

Deploying as Physical Appliance

- Prerequisites and Guidelines, on page 1
- Deploying Nexus Dashboard as Physical Appliance, on page 4

Prerequisites and Guidelines

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

- Review and complete the general prerequisites described in Deployment Overview and Requirements.

  If you are looking to completely re-image the server, for example in case you cannot sign in as the `rescue-user` for manual recovery, see the *Troubleshooting*.

- Ensure that you have enough physical nodes to support your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor and the specific services you plan to deploy. You can use the Nexus Dashboard Capacity Planning tool to verify that the virtual form factor satisfies your deployment requirements.

Note The capacity planning tool lists the total number of primary and worker nodes required. In addition, this form factor supports up to 2 standby nodes.

This document describes how to initially deploy the base Cisco Nexus Dashboard cluster. If you want to expand an existing cluster with another nodes (such as `worker` or `standby`), see the *Infrastructure Management* article instead, which is also available from the Cisco Nexus Dashboard UI.

- 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*.

  The physical appliance form factor is supported on the UCS-C220-M5 and UCS-C225-M6 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

Process 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

Process 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) 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.

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

  To configure a CIMC IP address:

  1. Power on the server.

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

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

  3. Provide the following information:

     - Set **NIC mode** to `Dedicated`.
     - 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.

     - Ensure that **NIC Redundancy** is set to `None`.
     - Press **F1** for more options such as hostname, DNS, default user passwords, port properties, and reset port profiles.

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

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

  SoL is required for the `connect host` command, which you use to connect to the node to provide basic configuration information. To use the SoL, you must first enable it on your CIMC. SSH into the node using the CIMC IP address and enter the sign-in credentials. Run the following 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


C220-WZP23150D4C /sol #
```

- 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 2.2.1 image pre-installed, but you want to deploy Release 3.0.1 instead, we recommend:

  1. First, bring up the Release 2.2.1 cluster, as described in the deployment guide for that release.
  2. Then upgrade to Release 3.0.1, as described in Upgrading Nexus Dashboard.



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 worker nodes can be added for horizontal scaling if required by the enter and number of services you deploy. For the maximum number of `worker` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

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 3-node Nexus Dashboard cluster.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, 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.

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

You will be prompted to run the first-time setup utility:

```
[ OK ] Started atomix-boot-setup.
       Starting Initial cloud-init job (pre-networking)...
       Starting logrotate...
       Starting logwatch...
       Starting keyhole...
[ OK ] Started keyhole.
[ OK ] Started logrotate.
[ OK ] Started logwatch.

Press any key to run first-boot setup on this console...
```

b) 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:
```

c) 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.

d) Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, choose `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, the initial setup configures the networking and brings up the UI, which you will use to add two other nodes and complete the cluster deployment.

