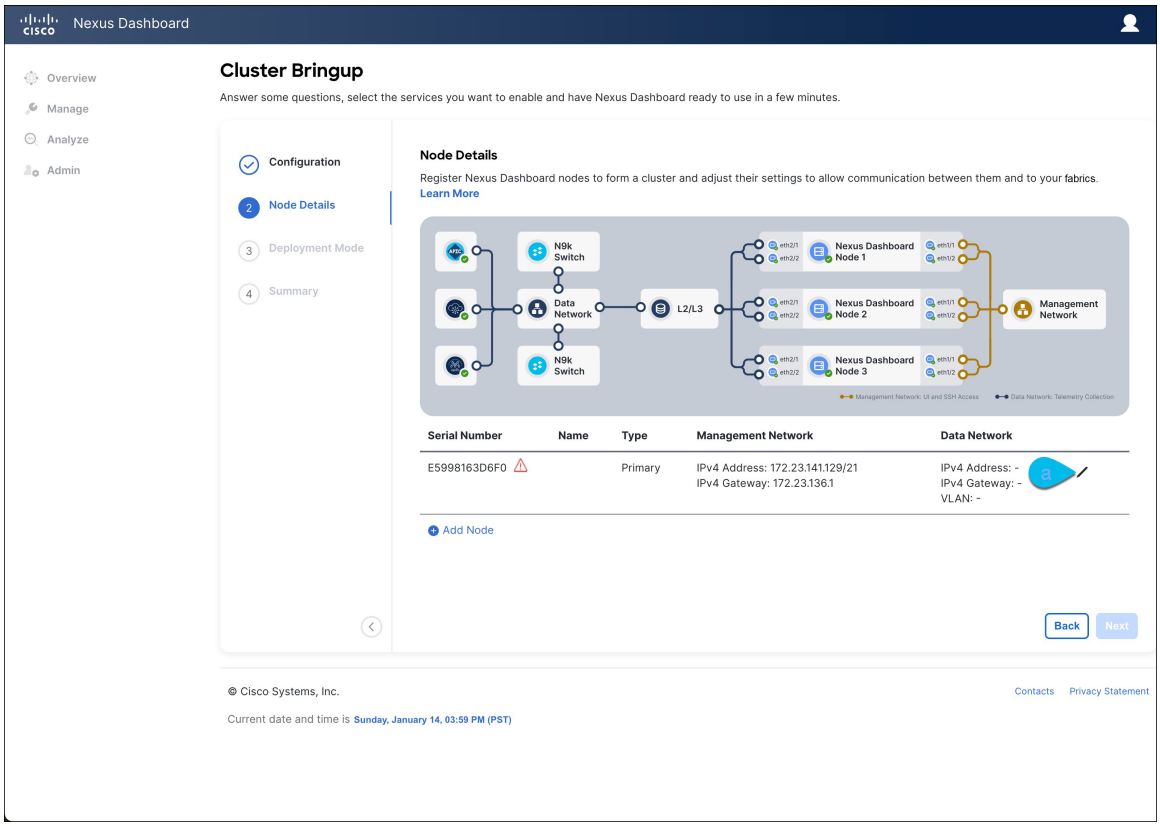
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

