Deploying in VMware ESX

Prerequisites and guidelines for deploying the Nexus Dashboard cluster in VMware ESX

Before you proceed with deploying the Nexus Dashboard cluster in VMware ESX, you must:

- Ensure that the ESX form factor supports 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** Some services (such as Nexus Dashboard Fabric Controller) may require only a single ESX virtual node for one or more specific use cases. In that case, the capacity planning tool will indicate the requirement and you can simply skip the additional node deployment step in the following sections.

- Review and complete the general prerequisites described in Prerequisites: Nexus Dashboard.

  Note that this document describes how to initially deploy the base Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `secondary` or `standby`), see the "Infrastructure Management" chapter of the *Cisco Nexus Dashboard User Guide* instead, which is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

- 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.

- When deploying in VMware ESX, you can deploy two types of nodes:

  - Data Node—Node profile with higher system requirements designed for specific services that require the additional resources.

• App Node—Node profile with a smaller resource footprint that can be used for most services.

Note Some larger scale Nexus Dashboard Fabric Controller deployments may require additional secondary nodes. If you plan to add secondary nodes to your NDFC cluster, you can deploy all nodes (the initial 3-node cluster and the additional secondary nodes) using the OVA-App profile. Detailed scale information is available in the *Verified Scalability Guide for Cisco Nexus Dashboard Fabric Controller* for your release.

Ensure you have enough system resources:

Table 1: Deployment requirements

Data node requirements	App node requirements
• VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2	• VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2
• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2 if deploying using VMware vCenter	• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2 if deploying using VMware vCenter
• Each VM requires the following:	• Each VM requires the following:
• 32 vCPUs with physical CPU reservation of at least 12,000 MHz	• 16 vCPUs with physical CPU reservation of at least 12,000 MHz
• 128GB of RAM with physical reservation	• 64GB of RAM with physical reservation
• 3TB SSD storage for the data volume and an additional 50GB for the system volume	• 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume
`Data` nodes must be deployed on storage with the following minimum performance requirements:	Some services require `App` nodes to be deployed on faster SSD storage while other services support HDD. Check the [Nexus Dashboard Capacity Planning](#) tool to ensure that you use the correct type of storage.
• The SSD must be attached to the data store directly or in JBOD mode if using a RAID Host Bus Adapter (HBA)	**Note** Beginning with Nexus Dashboard release 3.0(1i) and Nexus Dashboard Insights release 6.3(1), you can use the OVA-App node profile for the Insights service. However, you must change from the default 500GB disk requirement to 1536GB when deploying node VMs which will be used for hosting Insights.
• The SSDs must be optimized for `Mixed Use/Application` (not `Read-Optimized`)	
• 4K Random Read IOPS: `93000`	
• 4K Random Write IOPS: `31000`	
• We recommend that each Nexus Dashboard node is deployed in a different ESXi server.	• We recommend that each Nexus Dashboard node is deployed in a different ESXi server.

• If you plan to configure VLAN ID for the cluster nodes' data interfaces, you must enable VLAN 4095 on the data interface port group in VMware vCenter for Virtual Guest VLAN Tagging (VGT) mode.

If you specify a VLAN ID for Nexus Dashboard data interfaces, the packets must carry a Dot1q tag with that VLAN ID. When you set an explicit VLAN tag in a port group in the vSwitch and attach it to a Nexus Dashboard VM's VNIC, the vSwitch removes the Dot1q tag from the packet coming from the uplink before it sends the packet to that VNIC. Because the vND node expects the Dot1q tag, you must enable VLAN 4095 on the data interface port group to allow all VLANs.

- After each node's VM is deployed, ensure that the VMware Tools' periodic time synchronization is disabled as described in the deployment procedure in the next section.

- VMware vMotion is not supported for Nexus Dashboard cluster nodes.

- VMware Distributed Resource Scheduler (DRS) is not supported for Nexus Dashboard cluster nodes.

  If you have DRS enabled at the ESXi cluster level, you must explicitly disable it for the Nexus Dashboard VMs during deployment as described in the following section.

- Deploying via content library is not supported.

- VMware snapshots are only supported for Nexus Dashboard VMs that are powered off and must be done for all Nexus Dashboard VMs belonging to the same cluster.

  Snapshots of powered on VMs are not supported.

- You can choose to deploy the nodes directly in ESXi or using VMware vCenter.

