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 requirements.

  Scale support and co-hosting vary based on the cluster form factor 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 deployments 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 and Guidelines.

  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*

- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.
- Choose the type of node to deploy:
  - Data node—Node profile with higher system requirements designed for specific Nexus Dashboard features that require the additional resources.
  - App node—Node profile with a smaller resource footprint that can be used for most Nexus Dashboard features.

Note Some larger scale deployments may require additional secondary nodes. If you plan to add secondary nodes to your Nexus Dashboard 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 *Cisco Nexus Dashboard Verified Scalability Guide* 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, 8.0.3
• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2, 8.0.3 if deploying using VMware vCenter
• Each VM requires the following:
  • 32 vCPUs with physical CPU reservation of at least 35,200 MHz
  • 128GB of RAM with physical reservation
  • 3TB 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:
  • The SSD must be attached to the data store directly or in JBOD mode if using a RAID Host Bus Adapter (HBA)
  • 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.
	• VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2, 8.0.3
	• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0, 8.0.2, 8.0.3 if deploying using VMware vCenter
	• Each VM requires the following:
	  • 16 vCPUs with physical CPU reservation of at least 17,600 MHz
	  • 64GB of RAM with physical reservation
	  • 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume
	  Some features require `App` nodes to be deployed on faster SSD storage while other features support HDD. Check the Nexus Dashboard Capacity Planning tool to ensure that you use the correct type of storage.
	• 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 virtual Nexus Dashboard 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 using the content library is not supported.
- VMware snapshots are supported only 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 Deploy Nexus Dashboard Using VMware vCenter, on page 3.

  If you want to deploy directly in ESXi, following the steps described in Deploy Nexus Dashboard Directly in VMware ESXi, on page 15.

Deploy 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 Deploy Nexus Dashboard Directly in VMware ESXi, on page 15 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.

Procedure

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.

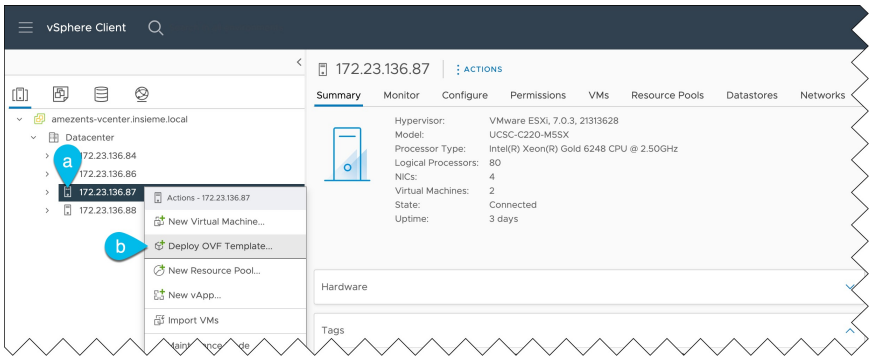

a) Right-click the ESX host where you want to deploy the VM.

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

The **Deploy OVF Template** wizard appears.

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

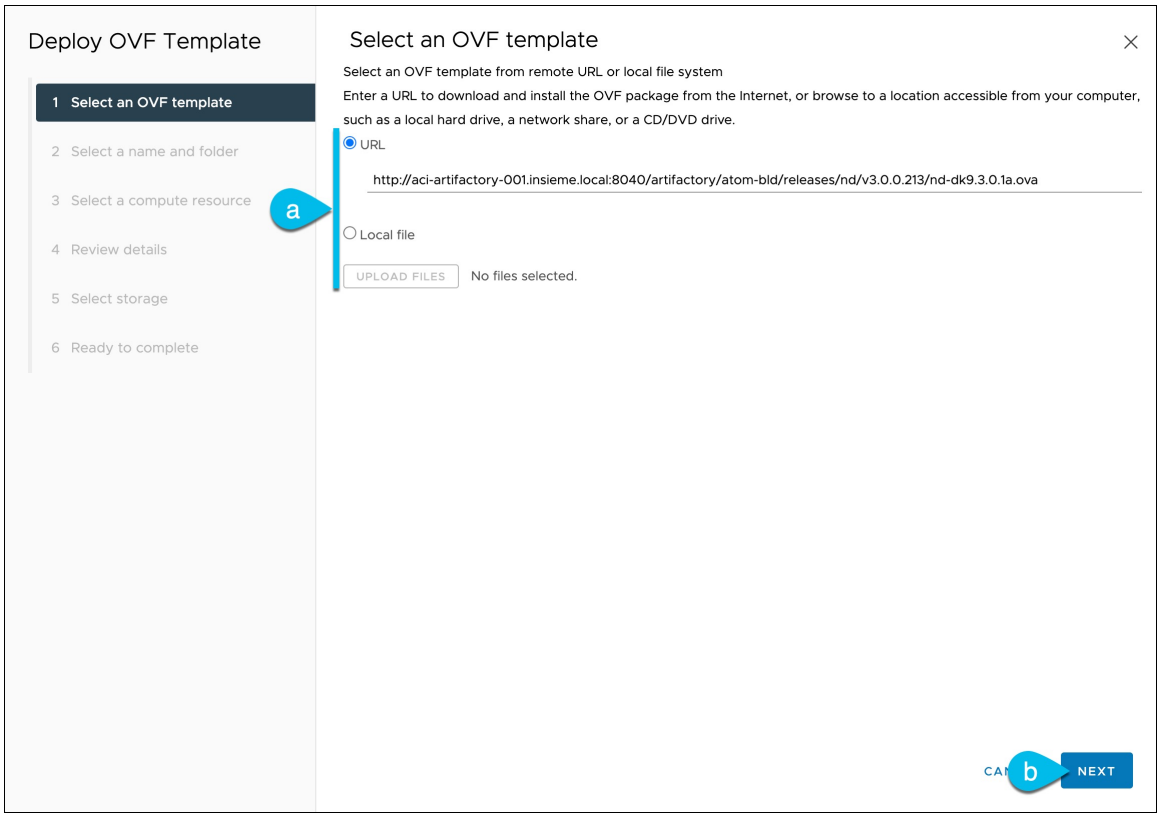

a) Provide the location of the image.