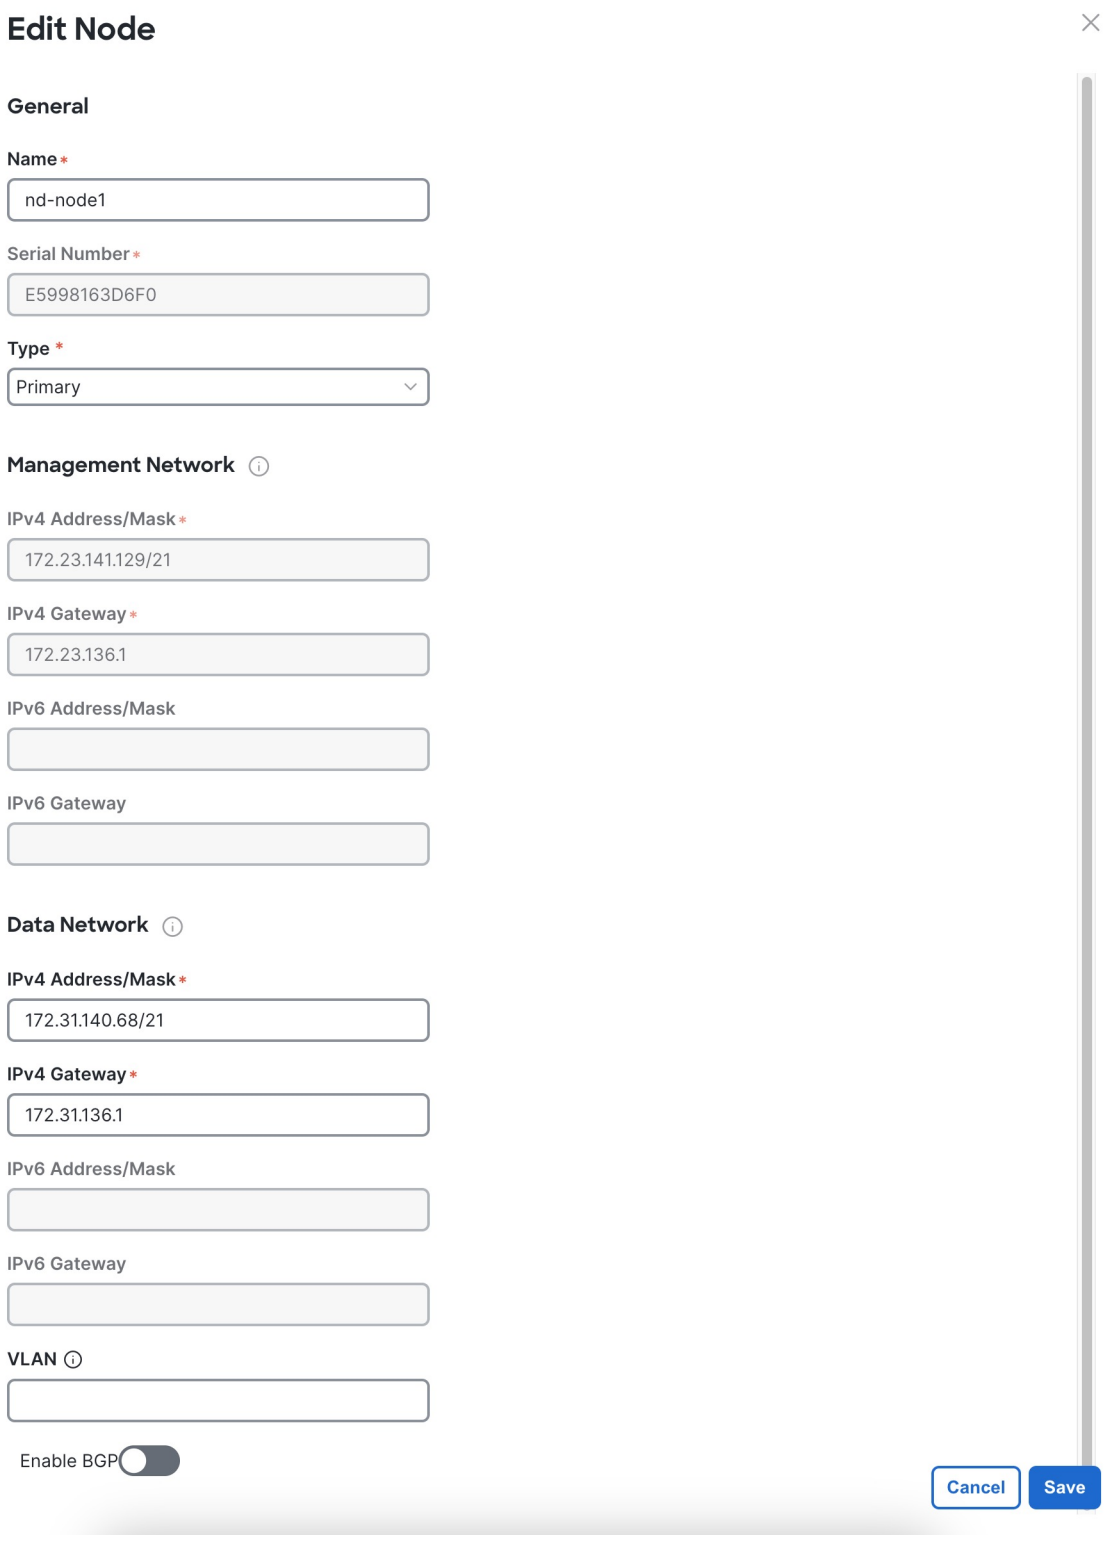
  Application and Services networks are described in the Prerequisites and guidelines for all enabled services section earlier in this document.

