



Deploying in Linux KVM (Release 3.2.2 and Later)

- [Prerequisites and guidelines for deploying the Nexus Dashboard cluster in Linux KVM, on page 1](#)
- [Deploying Nexus Dashboard in Linux KVM, on page 3](#)

Prerequisites and guidelines for deploying the Nexus Dashboard cluster in Linux KVM

Before you proceed with deploying the Nexus Dashboard cluster in a Linux KVM, the KVM must meet these prerequisites and you must follow these guidelines:

- The following prerequisites and guidelines apply only for Nexus Dashboard release 3.2.2:
 - Only CLI-based procedures are supported when creating the VMs for the nodes.
- The KVM host must be running one of the following supported Linux Distributions:
 - Red Hat Enterprise Linux (RHEL) 8.8 with KVM
 - Red Hat Enterprise Linux (RHEL) 8.10 with KVM
- Verify that you have already installed KVM, `qemu`, `virt-manager`, and `libvirt` on your RHEL host OS.
- Verify that `virt-install` is enabled for `virt-manager` deployments.
- The KVM form factor must support 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.
- To deploy Nexus Dashboard on Linux KVM, verify that KVM hypervisor meets the following requirements:
 - Requires CPU virtualization in order to bring up Nexus Dashboard virtual machines. Your host should have:

- An Intel processor with the Intel VT-x and Intel 64 virtualization extensions for x86-based systems, or
- An AMD processor with the AMD-V and the AMD64 virtualization extensions.

You can check for virtualization extensions by using following commands to determine whether your system has the hardware virtualization extensions, and that they are enabled:

```
grep -E 'svm|vmx' /proc/cpuinfo
```

- The KVM hosts must also meet the following compute resources requirements to host Nexus Dashboard virtual machines:
 - CPU: 16 vCPU
 - Memory: 64 GB of RAM
 - Storage: 550 GB of total disk storage.

The disk must have I/O latency of 20ms or less, so you should use SSD or NVME drives to allocate the storage. See [Verify the I/O latency of a Linux KVM storage device](#).
- Virtual machine deployments should also meet the following placement requirements if you are deploying a highly available cluster (for example, at least three virtual machines):
 - Virtual machines should not share same SSD.
 - We recommend that each Nexus Dashboard node is deployed in a different KVM hypervisor and each host has its own dedicated SSD disk.
- You must also configure the following required network bridges at the host level for Nexus Dashboard deployments:
 - Management Network Bridge (mgmt-bridge): The external network to manage Nexus Dashboard.
 - Data Network Bridge (data-bridge): The internal network used to form clustering within Nexus Dashboard.

Verify the I/O latency of a Linux KVM storage device

When you deploy a Nexus Dashboard cluster in a Linux KVM, the storage device of the KVM must have a latency under 20ms.

Follow these steps to verify the I/O latency of a Linux KVM storage device.

Procedure

-
- Step 1** Create a test directory.
- For example, create a directory named `test-data`.
- Step 2** Run the Flexible I/O tester (FIO).

```
# fio --rw=write --ioengine=sync --fdatasync=1 --directory=test-data --size=22m --bs=2300 --name=mytest
```

- Step 3** After you use the command, confirm that the `99.00th=[value]` in the `fsync/fdatasync/sync_file_range` section is under 20ms.

Deploying Nexus Dashboard in Linux KVM

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



Note KVM-based deployments are supported for Nexus Dashboard Fabric Controller services only.

Before you begin

- Ensure that you meet the requirements and guidelines described in [Prerequisites and guidelines for deploying the Nexus Dashboard cluster in Linux KVM](#).

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 `nd-dk9.3.2.1i.qcow2` into the `/home/nd-base` directory on each KVM host.

- Step 3** Make the following configurations on each KVM host:

a) Edit `/etc/libvirt/qemu.conf` and make sure the user and group is correctly configured based on the ownership of the storage that you plan to use for the Nexus Dashboard deployment.

This is only required if you plan to use disk storage paths that are different from the default `libvirtd`.

b) Edit `/etc/libvirt/libvirt.conf` and uncomment `uri_default`.

c) Restart the `libvirtd` service after updating the configuration using the `systemctl restart libvirtd` command from root.

- Step 4** Create the required disk images on each node.

