



Deploying in Linux KVM

- [Prerequisites and Guidelines, on page 1](#)
- [Deploying Nexus Dashboard in Linux KVM, on page 4](#)

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 [Deployment Overview and Requirements](#).
- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- Ensure you have enough system resources:

Table 1: Deployment Requirements

Orchestrator Version	Requirements
Release 2.3.x	

Orchestrator Version	Requirements
	<ul style="list-style-type: none"> • KVM deployments are supported for Nexus Dashboard Fabric Controller and Nexus Dashboard Orchestrator services only. Specific versions of the required OS and libraries for each service are listed below. • For Nexus Dashboard Fabric Controller: <ul style="list-style-type: none"> • You must deploy in CentOS 7.9 or Red Hat Enterprise Linux 8.6 • You must have the supported versions of Kernel and KVM: <ul style="list-style-type: none"> • For CentOS 7.9, Kernel version <code>3.10.0-957.el7.x86_64</code> and KVM version <code>libvirt-4.5.0-23.el7_7.1.x86_64</code> • For RHEL 8.6, Kernel version <code>4.18.0-372.9.1.el8.x86_64</code> and KVM version <code>libvirt 8.0.0</code> • For Nexus Dashboard Fabric Orchestrator: <ul style="list-style-type: none"> • You must deploy in CentOS 7.7 • You must have the supported versions of Kernel and KVM: <ul style="list-style-type: none"> • Kernel <code>3.10.0-1062.el7.x86_64</code> • KVM <code>libvirt 4.5.0</code> • 16 vCPUs • 64 GB of RAM • 550 GB disk Each node requires a dedicated disk partition • The disk must have I/O latency of 20ms or less. To verify the I/O latency: <ol style="list-style-type: none"> 1. Create a test directory. For example, <code>test-data</code>. 2. Run the following command: <pre># fio --rw=write --ioengine=sync --fdatasync=1 --directory=test-data --size=22m --bs=2300 --name=mytest</pre>

Orchestrator Version	Requirements
	<p>3. After the command is executed, confirm that the <code>99.00th=[<value>]</code> in the <code>fsync/fdatasync/sync_file_range</code> section is under 20ms.</p> <ul style="list-style-type: none"> 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](#).

Step 1

Download the Cisco Nexus Dashboard image.

- Browse to the Software Download page.

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

- Click **Nexus Dashboard Software**.
- From the left sidebar, choose the Nexus Dashboard version you want to download.
- 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.

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

- 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-node3/nd-node3-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
 - `/home/nd-node3/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.

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:

- Select **NIC <mac>**.
- For **Device model**, choose `e1000`.
- For **Network Source**, choose the bridge device and provide the name of the "mgmt" bridge.

In the VM details window, add a second NIC:

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

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

- Click **Add Hardware**.
- In the **Add New Virtual Hardware** screen, select **Storage**.
- Select **Select or create custom storage**, click **Manage**, and select the `nd-node1-disk2.qcow2` file you created.
- Click **Finish** to add the second disk.

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

Step 6

Repeat the previous step to create the VMs for the second and third nodes, then start all VMs.

Step 7

Open one of the node's console and configure the node's basic information.

- 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...

- Enter and confirm the `admin` password

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

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

- Enter the management network information.

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

- 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 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 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.