  If you want to deploy using VMware vCenter, following the steps described in Deploying Nexus Dashboard Using VMware vCenter, on page 4.

  If you want to deploy directly in ESXi, following the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 23.

> **Note**
> If you plan to deploy Nexus Dashboard Insights using the OVA-App node profile, you must deploy using VMware vCenter.
>
> Nexus Dashboard Insights requires a larger disk size than the default value for OVA-App node profiles. If you plan to deploy NDI using the OVA-App node profile, you must change the default disk size for OVA-App nodes from 500GB to 1.5TB during VM deployment. Disk size customization is supported when deploying through VMware vCenter only. For detailed Insights requirements, see the Nexus Dashboard Capacity Planning tool.

Deploying Nexus Dashboard Using VMware vCenter

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter. If you prefer to deploy directly in ESXi, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 23 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and guidelines for deploying the Nexus Dashboard cluster in VMware ESX, on page 1.

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

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

b) Choose the Nexus Dashboard release version you want to download.

c) Click the **Download** icon next to the Nexus Dashboard OVA image (`nd-dk9.<version>.ova`).

Step 2   Log in to your VMware vCenter.

Depending on the version of your vSphere client, the location and order of configuration screens may differ slightly. The following steps provide deployment details using VMware vSphere Client 7.0.
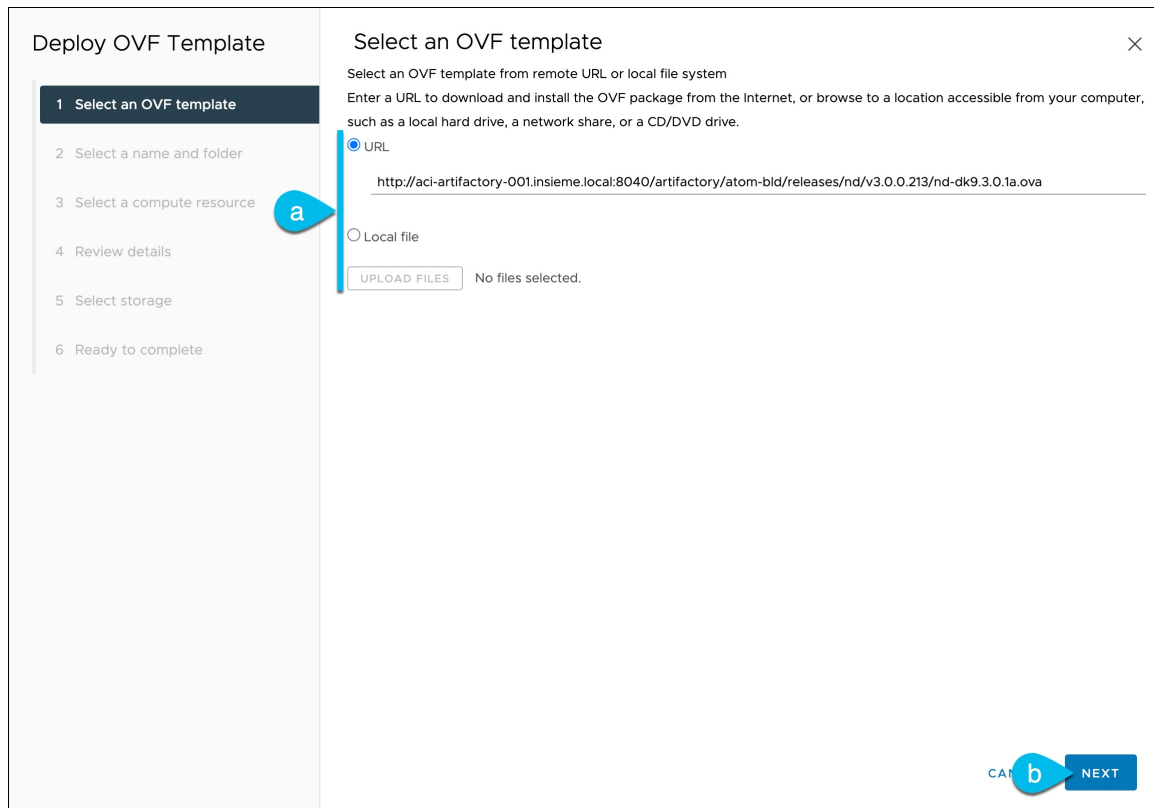
Step 3   Start the new VM deployment.
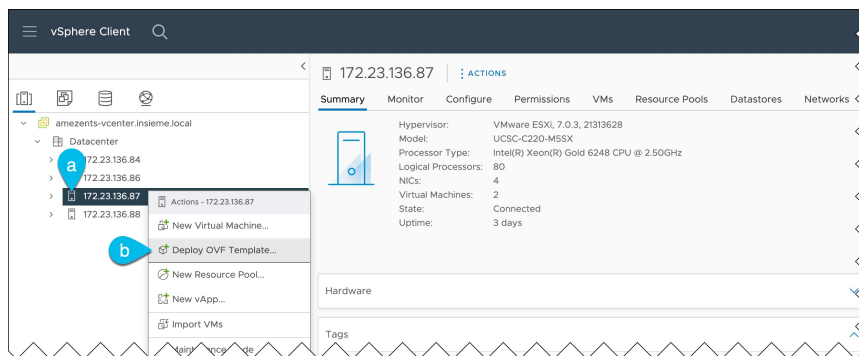
a) Right-click the ESX host where you want to deploy the VM.