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.

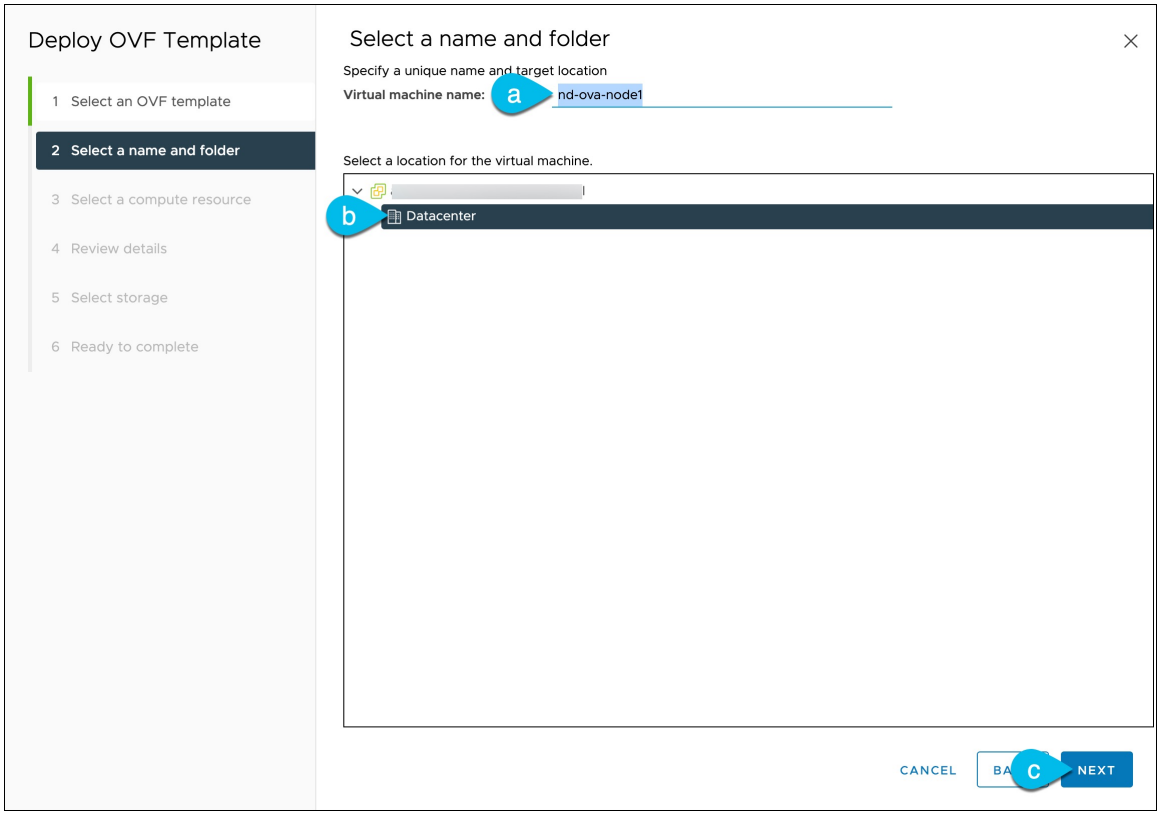

a) Provide the name for the virtual machine.

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.

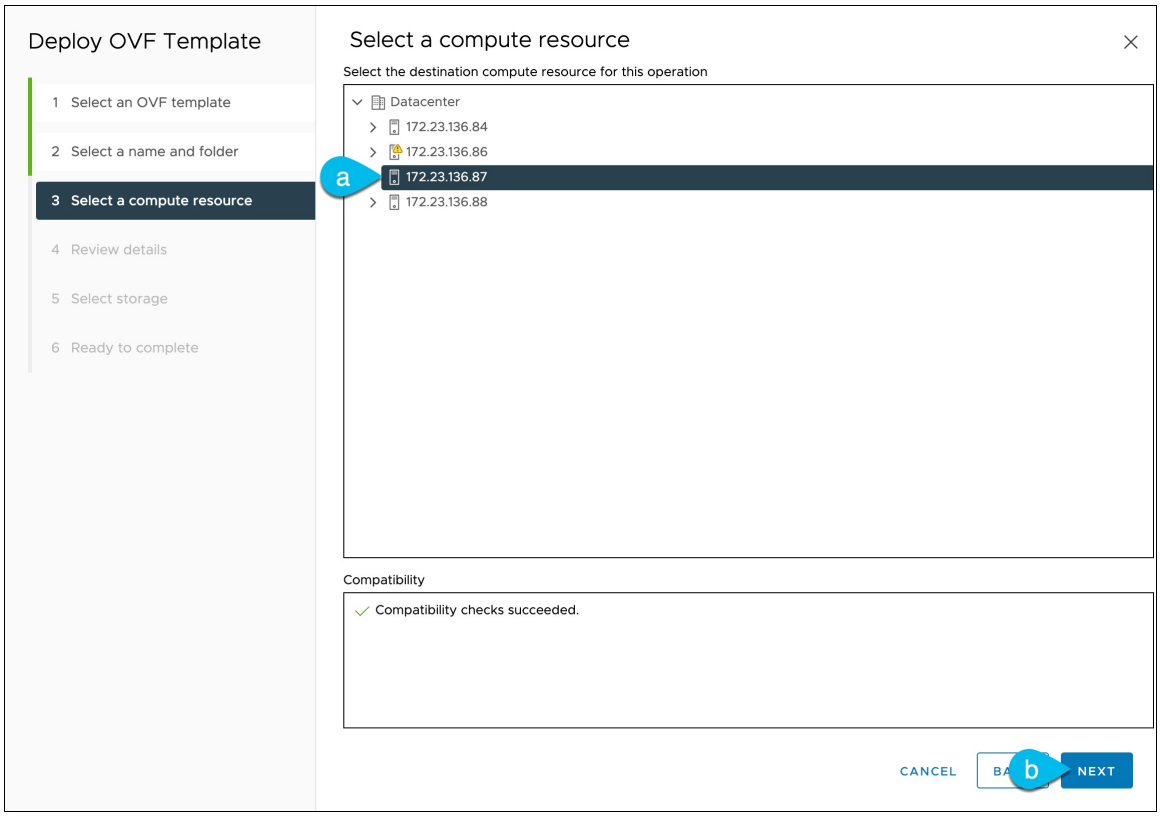

a) Select the vCenter data center and the ESX host for the virtual machine.
b) Click **Next** to continue

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.

