Cisco Nexus Dashboard Deployment Guide, Release 2.0.x

First Published: 2020-10-30

Last Modified: 2022-03-22

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

CONTENTS

CHAPTER 1

New and Changed Information

New and Changed Information

The following table provides an overview of the significant changes to the organization and features in this guide from the release in which the guide was first published to the current release. The table does not provide an exhaustive list of all changes made to the guide.

Table 1: Latest Updates

Release	New Feature or Update	Where Documented
2.0.2h	Deployment information for virtual and cloud form factors.	Deploying in VMware ESX, on page 29 Deploying in Amazon Web Services, on page 45 Deploying in Microsoft Azure, on page 53
2.0.2g	Additional information on upgrading to this release.	Upgrading Nexus Dashboard, on page 59
2.0.1	First release of this document.	--

New and Changed Information

CHAPTER 2

Deployment Overview and Requirements

Deployment Overview

Cisco Nexus Dashboard is a central management console for multiple data center sites and a common platform for hosting Cisco data center operation applications, such as Nexus Insights and Nexus Assurance Engine. These applications are universally available for all the data center sites and provide real time analytics, visibility, and assurance for network policies and operations. Cisco Multi-Site Orchestrator can also run on Nexus Dashboard as a hosted application.

Nexus Dashboard provides a common platform and modern technology stack for these micro-services-based applications, simplifying the life cycle management of the different modern applications and reducing the operational overhead to run and maintain these applications. It also provides a central integration point for external 3rd party applications with the locally hosted applications.

Each Nexus Dashboard cluster consists of 3 `master` nodes. For physical Nexus Dashboard clusters, you can also provision up to 4 `worker` nodes to enable horizontal scaling and up to 2 `standby` nodes for easy cluster recovery in case of a master node failure. For virtual and cloud clusters, only the base 3-node cluster is supported.

Note This document describes initial configuration of the 3-node cluster. After your cluster is up and running, you can configure and deploy additional nodes as described in the *Cisco Nexus Dashboard User Guide*, which is also available directly from the Nexus Dashboard GUI.

Hardware vs Software Stack

Nexus Dashboard is offered as a cluster of specialized Cisco UCS servers (Nexus Dashboard platform) with the software framework (Nexus Dashboard) pre-installed on it. The Cisco Nexus Dashboard software stack can be decoupled from the hardware and deployed in a number of virtual form factors. For the purposes of

this document, we will use "Nexus Dashboard platform" specifically to refer to the hardware and "Nexus Dashboard" to refer to the software stack and the GUI console.

This guide describes the initial deployment of the Nexus Dashboard software; hardware setup is described in the *Nexus Dashboard Hardware Setup Guide*, while other Nexus Dashboard operations procedures are described in the *Cisco Nexus Dashboard User Guide*.

Nexus Dashboard and Cisco DCNM

Nexus Dashboard may be used in context of Cisco DCNM. In this case, DCNM is not an application running in the Nexus Dashboard software stack. Instead, the DCNM image (`.iso`) is installed directly on the Nexus Dashboard physical servers in order to provide additional compute resources to the applications installed and running in Cisco DCNM thus enabling horizontal scaling of the DCNM platform. As this document deals with the Nexus Dashboard software stack deployments, see a Cisco DCNM Installation Guide appropriate for your deployment type for information related to installing DCNM on Nexus Dashboard hardware.

Available Form Factors

Cisco Nexus Dashboard, Release 2.0.1 and 2.0.2g can be deployed as a physical appliance only. This refers to software stack already deployed on the Nexus Dashboard platform hardware that you purchase

Cisco Nexus Dashboard, Release 2.0.2h can be deployed using a number of different form factors. Keep in mind however, you must use the same form factor for all nodes, mixing different form factors within the same cluster is not supported.

Note Nexus Dashboard, Release 2.0.2h supports virtual form factor clusters for Multi-Site Orchestrator application only. For other applications, such as Nexus Insights, you must deploy a physical cluster.

- Cisco Nexus Dashboard physical appliance (`.iso`)

  This form factor refers to the original physical appliance hardware that you purchased with the Cisco Nexus Dashboard software stack pre-installed on it.