b) Select **Deploy OVF Template...**

The **Deploy OVF Template** wizard appears.

Step 4   In the **Select an OVF template** screen, provide the OVA image.

a) Provide the location of the image.

If you hosted the image on a web server in your environment, select **URL** and provide the URL to the image as shown in the above screenshot.

If your image is local, select **Local file** and click **Choose Files** to select the OVA file you downloaded.

b) Click **Next** to continue.

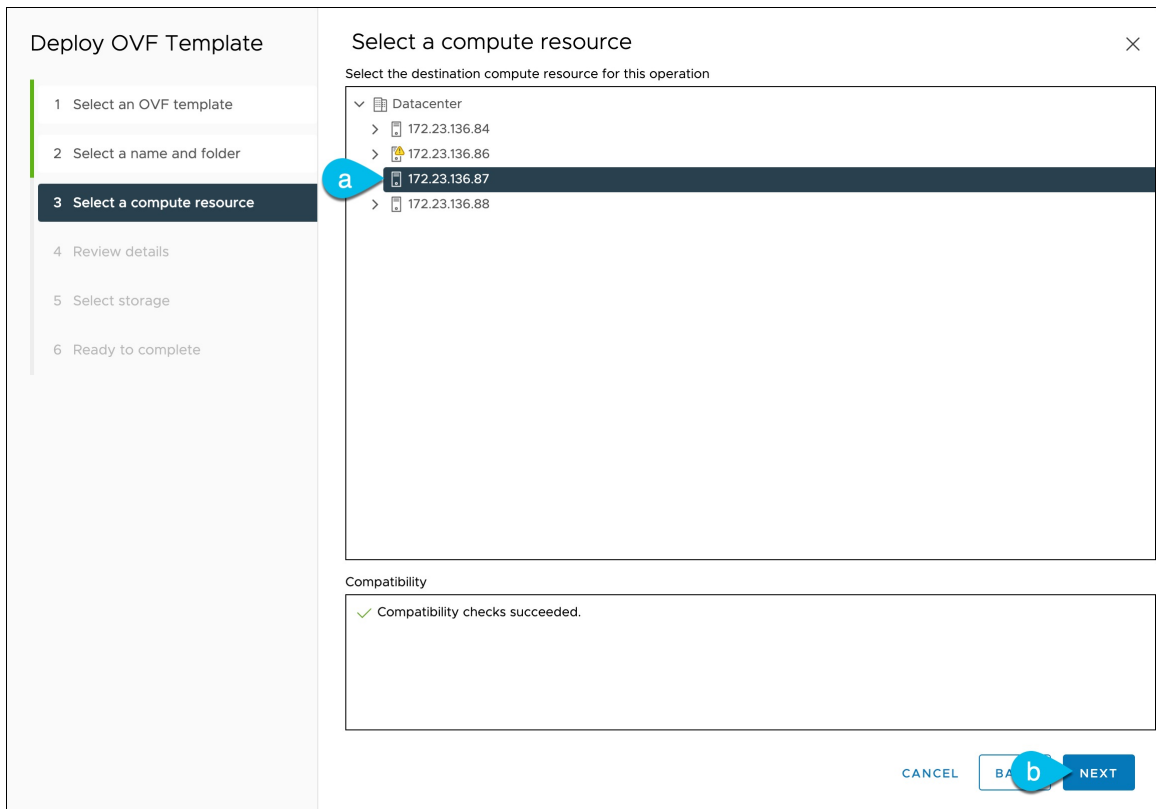**Step 5** In the **Select a name and folder** screen, provide a name and location for the VM.

a) Provide the name for the virtual machine.