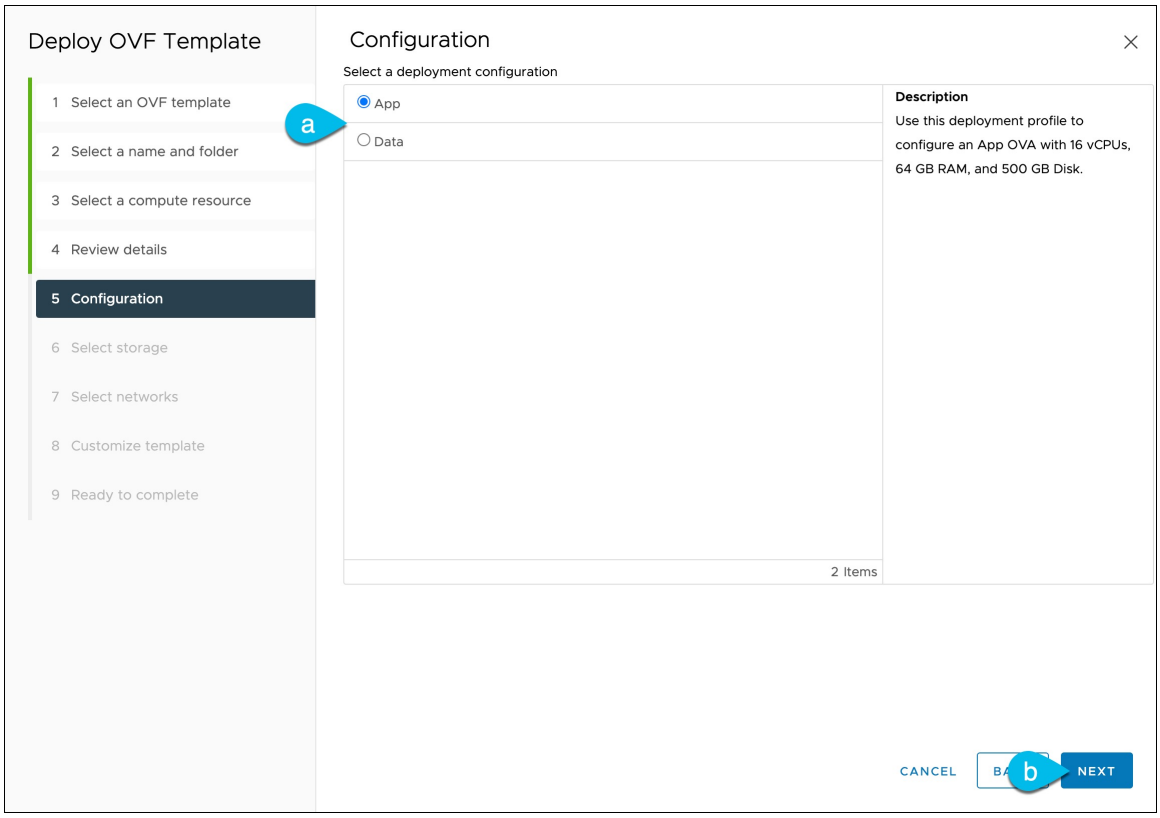

a) Select either `App` or `Data` node profile based on your use case requirements.

   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.

