

# Deploying in Linux KVM


- Prerequisites and Guidelines, on page 1
- Deploying Nexus Dashboard in Linux KVM, on page 2


# Prerequisites and Guidelines

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

- Ensure that the KVM form factor supports your scale and services requirements.

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

- Review and complete the general prerequisites described in Prerequisites: Nexus Dashboard.
- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.
- Cisco does not support the use of nested virtualization environments. Deploying Nexus Dashboard on a virtual machine that is itself running on a virtualized hypervisor (for example, KVM on ESXi) is not a supported configuration and may result in performance degradation or system instability.
- Ensure you have enough system resources:

*Table 1: Deployment Requirements*

| Requirements |
| --- |
| • KVM deployments are supported for Nexus Dashboard Fabric Controller services only.<br>• You must deploy in CentOS 7.9 or Red Hat Enterprise Linux 8.6<br>• You must have the supported versions of Kernel and KVM:<br>  • For CentOS 7.9, Kernel version `3.10.0-957.el7.x86_64` and KVM version `libvirt-4.5.0-23.el7_7.1.x86_64`<br>  • For RHEL 8.6, Kernel version `4.18.0-372.9.1.el8.x86_64` and KVM version `libvert 8.0.0`<br>• 16 vCPUs<br>• 64 GB of RAM<br>• 550 GB disk<br>Each node requires a dedicated disk partition<br>• The disk must have I/O latency of 20ms or less.<br>To verify the I/O latency:<br>1. Create a test directory.<br>For example, `test-data`.<br>2. Run the following command:<br>`# fio --rw=write --ioengine=sync --fdatasync=1 --directory=test-data --size=22m --bs=2300 --name=mytest`<br>3. After the command is executed, confirm that the `99.00th=[<value>]` in the `fsync/fdatasync/sync_file_range` section is under 20ms.<br>• We recommend that each Nexus Dashboard node is deployed in a different KVM hypervisor. |

# Deploying Nexus Dashboard in Linux KVM

This section describes how to deploy Cisco Nexus Dashboard cluster in Linux KVM.

### Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 1.

### Procedure

**Step 1** Download the Cisco Nexus Dashboard image.

a) Browse to the Software Download page.

https://software.cisco.com/download/home/286327743/type/286328258

b) Click **Nexus Dashboard Software**.

c) From the left sidebar, choose the Nexus Dashboard version you want to download.

d) Download the Cisco Nexus Dashboard image for Linux KVM (`nd-dk9.<version>.qcow2`).

**Step 2** Copy the image to the Linux KVM servers where you will host the nodes.

You can use `scp` to copy the image, for example:

```
# scp nd-dk9.<version>.qcow2 root@<kvm-host-ip>:/home/nd-base
```

The following steps assume you copied the image into the `/home/nd-base` directory.

**Step 3** Create the required disk images for the first node.

You will create a snapshot of the base `qcow2` image you downloaded and use the snapshots as the disk images for the nodes' VMs. You will also need to create a second disk image for each node.

a) Log in to your KVM host as the `root` user.

b) Create a directory for the node's snapshot.

The following steps assume you create the snapshot in the `/home/nd-node1` directory.

```
# mkdir -p /home/nd-node1/
# cd /home/nd-node1
```

c) Create the snapshot.

In the following command, replace `/home/nd-base/nd-dk9.<version>.qcow2` with the location of the base image you created in the previous step.

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.<version>.qcow2
/home/nd-node1/nd-node1-disk1.qcow2
```

**Note**
If you are deploying in RHEL 8.6, you may need to provide an additional parameter to define the destination snapshot's format as well. In that case, update the above command to the following:

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.2.1.1a.qcow2 /home/nd-node1/nd-node1-disk1.qcow2
-F qcow2
```

d) Create the additional disk image for the node.

Each node requires two disks: a snapshot of the base Nexus Dashboard `qcow2` image and a second 500GB disk.

```
# qemu-img create -f qcow2 /home/nd-node1/nd-node1-disk2.qcow2 500G
```

**Step 4** Repeat the previous step to create the disk images for the second and third nodes.

Before you proceed to the next step, you should have the following:

- For the first node, `/home/nd-node1/` directory with two disk images:

- `/home/nd-node1/nd-node1-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
- `/home/nd-node1/nd-node1-disk2.qcow2`, which is a new 500GB disk you created.

- For the second node, `/home/nd-node2/` directory with two disk images:
  - `/home/nd-node2/nd-node2-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
  - `/home/nd-node2/nd-node2-disk2.qcow2`, which is a new 500GB disk you created.

- For the third node, `/home/nd-node3/` directory with two disk images:
  - `/home/nd-node1/nd-node3-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
  - `/home/nd-node1/nd-node3-disk2.qcow2`, which is a new 500GB disk you created.

**Step 5** Create the first node's VM.

a) Open the KVM console and click **New Virtual Machine**.