For example, `nd-ova-node1`.

b) Select the location for the virtual machine.
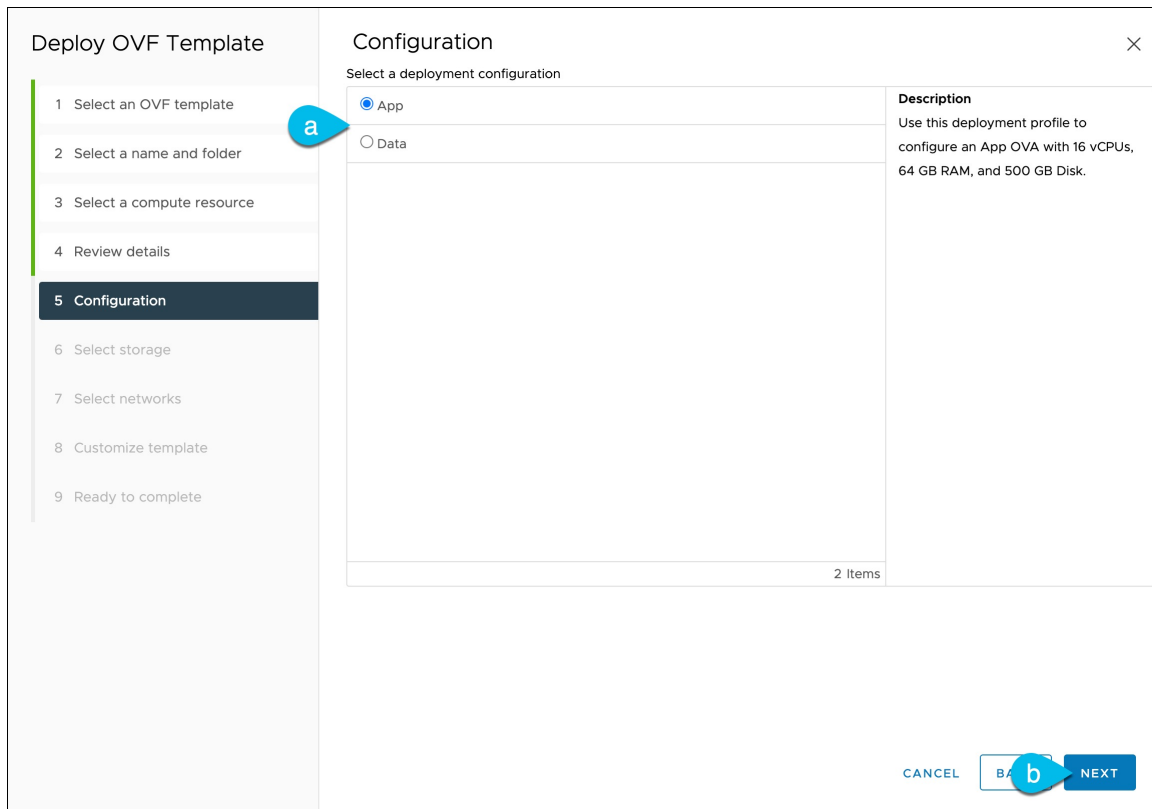c) Click **Next** to continue

Step 6 In the **Select a compute resource** screen, select the ESX host.

a) Select the vCenter data center and the ESX host for the virtual machine.

b) Click **Next** to continue

Step 7 In the **Review details** screen, click **Next** to continue.

Step 8 In the **Configuration** screen, select the node profile you want to deploy.

a) Select either `App` or `Data` node profile based on your use case requirements.

For more information about the node profiles, see Prerequisites and guidelines for deploying the Nexus Dashboard cluster in VMware ESX, on page 1.

b) Click **Next** to continue

Step 9   In the **Select storage** screen, provide the storage information.

a) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

b) Check the **Disable Storage DRS for this virtual machine** checkbox.

Nexus Dashboard does not support VMware DRS.

c) From the **Select virtual disk format** drop-down, choose `Thick Provisioning Lazy Zeroed`.

d) Click **Next** to continue.

Step 10   In the **Select networks** screen, choose the VM network for the Nexus Dashboard's Management and Data networks and click **Next** to continue.