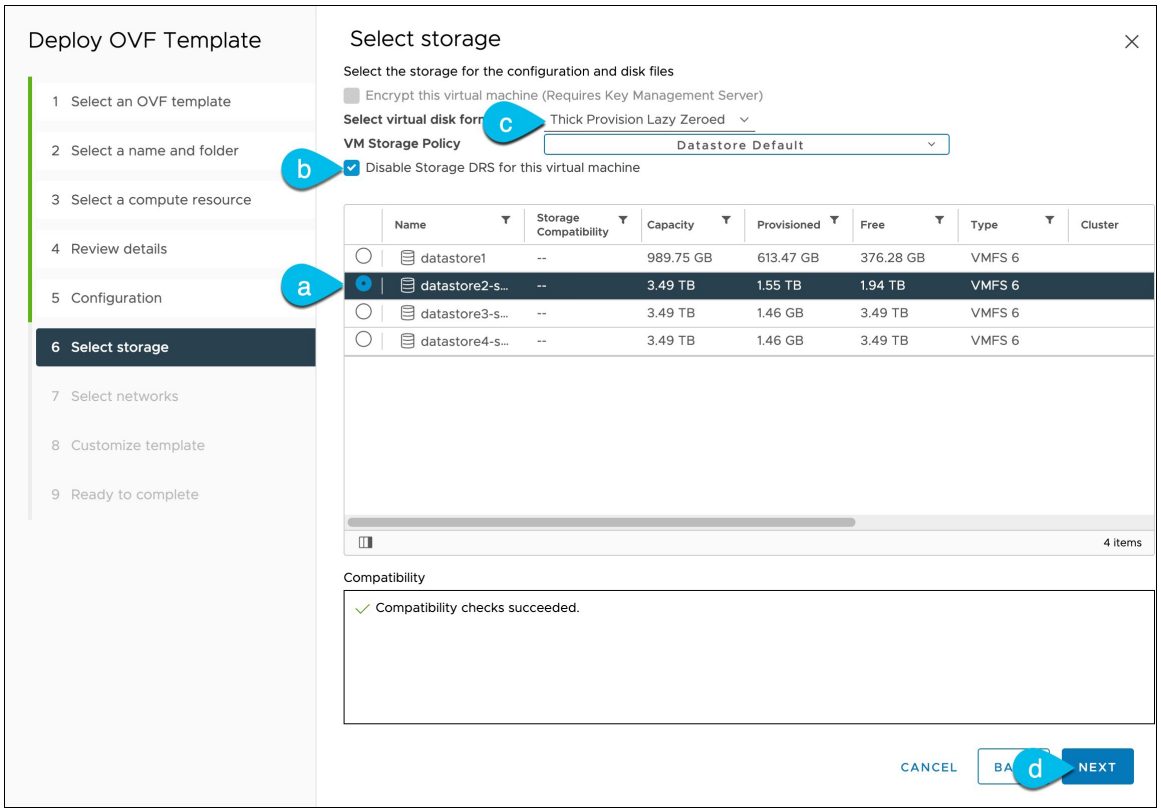

a) Select the datastore for the virtual machine.

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, both of which that have ports configured for high availability:

- **Data network**: The bonded ports **fabric0/fabric1** are used for the Nexus Dashboard cluster's data network.
- **Management network**: The bonded ports **mgmt0/mgmt1** are used for the Nexus Dashboard cluster's management network.

For more information about these networks, see General prerequisites and guidelines in the "Deployment Overview and Requirements" chapter.

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

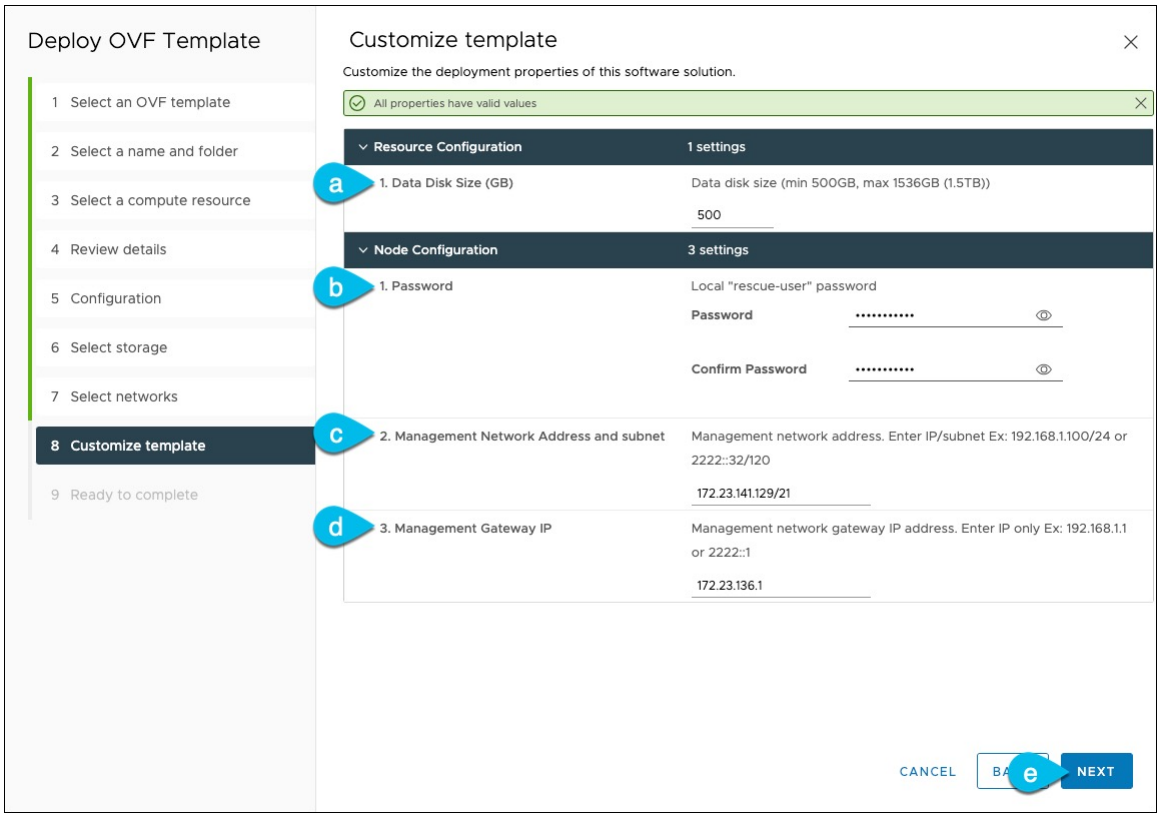

a) Choose the APP/Data type.
b) Provide and confirm the **Password**.

   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 entered in a previous step and click **Login**

Step 17 Enter the requested information in the **Basic Information** page of the **Cluster Bringup** wizard.

a) For **Cluster Name**, enter a name for this Nexus Dashboard cluster.

The cluster name must follow the RFC-1123 requirements.

b) For **Select the Nexus Dashboard Implementation type**, choose either **LAN** or **SAN** then click **Next**.

Step 18 In the **Node Details** page, 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 enter the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.

a) For **Cluster Connectivity**, if your cluster is deployed in L3 HA mode, choose **BGP**. Otherwise, choose **L2**.

BGP configuration is required for the persistent IP addresses feature used by telemetry. This feature is described in more detail in BGP configuration and persistent IP addresses 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. You must enable BGP now if the data network of nodes have different subnets.

b) 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 enter the other information.

c) For **Name**, enter a name for the node.

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

Note
If you need to change the name but the **Name** field is not editable, run the CIMC validation again to fix this issue.

d) For **Type**, choose **Primary**.

The first nodes of the cluster must be set to **Primary**. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note
If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address 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) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

g) 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.