You can open the KVM console from the command line using the `virt-manager` command.

If your Linux KVM environment does not have a desktop GUI, run the following command instead and proceed to step 6.

```
virt-install --import --name <node-name> --memory 65536 --vcpus 16 --os-type generic --disk
path=/path/to/disk1/nd-node1-d1.qcow2,format=qcow2,bus=virtio --disk
path=/path/to/disk2/nd-node1-d2.qcow2,format=qcow2,bus=virtio --network
bridge=<mgmt-bridge-name>,model=virtio --network bridge=<data-bridge-name>,model=virtio --console
pty,target_type=serial --noautoconsole --autostart
```

b) In the **New VM** screen, choose **Import existing disk image option** and click **Forward**.

c) In the **Provide existing storage path** field, click **Browse** and select the `nd-node1-disk1.qcow2` file.

We recommend that each node's disk image is stored on its own disk partition.

d) Choose `Generic` for the **OS type** and **Version**, then click **Forward**.

e) Specify 64GB memory and 16 CPUs, then click **Forward**.

f) Enter the **Name** of the virtual machine, for example `nd-node1` and check the **Customize configuration before install** option. Then click **Finish**.

**Note**
You must select the **Customize configuration before install** checkbox to be able to make the disk and network card customizations required for the node.

The VM details window will open.

In the VM details window, change the NIC's device model:

a) Select **NIC <mac>**.

b) For **Device model**, choose `e1000`.

c) For **Network Source**, choose the bridge device and provide the name of the "mgmt" bridge.

**Note**

Creating bridge devices is outside the scope of this guide and depends on the distribution and version of the operating system. Consult the operating system's documentation, such as Red Hat's *Configuring a network bridge*, for more information.

In the VM details window, add a second NIC:

a) Click **Add Hardware**.
b) In the **Add New Virtual Hardware** screen, select **Network**.
c) For **Network Source**, choose the bridge device and provide the name of the created "data" bridge.
d) Leave the default **Mac address** value.
e) For **Device model**, choose `e1000`.

In the VM details window, add the second disk image:

a) Click **Add Hardware**.
b) In the **Add New Virtual Hardware** screen, select **Storage**.
c) For the disk's bus driver, choose `IDE`.
d) Select **Select or create custom storage**, click **Manage**, and select the `nd-node1-disk2.qcow2` file you created.
e) Click **Finish** to add the second disk.

**Note**
Ensure that you enable the `Copy host CPU configuration` option in the Virtual Machine Manager UI.

Finally, click **Begin Installation** to finish creating the node's VM.

**Step 6** Repeat previous steps to deploy the second and third nodes, then start all VMs.

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

**Step 7** Open one of the node's console and configure the node's basic information. If your Linux KVM environment does not have a desktop GUI, run the `virsh console <node-name>` command to access the console of the node.

a) Press any key to begin initial setup.

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` SSH login as well as the initial GUI password.

**Note**
You must provide the same password for all nodes or the cluster creation will fail.

```
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
```

d) For the first node only, designate it as the "Cluster Leader".

You will log into the cluster leader node to finish configuration and complete cluster creation.

```
Is this the cluster leader?: y
```

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, 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
Cluster leader: yes

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

**Step 8** Repeat previous step to configure the initial information for the second and third nodes.

You do not need to wait for the first node configuration to complete, you can begin configuring the other two nodes simultaneously.

**Note**
You must provide the same password for all nodes or the cluster creation will fail.

The steps to deploy the second and third nodes are identical with the only exception being that you must indicate that they are not the **Cluster Leader**.

**Step 9** Wait for the initial bootstrap process to complete on all nodes.

After you provide and confirm management network information, the initial setup on the first node (`Cluster Leader`) configures the networking and brings up the UI, which you will use to add two other nodes and complete the cluster deployment.