There are two networks required by the Nexus Dashboard cluster:

   • **fabric0** is used for the Nexus Dashboard cluster's Data Network

   • **mgmt0** is used for the Nexus Dashboard cluster's Management Network.

For more information about these networks, see Prerequisites and guidelines for all enabled services in the "Deployment Overview and Requirements" chapter.

Step 11   In the **Customize template** screen, provide the required information.

a) Provide the size for the node's data volume.

The default values will be pre-populated based on the type of node you are deploying, with App node having a single 500GB disk and Data node having a single 3TB disk. In addition to the data volume, a second 50GB system volume will also be configured but cannot be customized.

Note   If you want to specify a custom disk size for your node, you must do so during VM deployment. Resizing the disk after the node is brought up is not supported by Nexus Dashboard.

If you plan to deploy Nexus Dashboard Insights using the OVA-App node profile, you must change the data disk size from the default 500GB value to 1536GB. For additional information about cluster sizing, system resource requirements, and node profile support, see the Nexus Dashboard Capacity Planning.

b) Provide and confirm the **Password**.

This password is used for the `rescue-user` account on each node.

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

c) Provide the **Management Network** IP address and netmask.

d) Provide the **Management Network** IP gateway.

e) Click **Next** to continue.

Step 12   In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the first node.

Step 13   Repeat previous steps to deploy the additional nodes.

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

For multi-node clusters, you must deploy two additional `Primary` nodes and as many `Secondary` nodes as required by your specific use case. The total number of required nodes is available in the Nexus Dashboard Capacity Planning tool.

You do not need to wait for the first node's VM deployment to complete, you can begin deploying the other two nodes simultaneously. The steps to deploy the second and third nodes are identical to the first node's.

Step 14   Wait for the VM(s) to finish deploying.

Step 15   Ensure that the VMware Tools periodic time synchronization is disabled, then start the VMs.

To disable time synchronization:

a)   Right-click the node's VM and select **Edit Settings**.
b)   In the **Edit Settings** window, select the **VM Options** tab.
c)   Expand the **VMware Tools** category and uncheck the **Synchronize time periodically** option.

Step 16   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 17   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 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 18 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.

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 19 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 *

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

Note	In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.
	You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 21   In the **Node Details** page, verify the provided information and click **Next** to continue.

Step 22 Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Note	Depending on the number of nodes in the cluster, some services or cohosting scenarios may not be supported. If you are unable to choose the desired number of services, click **Back** and ensure that you have provided enough secondary nodes in the previous step.
	The deployment mode cannot be changed after the cluster is deployed, so you must ensure that you have completed all service-specific prerequisites described in earlier chapters of this document:

  • Nexus Dashboard Fabric Controller prerequisites

  • Nexus Dashboard Orchestrator prerequisites

  • Nexus Dashboard Insights prerequisites

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

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

c) Click **Next** to proceed.

Step 23   In the **Summary** screen, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

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

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

Step 24   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 25 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.

Deploying Nexus Dashboard Directly in VMware ESXi

This section describes how to deploy Cisco Nexus Dashboard cluster directly in VMware ESXi. If you prefer to deploy using vCenter, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 23 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and guidelines for deploying the Nexus Dashboard cluster in VMware ESX, on page 1.

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

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

b) Choose the Nexus Dashboard release version you want to download.

c) Click the **Download** icon next to the Nexus Dashboard OVA image (`nd-dk9.<version>.ova`).

Step 2 Log in to your VMware ESXi.

Depending on the version of your ESXi server, the location and order of configuration screens may differ slightly. The following steps provide deployment details using VMware ESXi 7.0.

Step 3 Right-click the host and select **Create/Register VM**.
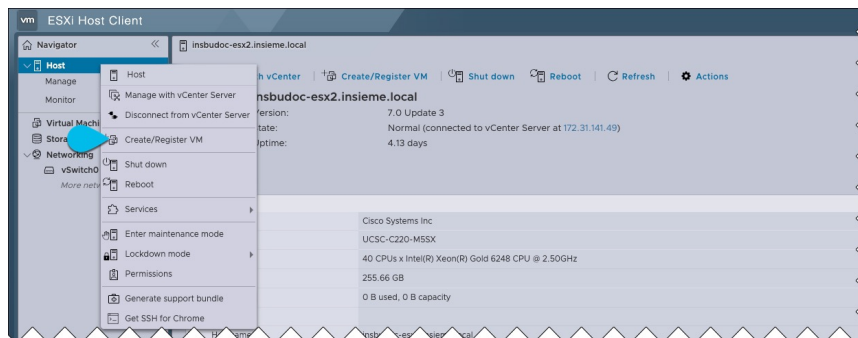
Step 4 In the **Select creation type** screen, choose `Deploy a virtual machine from an OVF or OVA file`, then click **Next**.