As mentioned earlier, you will need a total of 550GB of SSD storage to create two disk images:

- Boot disk based on QCOW2 image that you downloaded
- Data disk of 500GB

Step 5

Log in to your KVM host as the `root` user and perform the following steps on each node.

- Mount the storage disk (raw disk or LVM) to directory `/opt/cisco/nd`.
- Create the following script as `/root/create_vm.sh` under the root directory:

Note

If you manually type this information, verify that there are no empty spaces present after any of these lines.

```
#!/bin/bash -ex

# Configuration
# Name of Nexus Dashboard Virtual machine
name=ndl

# Path of Nexus Dashboard QCOW2 image.
nd_qcow2=/home/nd-base/nd-dk9.3.2.1i.qcow2

# Disk Path to storage Boot and Data Disks.
data_disk=/opt/cisco/nd/data

# Management Network Bridge
mgmt_bridge=mgmt-bridge

# Data Network bridge
data_bridge=data-bridge

# Data Disk Size
data_size=500G

# CPU Cores
cpus=16

# Memory in units of MB.
memory=65536

# actual script
rm -rf $data_disk/boot.img
/usr/bin/qemu-img convert -f qcow2 -O raw $nd_qcow2 $data_disk/boot.img
rm -rf $data_disk/disk.img
/usr/bin/qemu-img create -f raw $data_disk/disk.img $data_size
virt-install \
--import \
--name $name \
--memory $memory \
--vcpus $cpus \
--os-type generic \
--osinfo detect=on,require=off \
--check path_in_use=off \
--disk path=${data_disk}/boot.img,format=raw,bus=virtio \
--disk path=${data_disk}/disk.img,format=raw,bus=virtio \
--network bridge=$mgmt_bridge,model=virtio \
--network bridge=$data_bridge,model=virtio \
--console pty,target_type=serial \
--noautoconsole \
--autostart
```

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:

Cluster Bringup

Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.

1 Configuration

Configuration
Provide the cluster name and configure the DNS, NTP and Proxy to set up Nexus Dashboard and bring up the user interface

Nexus Dashboard Cluster Name *
nd-cluster

☐ Enable IPv6

DNS

DNS Provider IP Address *
171.70.168.183

+ Add DNS Provider

DNS Search Domain
+ Add DNS Search Domain

NTP

☐ NTP Authentication

NTP Host *	Key ID	Preferred
171.68.38.65		true

+ Add NTP Host Name/IP Address

Proxy [Skip Proxy](#)

Ignore Hosts
+ Add Ignore Host

Proxy Server *

☐ Authentication required for proxy

Advanced Settings

App Network *
172.17.0.1/16

Service Network *
100.80.0.0/16

App Network IPv6
2000::/108

Service Network IPv6
3000::/108

[Next](#)

- 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:

NTP Host*	Key ID	Preferred	
2001:420:28e:202a:5054:ff:fe6f:b3f6		true	 
 Add NTP Host Name/IP Address			

 Could not validate one or more hosts Can not reach NTP on Management Network

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 12

Provide the **Cluster Details**.

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

Cluster Bringup

Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.

1 Configuration

Configuration

Provide the cluster name and configure the DNS, NTP and Proxy to set up Nexus Dashboard and bring up the user interface

Nexus Dashboard Cluster Name *

nd-cluster

Enable IPv6

DNS

DNS Provider IP Address *

171.70.168.183

DNS Search Domain

NTP

NTP Authentication

NTP Host * **Key ID** **Preferred**

171.68.38.65 true

Proxy

Ignore Hosts

Proxy Server *

Advanced Settings

App Network *

172.17.0.1/16

Service Network *

100.80.0.0/16

App Network IPv6

2000::/108

Service Network IPv6

3000::/108

Next

- 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:

NTP Host*	Key ID	Preferred	
2001:420:28e:202a:5054:ff:fe6f:b3f6		true	 

 Add NTP Host Name/IP Address

 Could not validate one or more hosts Can not reach NTP on Management Network

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 13

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.

Overview

Manage

Analyze

Admin

Nexus Dashboard

Cluster Bringup

Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.

Configuration

Node Details

Deployment Mode

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](#)

Management Network: UI and SSH Access

Data Network: Telemetry Collection

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.