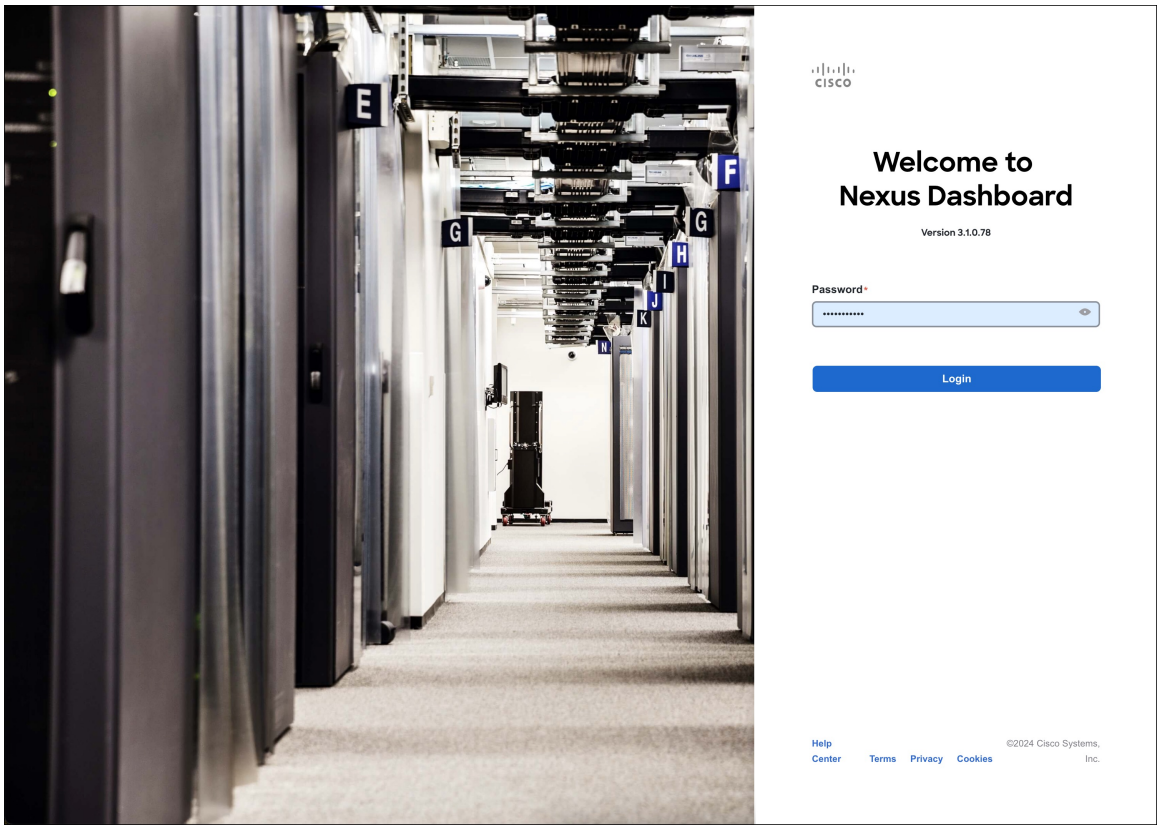
```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

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

**Step 10** Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

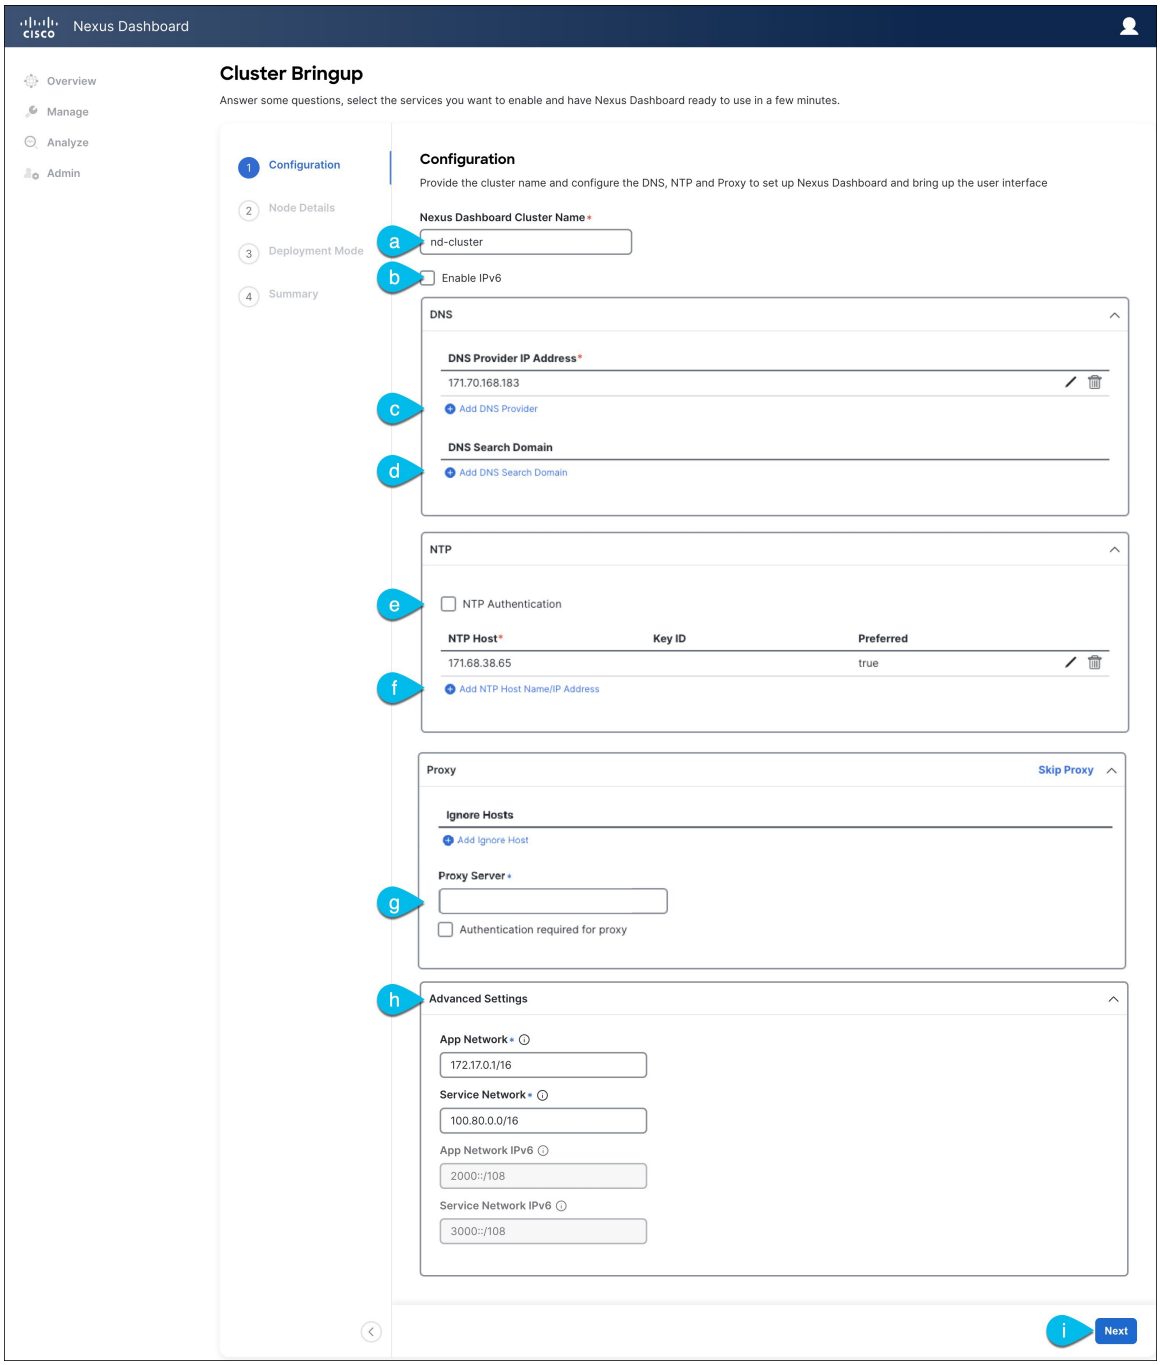
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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




**Step 11** 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) 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.

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