Step 5 In the **Select OVF and VMDK files** screen, provide the virtual machine name (for example, nd-ova-node1) and the OVA image you downloaded in the first step, then click **Next**.

Step 6 In the **Select storage** screen, choose the datastore for the VM, then click **Next**.

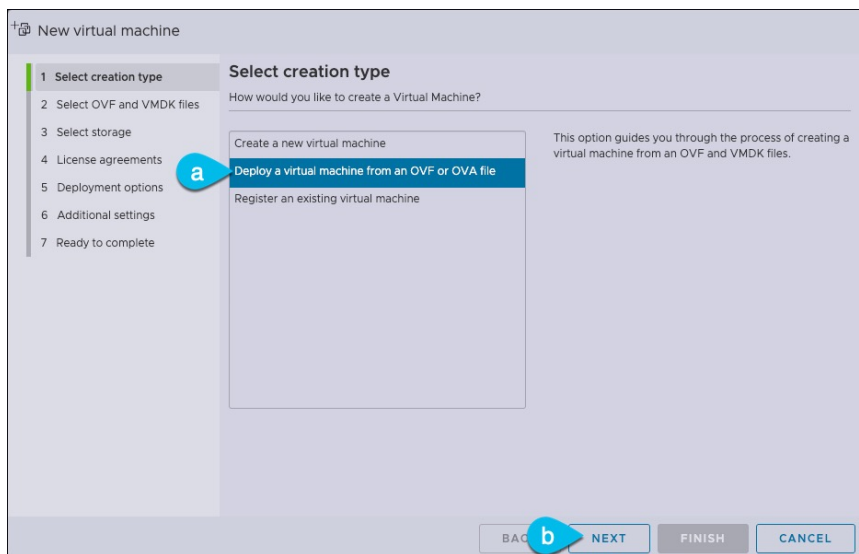**Step 7** In the **Select OVF and VMDK files** screen, provide the virtual machine name (for example, nd-node1) and the OVA image you downloaded in the first step, then click **Next**.

Step 8 Specify the **Deployment options**.

In the **Deployment options** screen, provide the following:

- From the **Network mappings** dropdowns, choose the networks for the Nexus Dashboard management (mgmt0) and data (fabric0) interfaces.

  Nexus Dashboard networks are described in Prerequisites: Nexus Dashboard.

- From the **Deployment type** dropdown, choose the node profile (App or Data).

  Node profiles are described in Prerequisites and guidelines for deploying the Nexus Dashboard cluster in VMware ESX, on page 1.

- For **Disk provisioning** type, choose Thick.

- Disable the **Power on automatically** option.

Step 9 In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the first node.

Step 10 Repeat previous steps to deploy the second and third nodes.

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

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

Step 11 Wait for the VM(s) to finish deploying.

Step 12 Ensure that the VMware Tools periodic time synchronization is disabled, then start the VMs.

To disable time synchronization:

a) Right-click the node's VM and select **Edit Settings**.

b) In the **Edit Settings** window, select the **VM Options** tab.

c) Expand the `VMware Tools` category and uncheck the `Synchronize guest time with host` option.

Step 13 Open one of the node's console and configure the node's basic information.

a) 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: no

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

Step 14 Repeat previous steps to deploy the additional nodes.

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

For multi-node clusters, you must deploy two additional `Primary` nodes and as many `Secondary` nodes as required by your specific use case. The total number of required nodes is available in the Nexus Dashboard Capacity Planning tool.

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 additional nodes are identical with the only exception being that you must indicate that they are not the **Cluster Leader**.

Step 15   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 16   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 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 17   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.