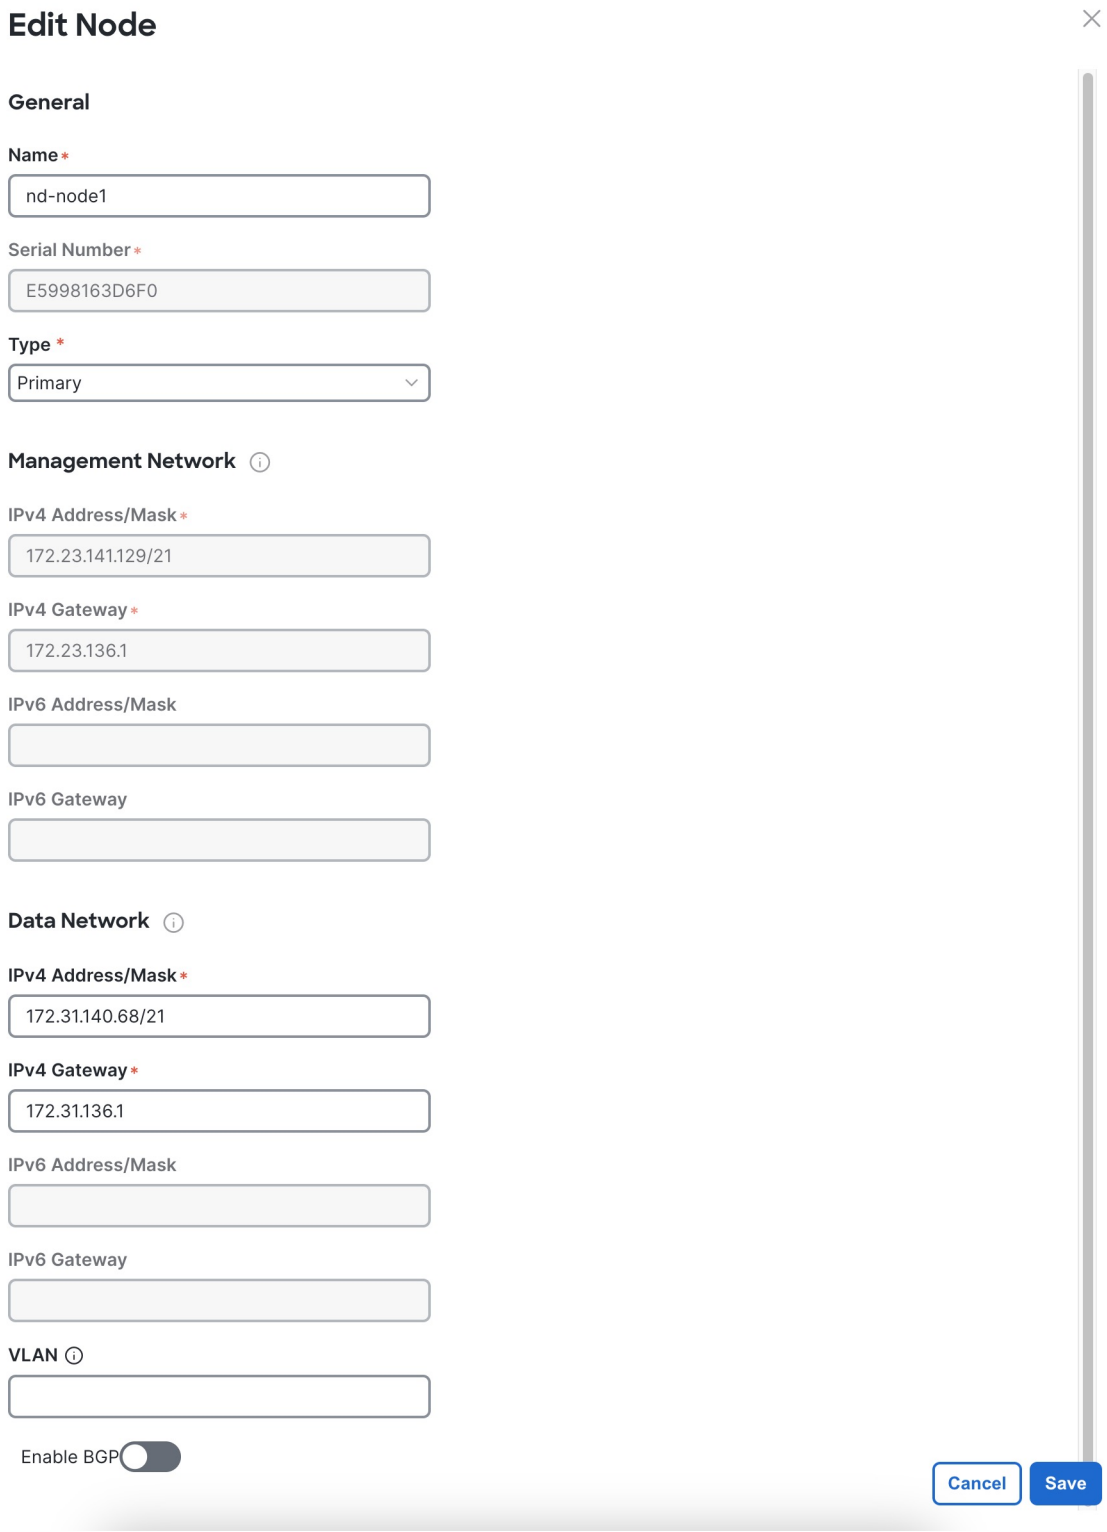

a) In the **Deployment Details** area, provide the **Management IP Address** and **Password** for the second node

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 required for 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 IP addresses feature. This feature is described in more detail in BGP configuration and persistent IP addresses 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 enter information about any additional secondary or standby nodes.

Note
To support higher scale, you must provide a 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 information that you entered, then click **Next**.