Current date and time is Sunday, January 14, 03:59 PM (PST)

Contacts

Privacy Statement

Edit Node

General

Name *

nd-node1

Serial Number *

E5998163D6F0

Type *

Primary

Management Network ⓘ

IPv4 Address/Mask *

172.23.141.129/21

IPv4 Gateway *

172.23.136.1

IPv6 Address/Mask

IPv6 Gateway

Data Network ⓘ

IPv4 Address/Mask *

172.31.140.68/21

IPv4 Gateway *

172.31.136.1

IPv6 Address/Mask

IPv6 Gateway

VLAN ⓘ

Enable BGP ☐

Cancel

Save

- 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 14

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.

Edit Node

General

Name *

Serial Number *

Type *

Management Network ⓘ

IPv4 Address/Mask *

IPv4 Gateway *

IPv6 Address/Mask

IPv6 Gateway

Data Network ⓘ

IPv4 Address/Mask *

IPv4 Gateway *

IPv6 Address/Mask

IPv6 Gateway

VLAN ⓘ

Enable BGP ☐

Cancel

Save

- 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 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.

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

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

Step 15

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

Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.

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](#)

The diagram illustrates a network topology where three Nexus Dashboard nodes (nd-node1, nd-node2, nd-node3) are connected to a central L2/L3 switch. This switch is connected to two N9k switches, which are in turn connected to a Data Network. The nodes are also connected to a Management Network. The table below lists the details for each node.

Serial Number	Name	Type	Management Network	Data Network
E5998163D6F0	nd-node1	Primary	IPv4 Address: 172.23.141.129/21 IPv4 Gateway: 172.23.136.1	IPv4 Address: 172.31.140.68/21 IPv4 Gateway: 172.31.136.1 VLAN: -
B24A80654FA1	nd-node2	Primary	IPv4 Address: 172.23.141.130/21 IPv4 Gateway: 172.23.136.1	IPv4 Address: 172.31.140.70/21 IPv4 Gateway: 172.31.136.1 VLAN: -
F372DC08B069	nd-node3	Primary	IPv4 Address: 172.23.141.131/21 IPv4 Gateway: 172.23.136.1	IPv4 Address: 172.31.140.72/21 IPv4 Gateway: 172.31.136.1 VLAN: -

[Add Node](#) [Next](#)

Step 16 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 for all enabled services](#) section.

c) Click **Next** to proceed.

Step 17 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.

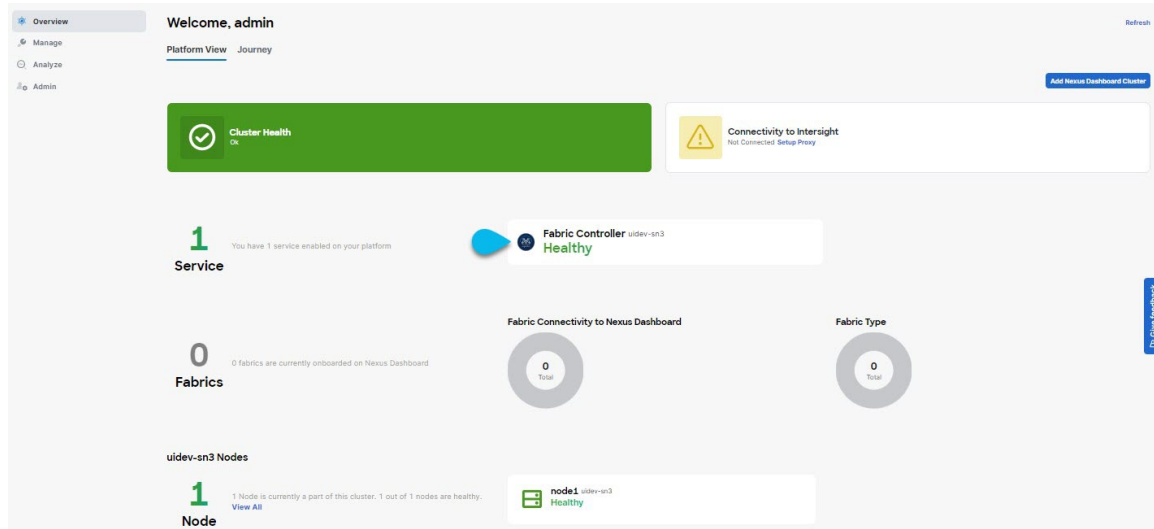
It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 18 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 19

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](#).
-