j) Click **Next** to continue.

Step 5 In the **Node Details** screen, 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 provide the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.


Nexus Dashboard
Overview
Manage
Analyze
Admin
Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.
Configuration
2 Node Details
3 Deployment Mode
4 Summary
Node Details
Register Nexus Dashboard nodes to form a cluster and adjust their settings to allow communication between them and to your fabrics.
Learn More
N9k Switch
Data Network
N9k Switch
L2/L3
Nexus Dashboard Node 1
Nexus Dashboard Node 2
Nexus Dashboard Node 3
Management Network
Serial Number
Name
Type
Management Network
Data Network
E5998163D6F0
Primary
IPv4 Address: 172.23.141.129/21
IPv4 Gateway: 172.23.136.1
IPv4 Address: -
IPv4 Gateway: -
VLAN: -
Add Node
Back
Next
© Cisco Systems, Inc.
Contacts
Privacy Statement
Current date and time is Sunday, January 14, 03:59 PM (PST)



a) 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 provide other information.

b) Provide the **Name** for the node.

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

c) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

d) In the **Data Network** area, provide the node's **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP 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.

e) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and guidelines for all enabled services and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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.

If you choose to enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

  The router ID must be an IPv4 address, for example `1.1.1.1`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

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

Step 6 In the **Node Details** screen, click **Add Node** to add the second node to the cluster.

If you are deploying a single-node cluster, skip this step.

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

Note
For the **Username** field for the second node, enter the admin user ID (the ID 1 of the local user management).

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

The node's **Serial Number** is automatically populated after CIMC connectivity is validated.

c) Provide the **Name** for the node.

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

d) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

e) In the **Management Network** area, provide the node's **Management Network** information.

You must provide the Management network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, you must also provide 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, provide the node's **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP 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.

g) (Optional) If required, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and guidelines for all enabled services and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

Note
You can enable BGP at this time or in the Nexus Dashboard GUI after the cluster is deployed.

If you choose to enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

  The router ID must be an IPv4 address, for example `1.1.1.1`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