  The later sections in this document describe how to configure the software stack on the existing physical appliance hardware to deploy the cluster. Setting up the original Cisco Nexus Dashboard platform hardware is described in *Cisco Nexus Dashboard Hardware Setup Guide*.

- VMware ESX (`.ova`)

  Virtual form factor that allows you to deploy a Nexus Dashboard cluster using three VMware ESX virtual machines.

- Amazon Web Services (`.ami`)

  Cloud form factor that allows you to deploy a Nexus Dashboard cluster using three AWS instances.

- Microsoft Azure (`.arm`)

  Cloud form factor that allows you to deploy a Nexus Dashboard cluster using three Azure instances.

Upgrading From Previous Versions of Nexus Dashboard

If you are already running a Nexus Dashboard, Release 2.0.1 or later, you can upgrade to the latest release while retaining the cluster configuration and applications, as described in Upgrading Nexus Dashboard, on page 59

Upgrading From Application Services Engine

If you are running Application Services Engine, Release 1.1.3d as a physical appliance, you can upgrade to Nexus Dashboard to retain the cluster configuration and applications, as described in Upgrading Nexus Dashboard, on page 59

If you are running Application Services Engine, Release 1.1.3d as a virtual appliance or a release prior to Release 1.1.3d, stateful upgrade or migration of the cluster is supported to Nexus Dashboard, Release 2.0.2h or later only. If you want to deploy Release 2.0.1 or 2.0.2g, you would need to deploy a brand new physical appliance cluster and reinstall all the applications.

Cluster Sizing Guidelines

Nexus Dashboard supports co-hosting of applications. Depending on the type and number of applications you choose to run, you may be required to deploy additional worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see Cisco Nexus Dashboard Cluster Sizing.

After your initial 3-node cluster is up and running, you can configure and deploy additional nodes as described in the *Cisco Nexus Dashboard User Guide*, which is also available directly from the Nexus Dashboard GUI.

Supported Applications

For the full list of supported applications and the associated compatibility and interoperability information, see the Cisco Day-2 Operations Apps Support Matrix.

The following table provides a reference for the recommended application release versions for Nexus Dashboard, Release 2.x:

Table 2: Recommended Application Versions

Nexus Dashboard Release and Form Factor	Nexus Insights	Multi-Site Orchestrator	Network Assurance Engine
Nexus Dashboard, Release 2.0.1 Physical cluster	5.0(1)	3.2(1)	5.1(1a)
Nexus Dashboard, Release 2.0.2g Physical cluster	5.1(1)	3.2(1)	5.1(1b)
Nexus Dashboard, Release 2.0.2h Physical cluster	5.1(1)	3.3(1)	5.1(1b)

Nexus Dashboard Release and Form Factor	Nexus Insights	Multi-Site Orchestrator	Network Assurance Engine
Nexus Dashboard, Release 2.0.2h Virtual cluster	Not supported	3.3(1)	Not supported

Prerequisites and Guidelines

Network Time Protocol (NTP)

The Nexus Dashboard nodes use NTP for clock synchronization, so you must have an NTP server configured in your environment.

Nexus Dashboard External Networks

Cisco Nexus Dashboard is deployed as a cluster, connecting each service node to two networks. When first configuring Nexus Dashboard, you will need to provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management Network.

Individual applications installed in the Nexus Dashboard may utilize the two networks for additional purposes, so we recommend consulting the specific application's documentation in addition to this document for your deployment planning.

- **Data Network** is used for:
  - Nexus Dashboard node clustering
  - Application to application communication
  - Nexus Dashboard nodes to Cisco APIC, Cloud APIC, and DCNM communication

    For example, the network traffic for Day-2 Operations applications such as NAE.
- **Management Network** is used for:
  - Accessing Nexus Dashboard GUI
  - Accessing Nexus Dashboard CLI via SSH
  - DNS and NTP communication
  - Nexus Dashboard firmware upload
  - Accessing Cisco DC App Center (AppStore)

    If you want to use the Nexus Dashboard App Store to install applications, https://dcappcenter.cisco.com must be reachable via the Management Network
  - Intersight device connector

The two networks have the following requirements:

- The two interfaces can be in the same or different subnets.