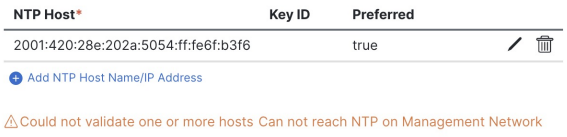
In the additional fields, provide the following information:

- **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.

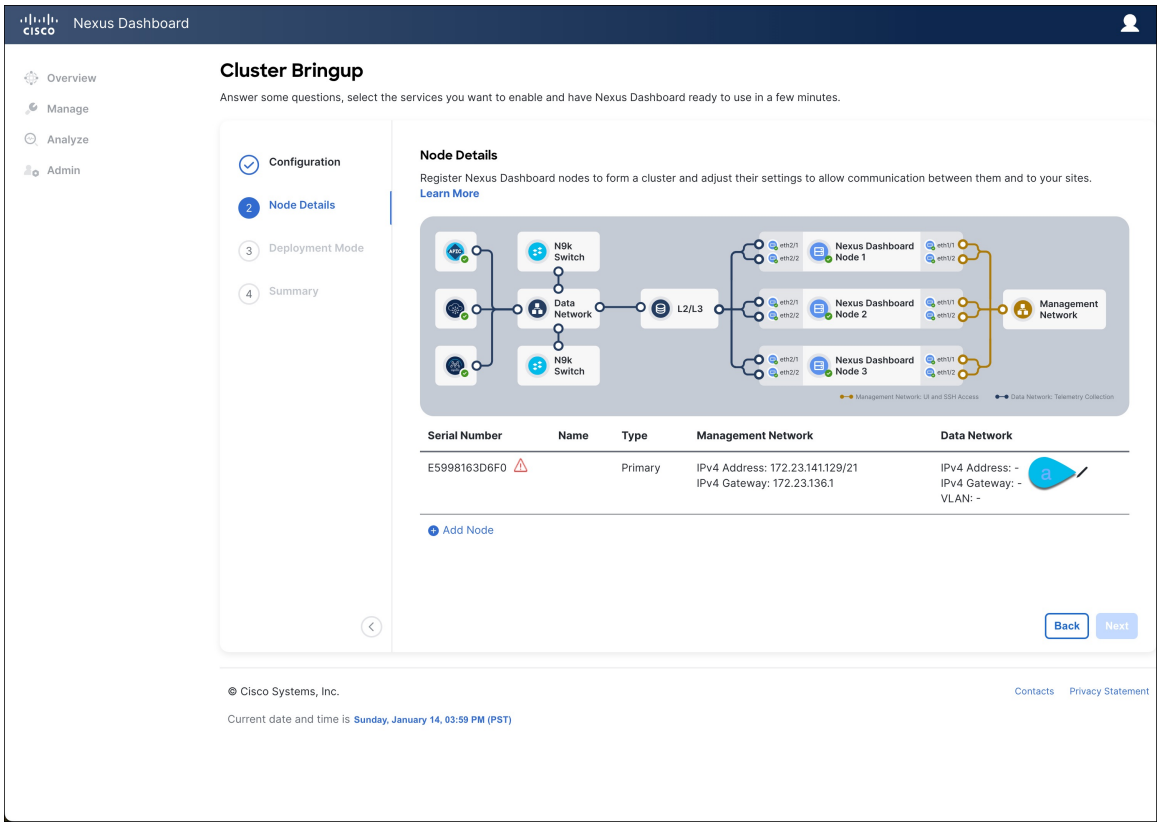
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

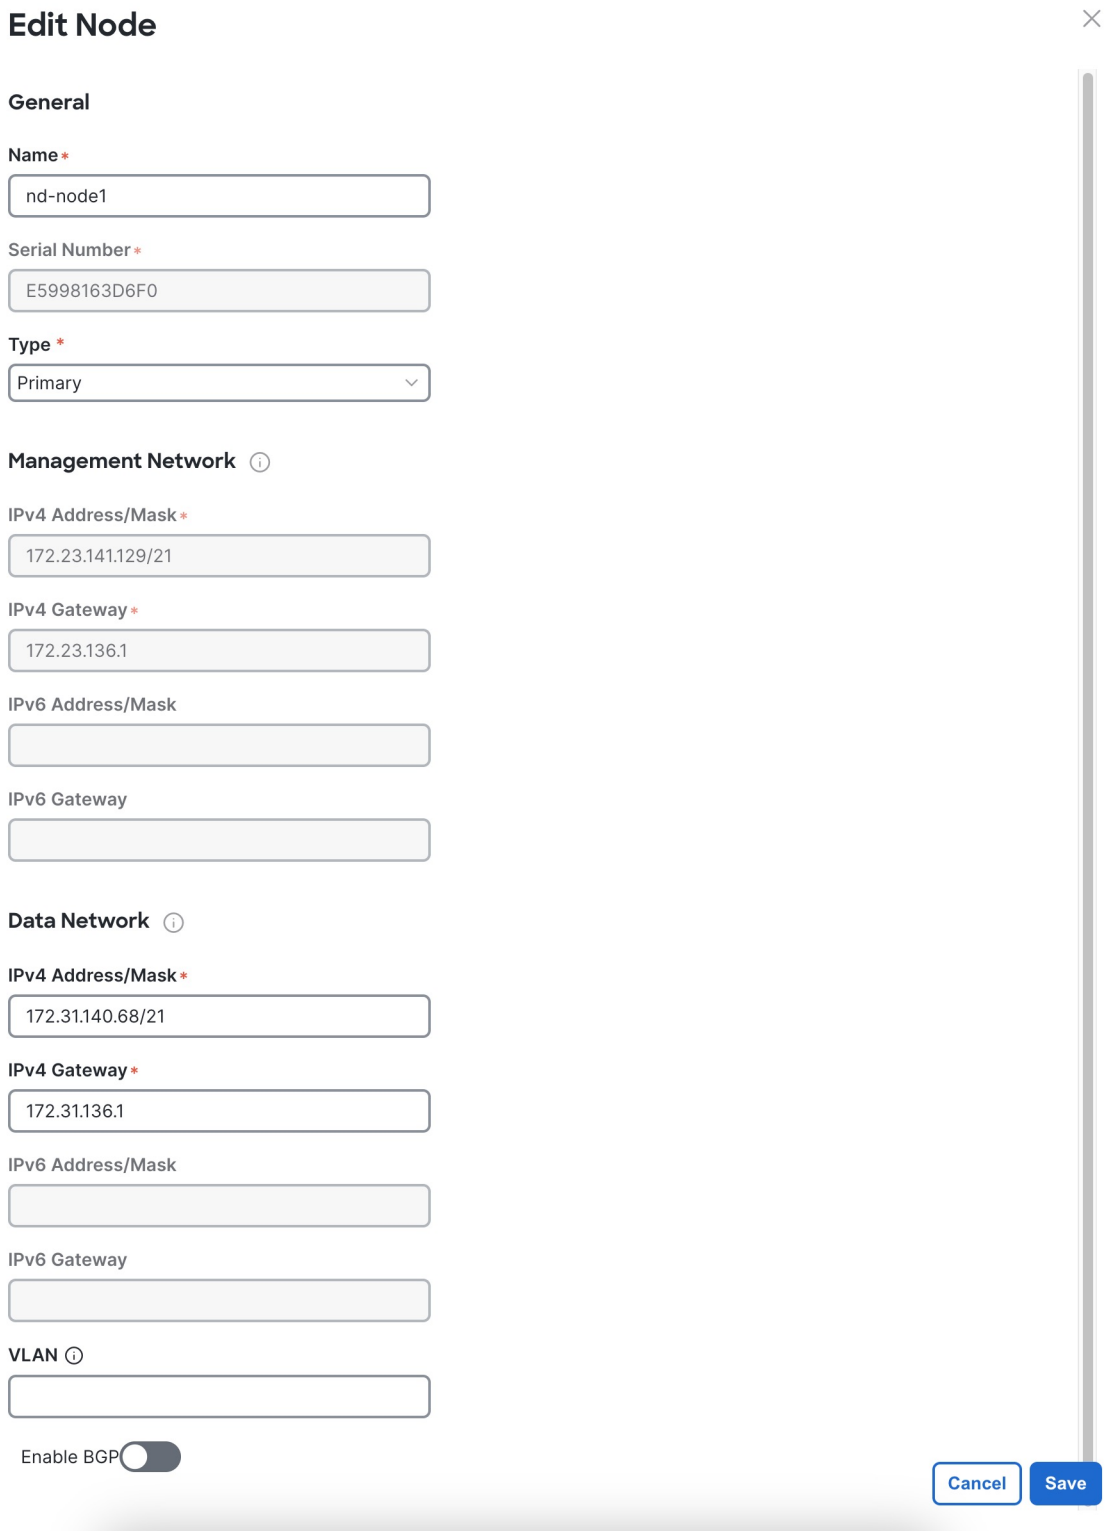
  Application and Services networks are described in the Prerequisites and Guidelines section earlier in this document.

j) Click **Next** to continue.