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 18 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*

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

Note	In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.
	You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 20   In the **Node Details** page, verify the provided information and click **Next** to continue.
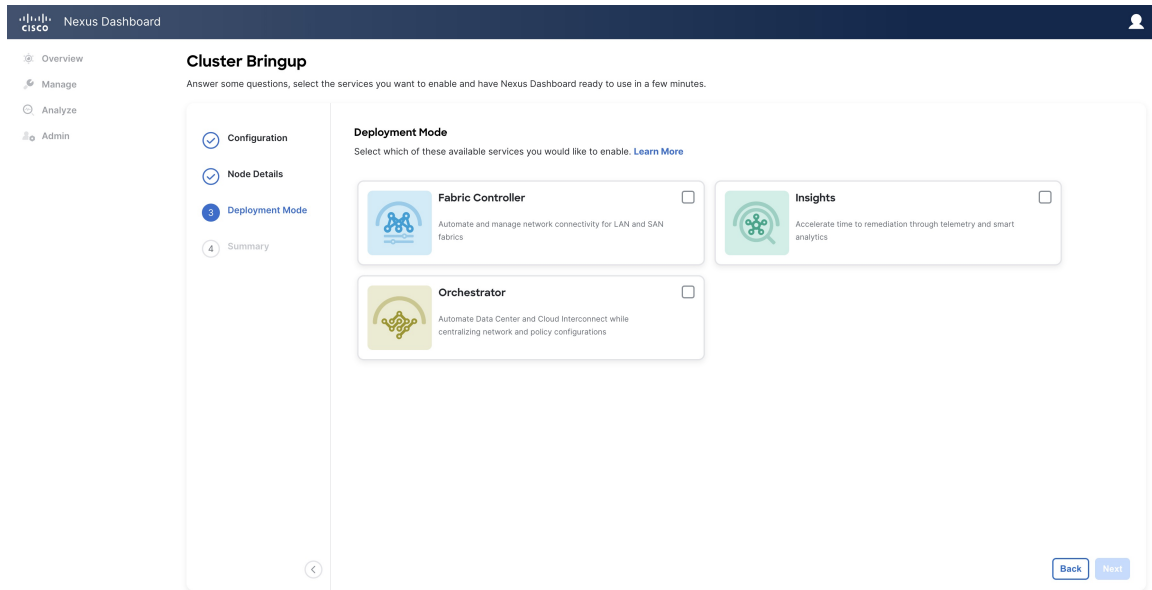
Step 21   Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

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

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

• Nexus Dashboard Fabric Controller prerequisites

• Nexus Dashboard Orchestrator prerequisites

• Nexus Dashboard Insights prerequisites

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

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

c) Click **Next** to proceed.

Step 22 In the **Summary** screen, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

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

It may take up to 30 minutes for the cluster to form and all the services to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.
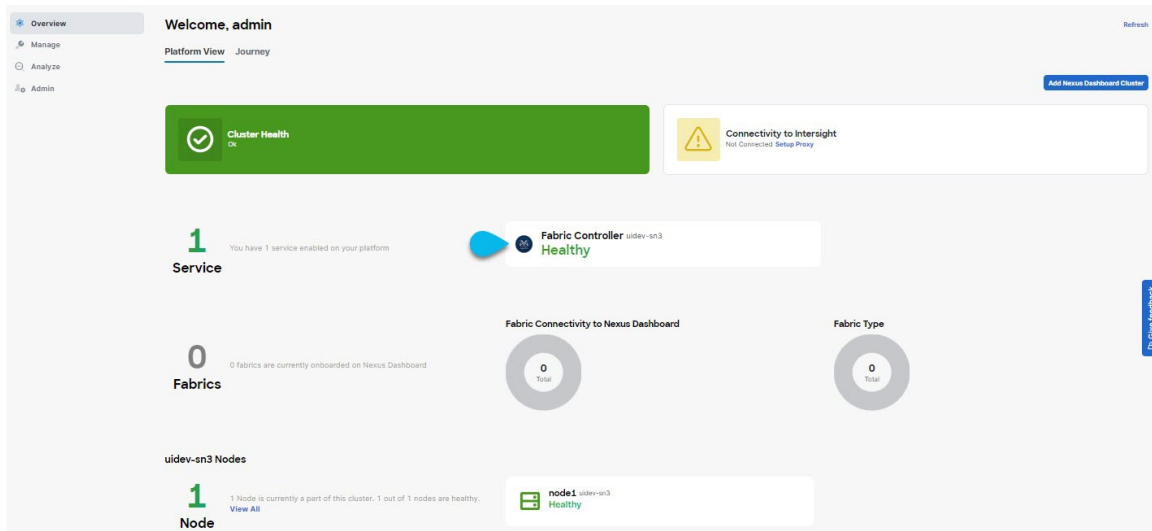
Step 23 Verify that the cluster is healthy.

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 24   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.