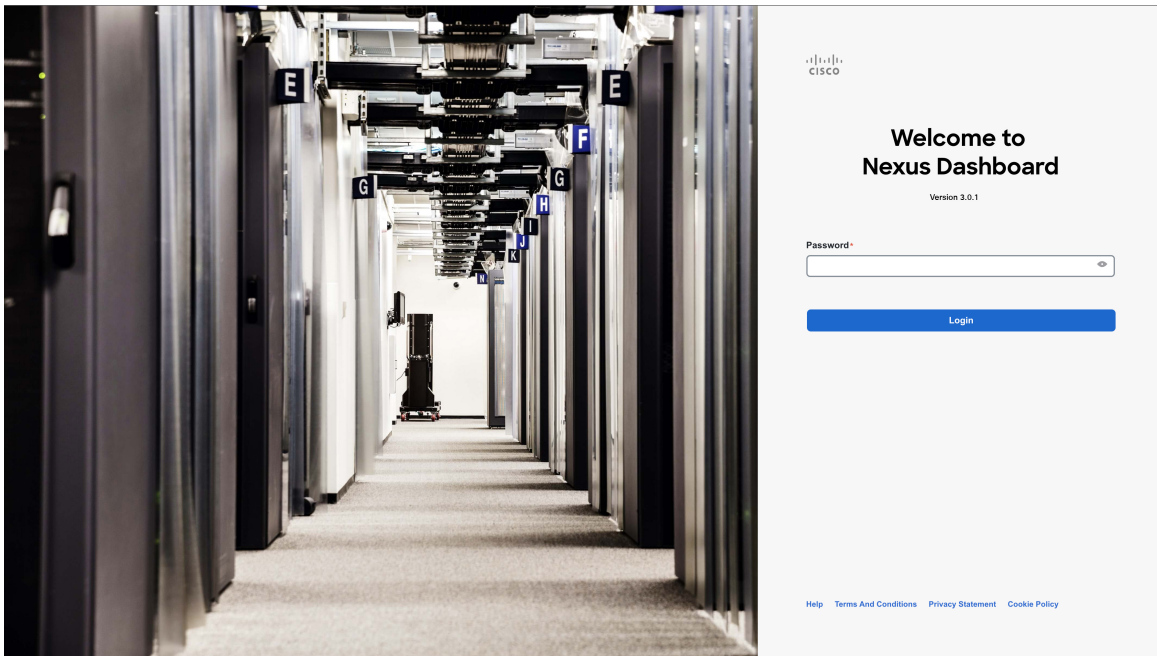
```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

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.

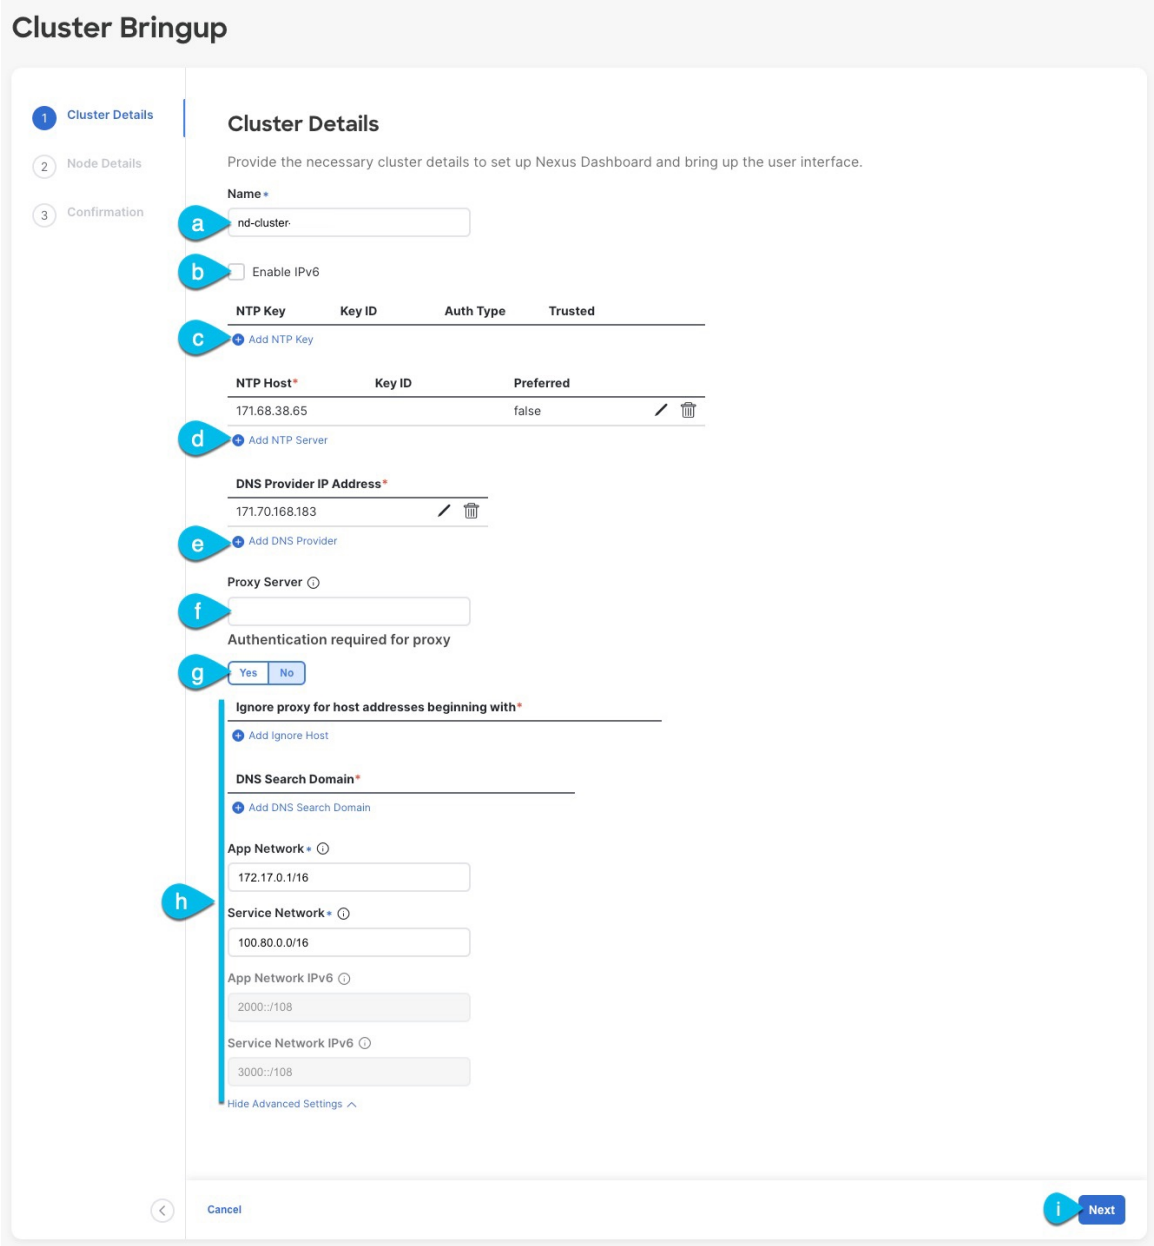
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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.

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

c) (Optional) If you want to enable NTP server authentication, 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
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines.

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

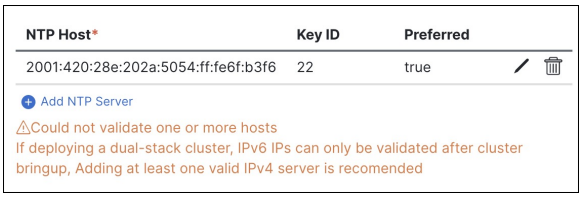
d) Click **+Add NTP Host** to add one or more NTP servers.

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.
- 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.

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

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