Step 22 In the **Persistent IPs** page, if you want to add more persistent IP addresses, click **+ Add Data Service IP Address**, enter the IP address, and click the checkmark icon. Repeat this step as many times as desired, then click **Next**.

You must configure the minimum number of required persistent IP addresses during the bootstrap process. This step enables you to add more persistent IP addresses if desired.

Step 23 In the **Summary** page, 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 60 minutes or more for the cluster to form, depending on the number of nodes in the cluster, and all the features to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 24 Verify that the cluster is healthy.

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 look at the **Anomaly Level** on the **Home** > **Overview** page to ensure the cluster is healthy:

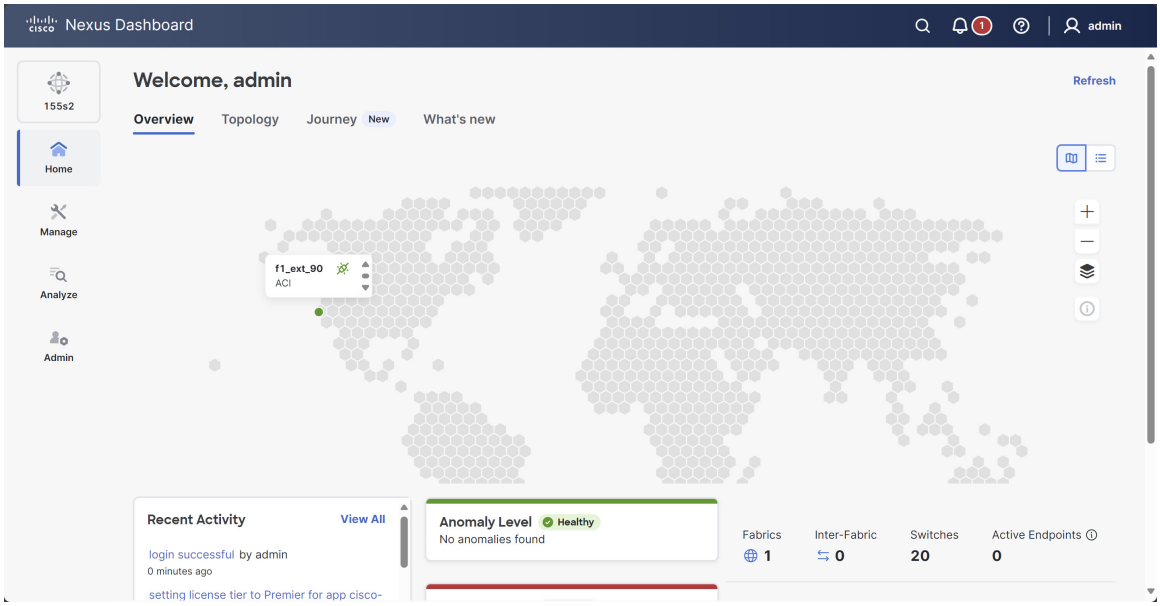

Alternatively, you can log in to any one node using SSH as the `rescue-user` using the password you entered during node deployment and using the `acs health` command to see the status:

- While the cluster is converging, you may see the following output:

```
$ 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 physical Nexus Dashboard cluster.

To resolve the issue, enter the `acs reboot clean` command on the affected node.

Step 25 After you have deployed Nexus Dashboard, see the collections page for this release for configuration information.

What to do next

The next task is to create the fabrics and fabric groups. See the *Creating Fabrics and Fabric Groups* article for this release on the Cisco Nexus Dashboard collections page.

Deploy 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 Deploy Nexus Dashboard Directly in VMware ESXi, on page 15 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.

Procedure

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**.

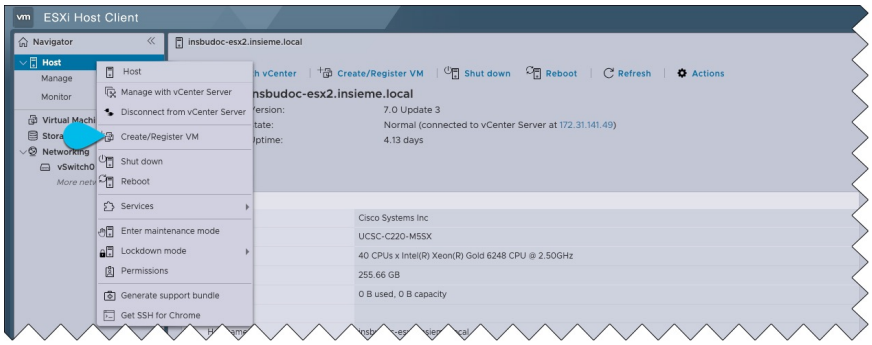

Step 4 In the **Select creation type** screen, choose `Deploy a virtual machine from an OVF or OVA file`, then click **Next**.

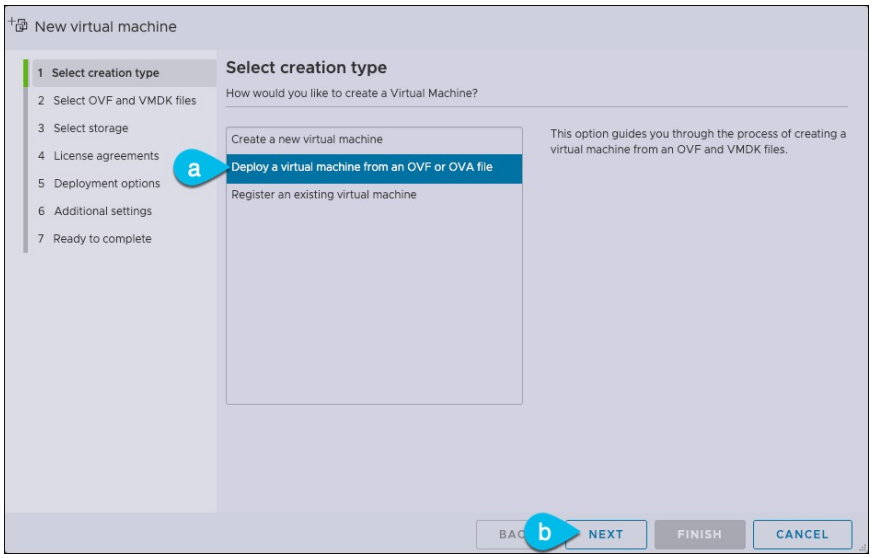

Step 5 In the **Select OVF and VMDK files** screen, provide the virtual machine name (for example, `nd-ova-node1`) and the OVA image you downloaded in the first step, then click **Next**.

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 General prerequisites and guidelines.

- 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 entered in a previous step and click **Login**

Step 16 Enter the requested information in the **Basic Information** page of the **Cluster Bringup** wizard.

a) For **Cluster Name**, enter a name for this Nexus Dashboard cluster.

The cluster name must follow the RFC-1123 requirements.

b) For **Select the Nexus Dashboard Implementation type**, choose either **LAN** or **SAN** then click **Next**.

Step 17 In the **Node Details** page, 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 enter the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.

a) For **Cluster Connectivity**, if your cluster is deployed in L3 HA mode, choose **BGP**. Otherwise, choose **L2**.

BGP configuration is required for the persistent IP addresses feature used by telemetry. This feature is described in more detail in BGP configuration and persistent IP addresses 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. You must enable BGP now if the data network of nodes have different subnets.

b) 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 enter the other information.

c) For **Name**, enter a name for the node.

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

Note
If you need to change the name but the **Name** field is not editable, run the CIMC validation again to fix this issue.

d) For **Type**, choose **Primary**.

The first nodes of the cluster must be set to **Primary**. You will add the secondary nodes in a later step if required for higher scale.

e) In the **Data Network** area, enter the node's data network information.