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

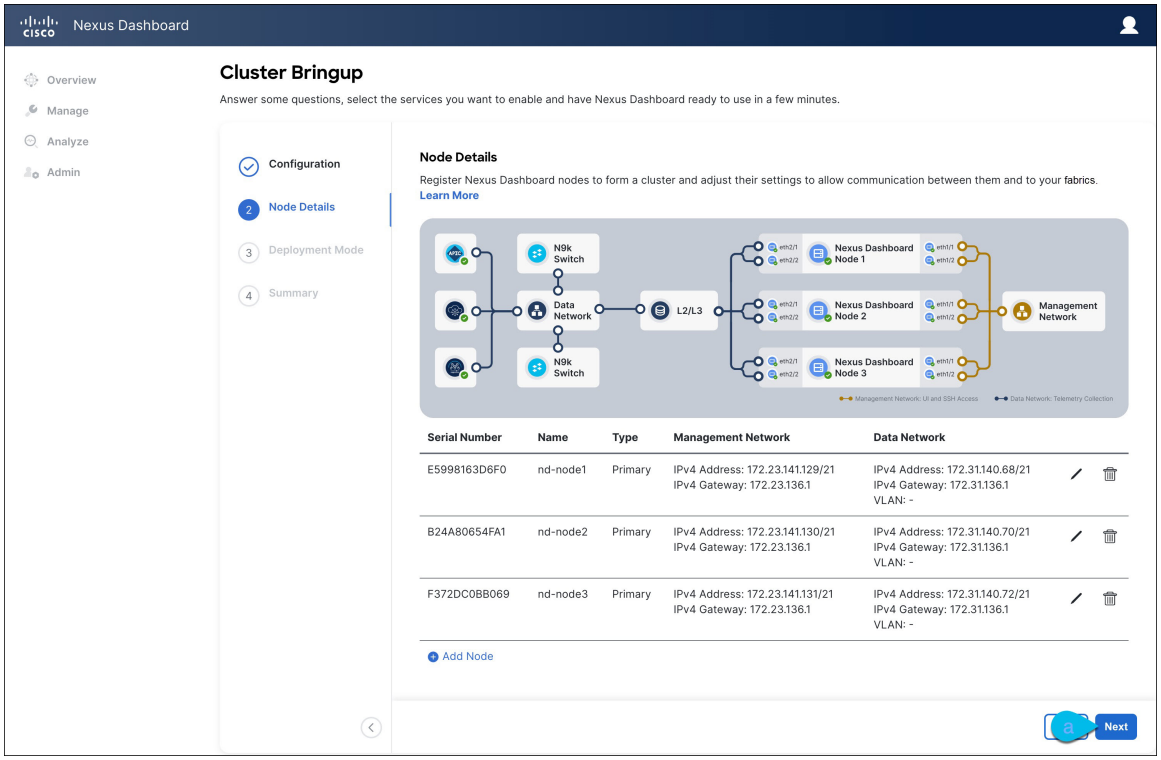
i) Repeat this step for the final (third) primary node of the cluster.

Step 7 (Optional) Repeat the previous step to provide information about any additional secondary or standby nodes.

Note
In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide 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 8 In the **Node Details** page, verify the provided information and click **Next** to continue.



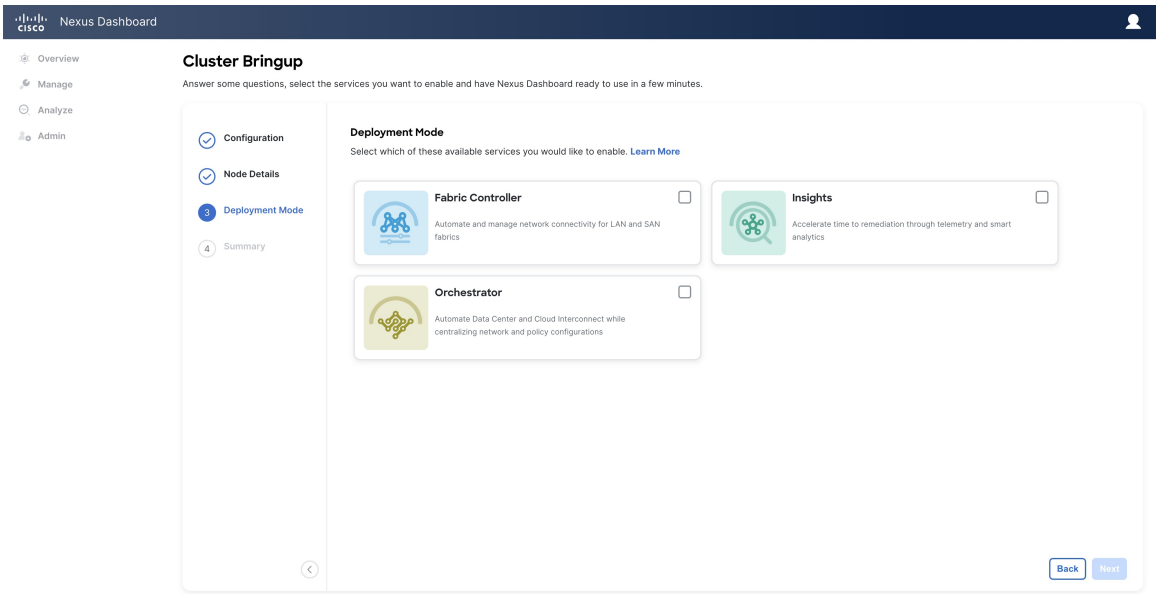
Step 9 Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Note
Depending on the number of nodes in the cluster, some services or cohosting scenarios may not be supported. If you are unable to choose the desired number of services, click **Back** and ensure that you have provided enough secondary nodes in the previous step.