f) Provide 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.

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, mouse over the information (**i**) icon next to the field, then click **Skip**.

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

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

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

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

- 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.

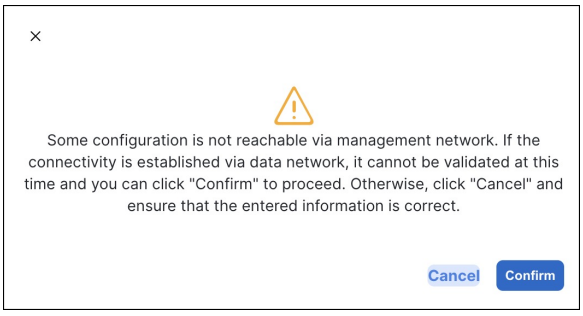
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  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 section earlier in this document.

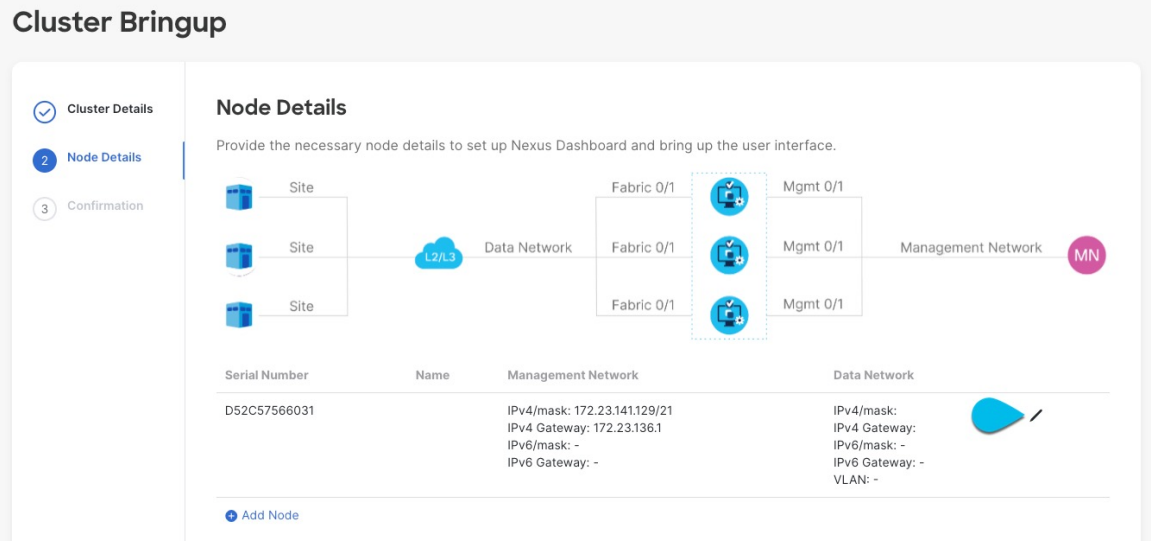
i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



Step 5 In the **Node Details** screen, update the current 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.