**Step 12** 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 sites.
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 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.

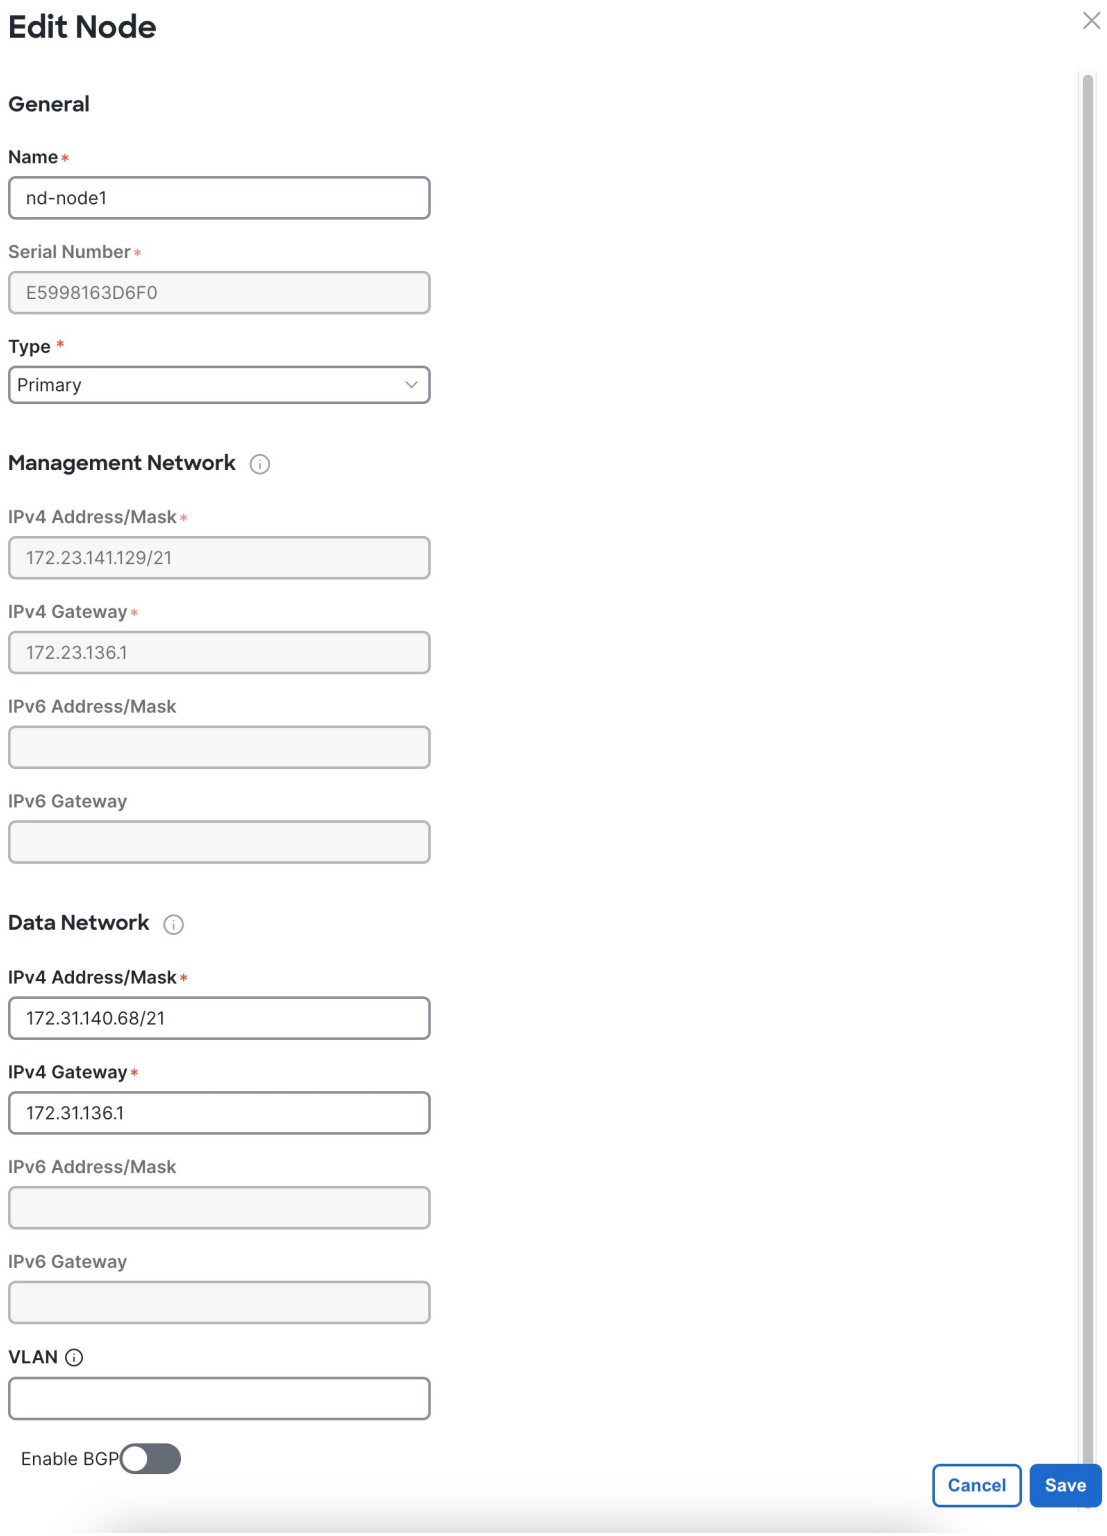
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

**Step 13** 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 **Management IP Address** and **Password** for the second node

You defined the management network information and the password during the initial node configuration steps.