Enter the data network IP address, netmask, and gateway. Optionally, you can also enter the VLAN ID for the network. Leave the VLAN ID field blank if your configuration does not require VLAN. If you chose **BGP** for **Cluster Connectivity**, enter the ASN.

If you enabled IPv6 functionality in a previous page, you must also enter the IPv6 address, netmask, and gateway.

Note
If you want to enter IPv6 information, you must do so during the cluster bootstrap process. To change the IP address 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) If you chose **BGP** for **Cluster Connectivity**, then in the **BGP peer details** area, enter the peer's IPv4 address and ASN.

You can click **+ Add IPv4 BGP peer** to add addition peers.

If you enabled IPv6 functionality in a previous page, you must also enter the peer's IPv6 address and ASN.

g) 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.

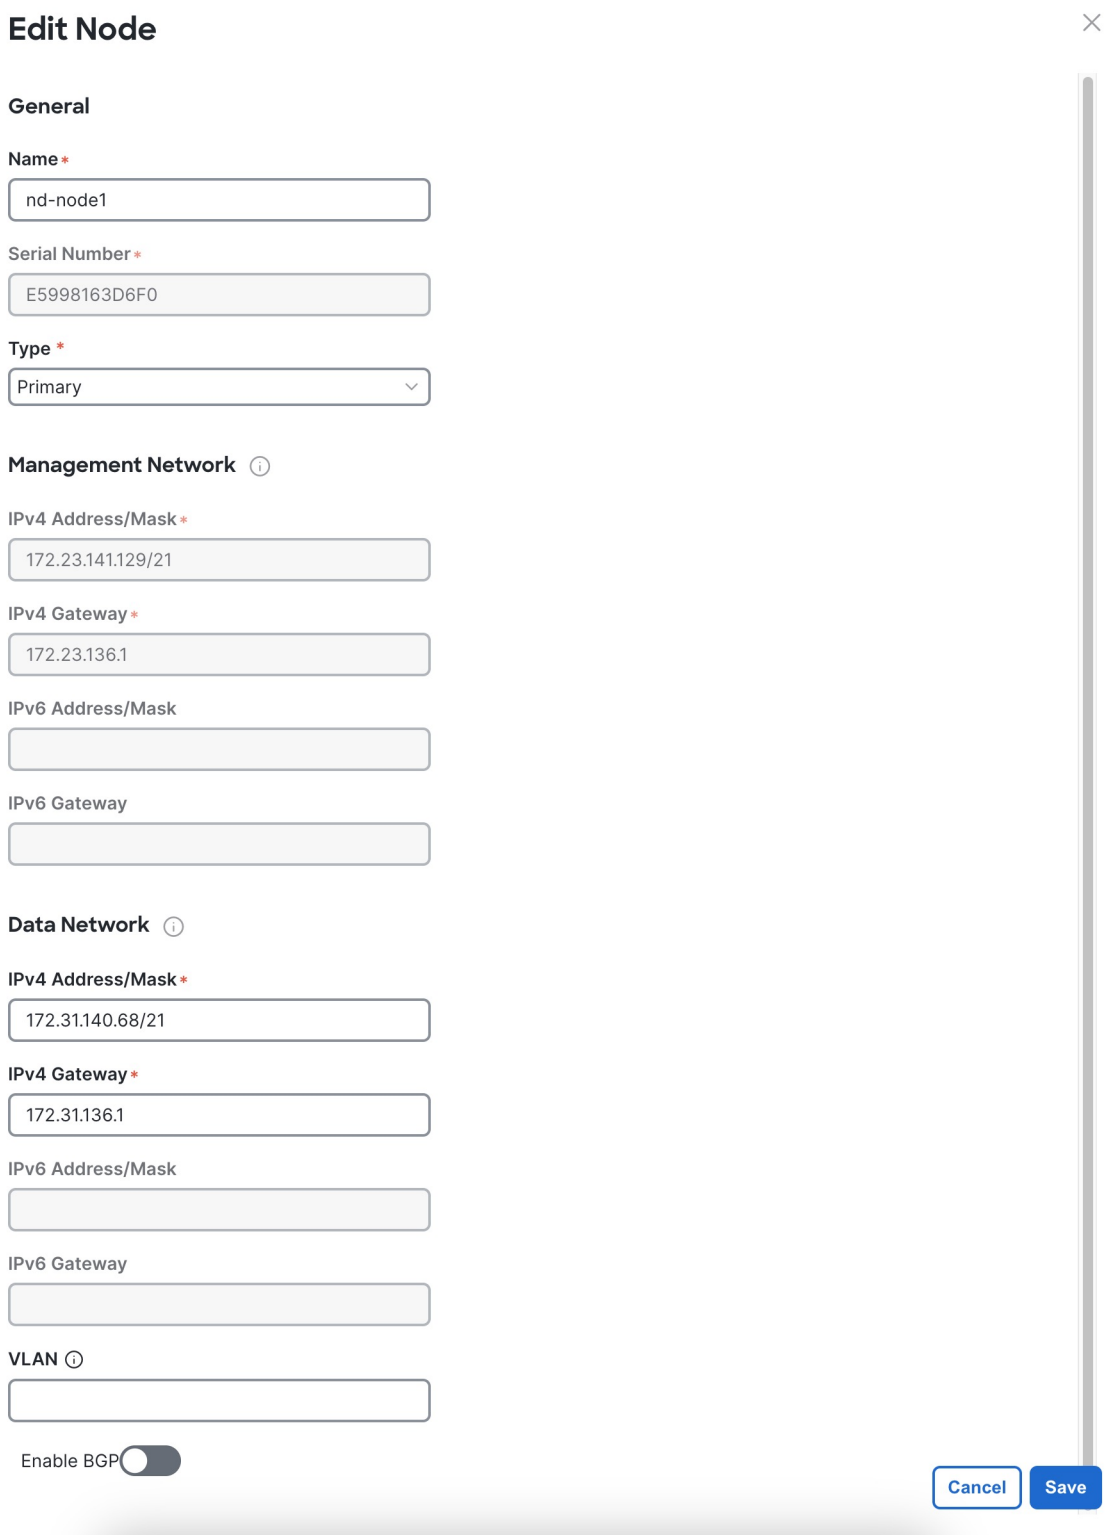

a) In the **Deployment Details** area, provide the **Management IP Address** and **Password** for the second node

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 required for 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 IP addresses feature. This feature is described in more detail in BGP configuration and persistent IP addresses 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 enter information about any additional secondary or standby nodes.

Note
To support higher scale, you must provide a 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 information that you entered, then click **Next**.

Step 21 In the **Persistent IPs** page, if you want to add more persistent IP addresses, click **+ Add Data Service IP Address**, enter the IP address, and click the checkmark icon. Repeat this step as many times as desired, then click **Next**.

You must configure the minimum number of required persistent IP addresses during the bootstrap process. This step enables you to add more persistent IP addresses if desired.

Step 22 In the **Summary** page, 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 60 minutes or more for the cluster to form, depending on the number of nodes in the cluster, and all the features to start. When cluster configuration is complete, the page will reload to the Nexus Dashboard GUI.

Step 23 Verify that the cluster is healthy.

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 look at the **Anomaly Level** on the **Home** > **Overview** page to ensure the cluster is healthy:

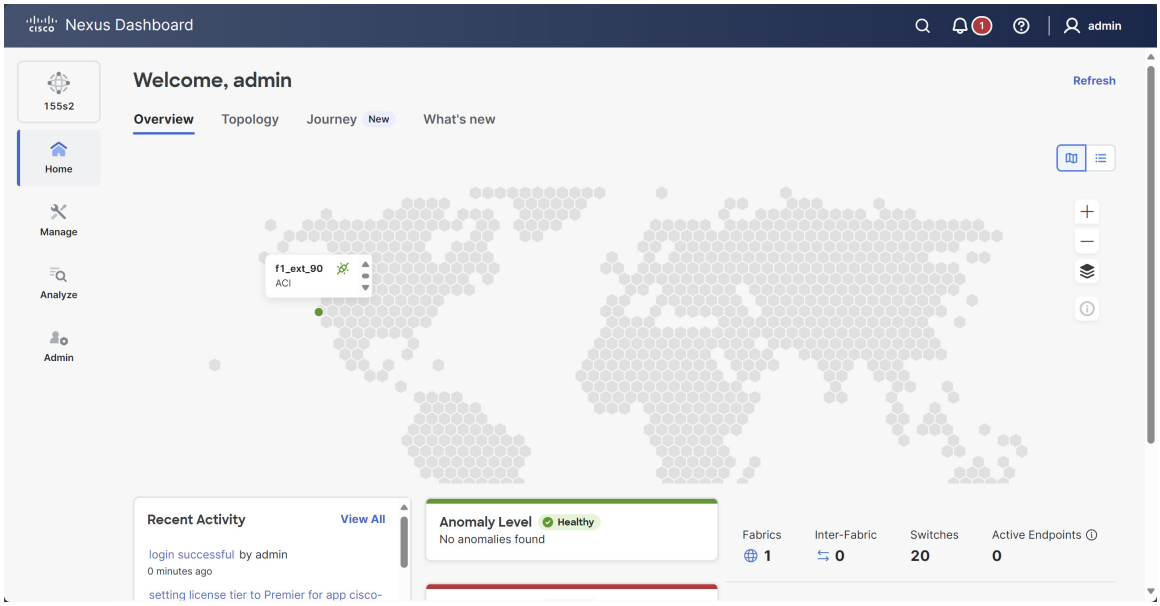

Alternatively, you can log in to any one node using SSH as the `rescue-user` using the password you entered during node deployment and using the `acs health` command to see the status:

- While the cluster is converging, you may see the following output:

```
$ 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 physical Nexus Dashboard cluster.

To resolve the issue, enter the `acs reboot clean` command on the affected node.

Step 24 After you have deployed Nexus Dashboard, see the collections page for this release for configuration information.