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

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

The node's **Serial Number** and the **Management Network** information are automatically populated.

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

c) 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, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines 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.

When you 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.

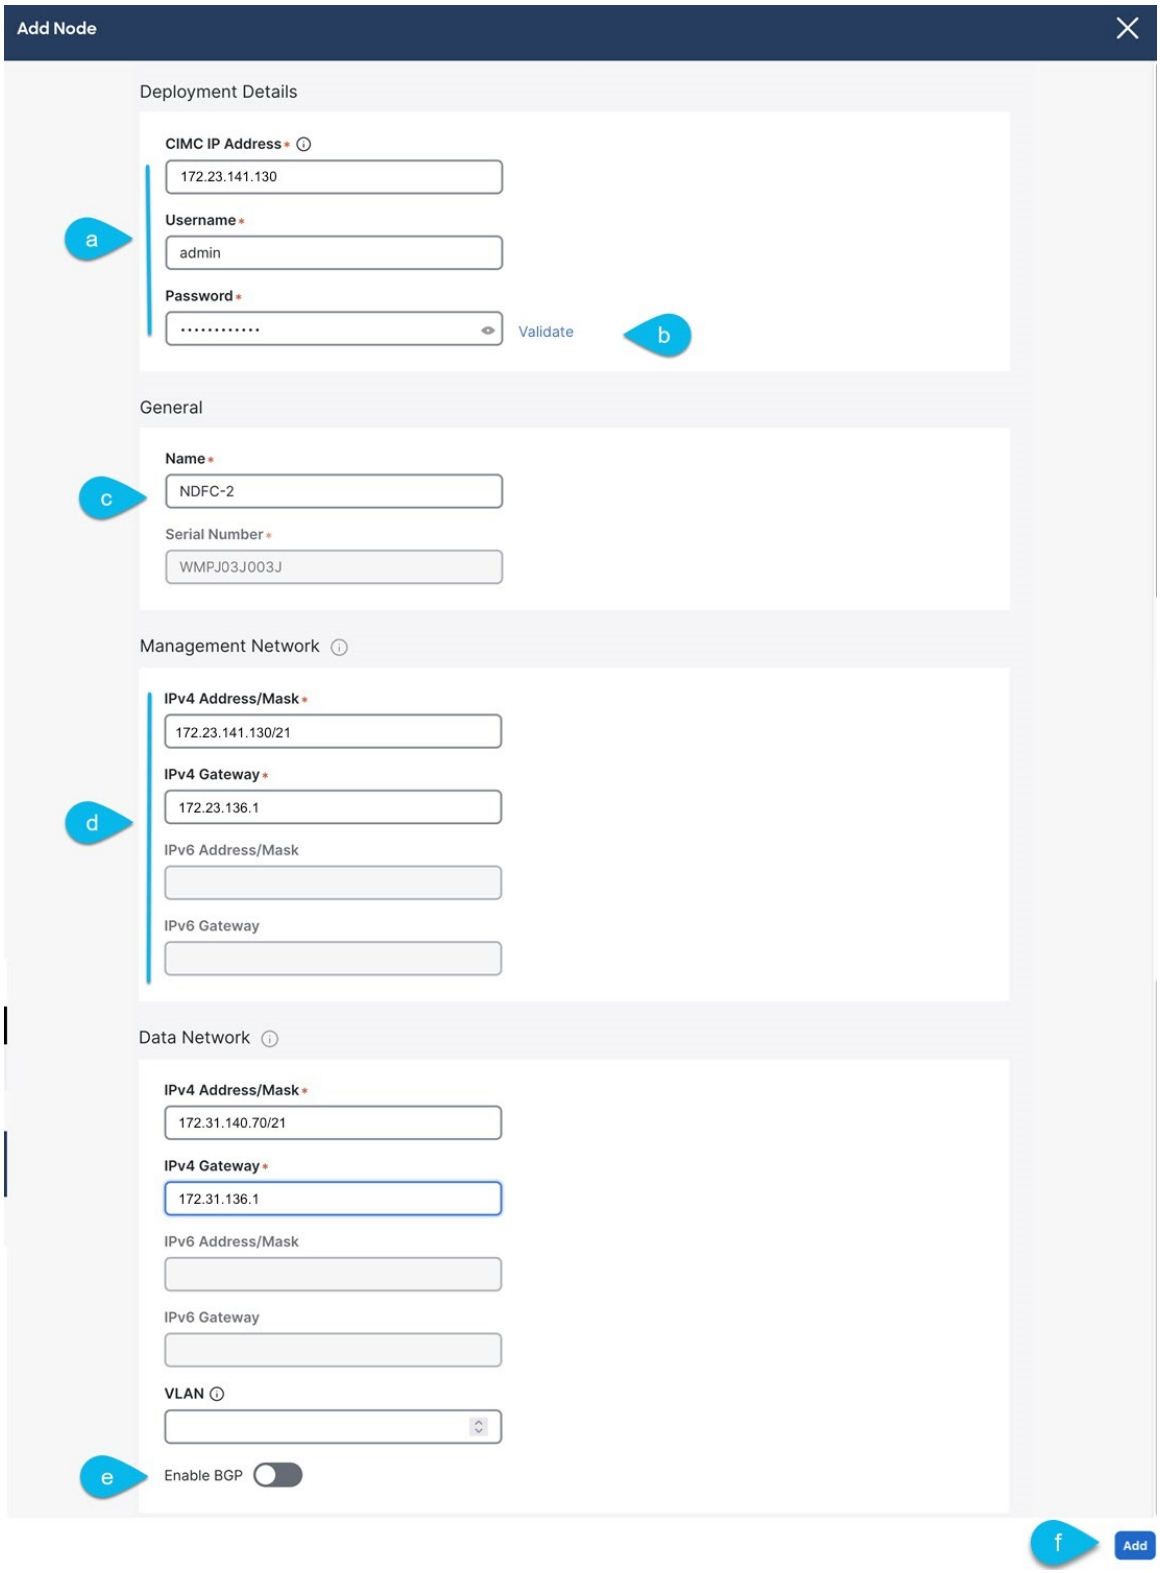
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

After network connectivity is validated, you can provide the other required information for the node.

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

The node's **Serial Number** and the **Management Network** information are automatically populated after CIMC connectivity is validated.