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

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

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

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 **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.

f) (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 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.

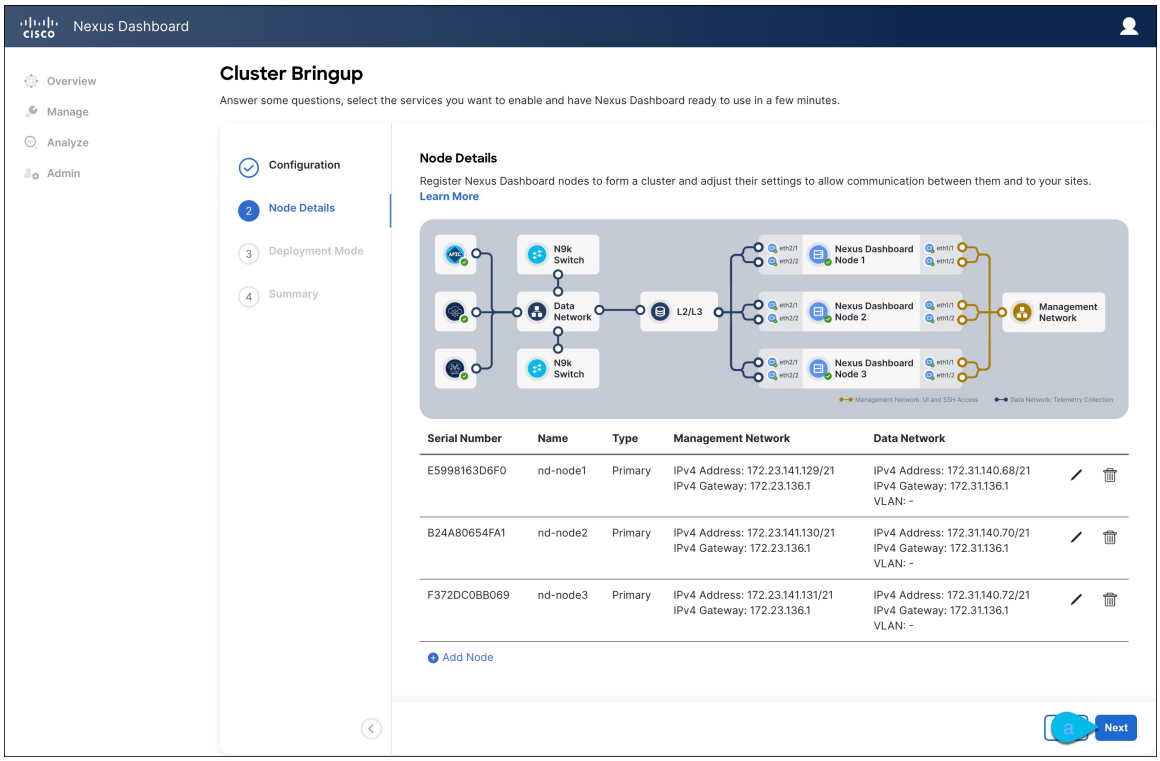
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

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