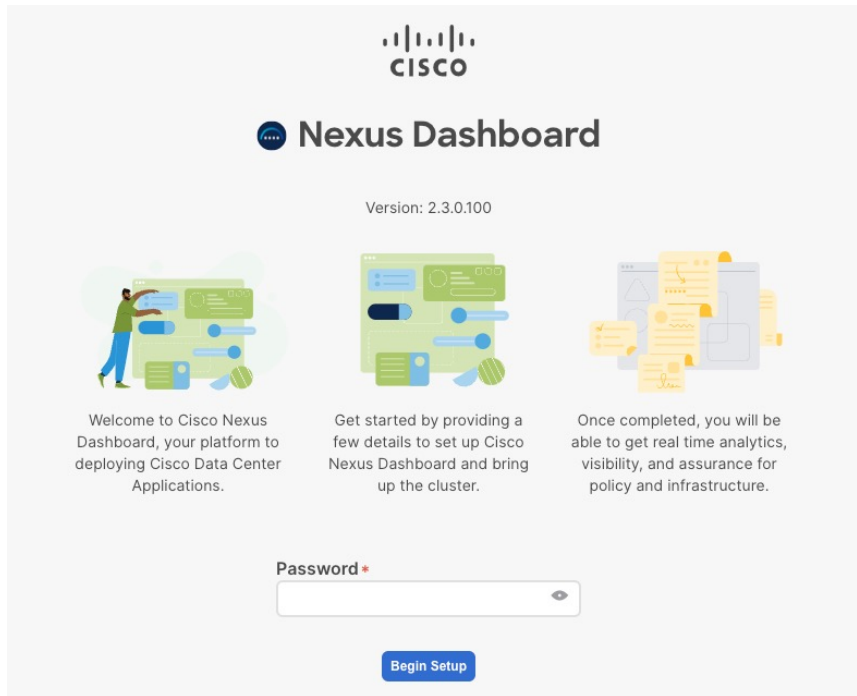
```
Please wait for system to boot: [#####] 100%
System up, please wait for UI to be online.
```

```
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 **Begin Setup**



Step 11 Provide the **Cluster Details**.

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

- a) Provide the **Cluster Name** for this Nexus Dashboard cluster.
- b) Click **+Add NTP Host** to add one or more NTP servers.

You must provide an IP address, fully qualified domain name (FQDN) are not supported.

After you enter the IP address, click the green checkmark icon to save it.

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

After you enter the IP address, click the green checkmark icon to save it.

- d) 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, which will allow you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

If you want to skip proxy configuration, click the information (i) icon next to the field, then click **Skip**.

- e) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to **Yes** and provide the login credentials.
- f) (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 enter the IP address, click the green 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.

Application and Services networks are described in the [Prerequisites and Guidelines](#) section earlier in this document.

- g) Click **Next** to continue.

Step 12

In the **Node Details** screen, provide the node's information.

- a) Click the **Edit** button next to the first node.
b) In the **Password** field, enter the password for this node and click **Validate**.

This will auto-populate the **Serial Number** and **Management Network** information for the node.

- c) Provide the node's **Name**.
d) Provide the node's **Data Network** information.

The **Management Network** information is already pre-populated with the information you provided for the first node.

You must provide the data network IP address/netmask (for example, `172.31.140.58/24`) and gateway (for example, `172.31.140.1`). Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

- e) (Optional) Provide IPv6 addresses for the management and data networks.

Nexus Dashboard supports either IPv4 or dual stack IPv4/IPv6 for the management and data networks.

Note If you want to provide IPv6 information, you must do that now during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

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

BGP configuration is required for the Persistent IPs feature required by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in detail in the "Persistent IP Addresses" sections of the *Nexus Dashboard User's 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.
- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

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

Step 13

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

The **Node Details** window opens.

- a) In the **Deployment Details** section, provide the node's **Management IP Address** and **Password** for the `rescue-user` you configured when deploying the node's VM, then click **Verify**.

This will auto-populate the **Serial Number** and **Management Network** information for the node.

- b) Provide the node's **Name**.
- c) Provide the node's **Data Network IP** address and gateway.

The **Management Network** information will be pre-populated with the information pulled from the node based on the management IP address and credentials you provided in the previous sub-step.

You must provide the data network IP address/netmask (for example, `172.31.141.58/24`) and gateway (for example, `172.31.141.1`). Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

- d) (Optional) Provide IPv6 information for the management and data networks.

Nexus Dashboard supports either IPv4 or dual stack IPv4/IPv6 for the management and data networks.

Note If you want to provide IPv6 information, you must do it during cluster bootstrap process. If you deploy the cluster using only IPv4 stack and want to add IPv6 information later, you would need to redeploy the cluster.

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

- e) (Optional) If required, **Enable BGP** for the data network.
- f) Click **Save** to save the changes.

Step 14 Repeat the previous step to add the 3rd node.

Step 15 In the **Node Details** screen, click **Next** to continue.

After you've provided the information for all 3 nodes in the cluster, continue to the next screen of the bootstrap process.

Node Details

Provide the necessary node details to set up Nexus Dashboard and bring up the User Interface.

The diagram shows three sites connected to a central Data Network (L2/L3). Each site is connected to the Data Network via Fabric 0/1. The Data Network is connected to a Management Network (MN) via Mgmt 0/1.

General

Serial Number	Name	Management Network	Data Network	
EA986C528737	node-ova-app1	IPv4/mask: 172.31.140.46/24 IPv4 Gateway: 172.31.140.1 IPv6/mask: - IPv6 Gateway: -	IPv4/mask: 172.31.141.58/24 IPv4 Gateway: 172.31.141.1 IPv6/mask: - IPv6 Gateway: - VLAN: -	✎ 🗑
B734BC2033AD	node-ova-app2	IPv4/mask: 172.31.140.60/24 IPv4 Gateway: 172.31.140.1 IPv6/mask: - IPv6 Gateway: -	IPv4/mask: 172.31.141.68/24 IPv4 Gateway: 172.31.141.1 IPv6/mask: - IPv6 Gateway: - VLAN: -	✎ 🗑
AED5046A16E2	node-ova-app3	IPv4/mask: 172.31.140.70/24 IPv4 Gateway: 172.31.140.1 IPv6/mask: - IPv6 Gateway: -	IPv4/mask: 172.31.141.72/24 IPv4 Gateway: 172.31.141.1 IPv6/mask: - IPv6 Gateway: - VLAN: -	✎ 🗑

Previous Next

Step 16 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 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 all three nodes are ready, you can log in to any one node via SSH 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.