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, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines 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.

When you 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 `2.2.2.2`

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

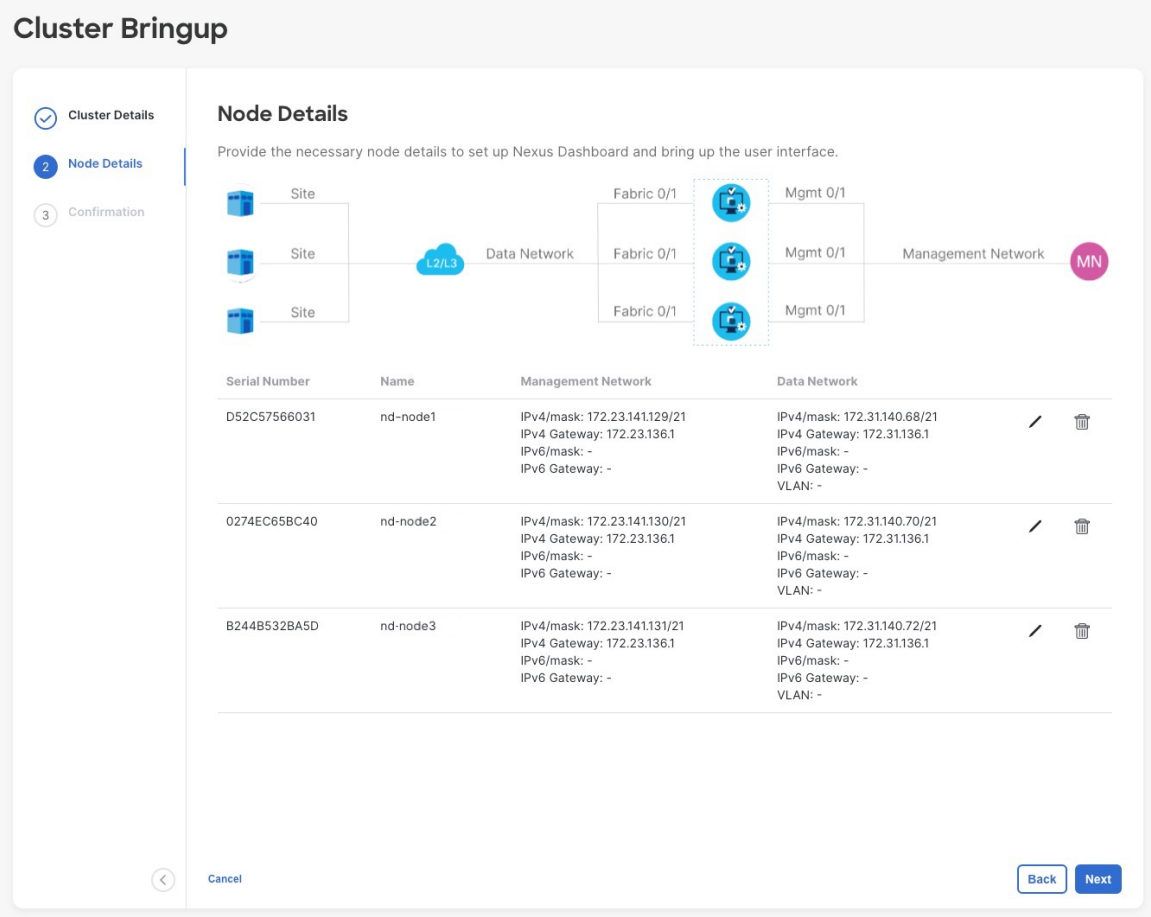
f) Click **Add** to save the changes.

Step 7 Repeat the previous step to add the 3rd node.

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

Step 8 In the **Node Details** page, click **Next** to continue.

After you have provided the management and data network information for all nodes, you can proceed to the final **Confirmation** screen.



Step 9 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create 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 10 Verify that the cluster is healthy.

It may take up to 30 minutes for the cluster to form and all the services to start.

After all three nodes are ready, you can log in to any one node via SSH as the `rescue-user` using the password you provided during node deployment and run the following command to verify cluster health:

a) Verify that the cluster is up and running.

You can check the current status of cluster deployment by logging in to any of the nodes and running the `acs health` command.

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
```

b) Log in to the Nexus Dashboard GUI.

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 of the Nexus Dashboard cluster.

Step 11 Configure the **Network Scale** parameters for your cluster.

This is described in the **Infrastructure Management** > **Cluster Configuration** section of the *Cisco Nexus Dashboard User Guide*, which is also available directly from your Nexus Dashboard's Help Center.