**Step 14** In the **Node Details** page, verify the provided information and click **Next** to continue.




**Step 15** Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

**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.

b) Click **Add Persistent Service IPs/Pools** to provide one or more persistent IPs required by Insights or Fabric Controller services.

For more information about persistent IPs, see the Prerequisites and Guidelines section.

c) Click **Next** to proceed.

**Step 16** In the **Summary** screen, review and verify the configuration information and click **Save** to build the cluster.

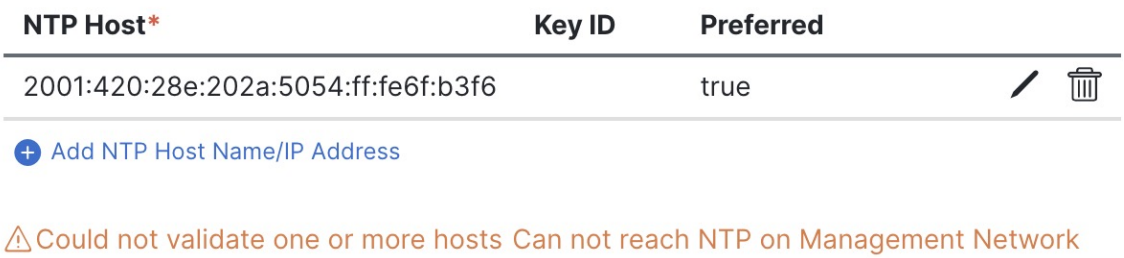
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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.

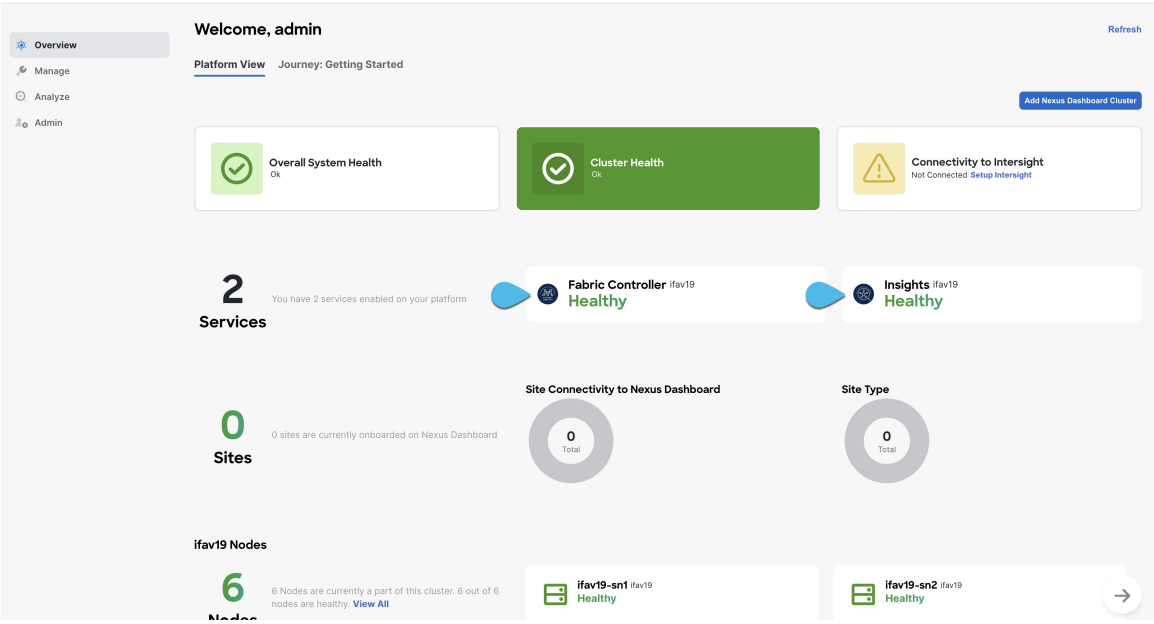
**Step 17** Verify that the cluster is healthy.

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

After 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 18** 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.