The deployment mode cannot be changed after the cluster is deployed, so you must ensure that you have completed all service-specific prerequisites described in earlier chapters of this document:

- Nexus Dashboard Fabric Controller prerequisites
- Nexus Dashboard Orchestrator prerequisites
- Nexus Dashboard Insights prerequisites



b) If you chose a deployment mode that includes Fabric Controller or Insights, click **Add Persistent Service IPs/Pools** to provide one or more persistent IPs required by Insights or Fabric Controller services.

Note that you can skip this step at this point if necessary, and you can add persistent IPs after the cluster comes up. For more information about persistent IPs, see the Prerequisites and guidelines for all enabled services section and the service-specific requirements chapters.

c) Click **Next** to proceed.

Step 10 In the **Summary** screen, 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.

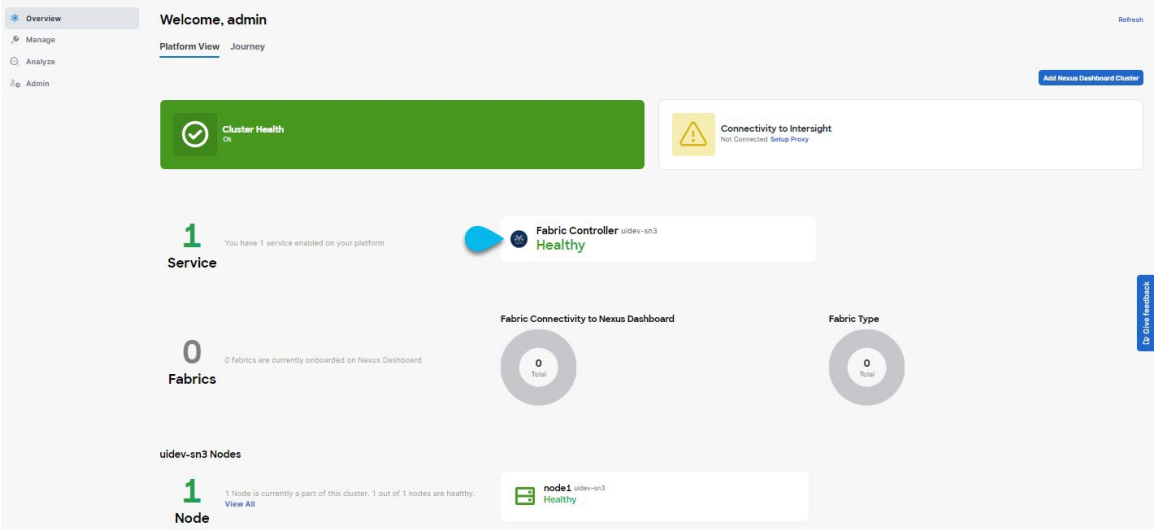
It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 11 Verify that the cluster is healthy.

Depending of the deployment mode, it may take more than 30 minutes for the cluster to form and all the services to start.

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 check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ 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 pND (Physical Nexus Dashboard) cluster.

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

Step 12 After you have deployed your Nexus Dashboard and services, you can configure each service as described in its configuration and operations articles.

- For Fabric Controller, see the *NDFC persona configuration* white paper and documentation library.
- For Orchestrator, see the documentation page.
- For Insights, see the documentation library.