In addition, each network's interfaces across different nodes in the cluster can also be in different subnets.

- The management network must provide IP reachability to each node's CIMC via TCP ports 22/443.

  Nexus Dashboard cluster configuration uses each node's CIMC IP address to configure the node.

- For Nexus Insights and Network Assurance Engine applications, the data network must provide IP reachability to the in-band network of each fabric and of the APIC.
- For Nexus Insights and AppDynamics integration, the data network must provide IP reachability to the AppDynamics controller.
- For Multi-Site Orchestrator application, the data network can have in-band and/or out-of-band IP reachability for Cisco APIC sites but must have in-band reachability for Cisco DCNM sites.
- The data network interface requires a minimum MTU of 1500 to be available for the Nexus Dashboard traffic.

  Higher MTU can be configured if desired.

- Connectivity between the nodes is required on both networks with the following additional round trip time (RTT) requirements.

  You must always use the lowest RTT requirement when deploying the Nexus Dashboard cluster and applications. For example, if you plan to co-host MSO and NI apps, site connectivity RTT must not exceed 50ms.

Table 3: RTT Requirements

Application	Connectivity	Maximum RTT
Nexus Dashboard cluster	Between nodes	150 ms
Multi-Site Orchestrator (MSO)	Between nodes	150 ms
	To sites	500 ms
Nexus Insights (NI)	Between nodes	50 ms
	To sites	50 ms
Network Assurance Engine (NAE)	Between nodes	50 ms
	To sites	50 ms

Nexus Dashboard Internal Networks

Two additional internal networks are required for communication between the containers used by the Nexus Dashboard:

- **Application overlay** is used for applications internally within Nexus Dashboard

  Application overlay must be a `/16` network and a default value is prepopulated during deployment.

- **Service overlay** is used internally by the Nexus Dashboard.

  Service overlay must be a `/16` network and a default value is prepopulated during deployment.

Note Communications between containers deployed in different Nexus Dashboard nodes is VXLAN-encapsulated and uses the data interfaces IP addresses as source and destination. This means that the Application Overlay and Service Overlay addresses are never exposed outside the data network and any traffic on these subnets is routed internally and does not leave the cluster nodes. For example, if you had another service (such as DNS) on the same subnet as one of the overlay networks, you would not be able to access it from your Nexus Dashboard as the traffic on that subnet would never be routed outside the cluster. As such, when configuring these networks, ensure that they are unique and do not overlap with any existing networks or services external to the cluster, which you may need to access from the Nexus Dashboard cluster nodes.

Communication Ports

The following ports are required by the Nexus Dashboard cluster and its applications:

Table 4:

Interface	Port Number	Port Type
Management Interface	--	ICMP
	22	TCP
	67	UDP
	69	UDP
	443	TCP
	5555	TCP
	9880	TCP
	30012	TCP
	30021	TCP
	30500-30600	TCP/UDP

Interface	Port Number	Port Type
Data Interface between ND nodes	53	TCP/UDP
	443	TCP
	3379	TCP
	3380	TCP
	4789	UDP
	9969	TCP
	9979	TCP
	9989	TCP
	15223	TCP
	30002-30006	TCP
	30009-30010	TCP
	30012	TC
	30015-30019	TCP
	30017	UDP
	30025	TCP
	30500-30600	TCP/UDP
Data Interface on APICs	22	TCP
	443	TCP
Data Interface between ND nodes and fabrics	443	TCP
	2022	TCP
	5640-5671	UDP
	5965	UDP
	8884	TCP
	9989	TCP
	30000-30001	TCP

Fabric Connectivity

The following sections describe how to connect your Nexus Dashboard cluster to your fabrics.

For on-premises APIC or DCNM fabrics, you can connect the Nexus Dashboard cluster in one of two ways:

- The Nexus Dashboard cluster connected to the fabric via a Layer 3 network.
- The Nexus Dashboard nodes connected to the leaf switches as typical hosts.

For Cloud APIC fabrics, you will need to connect via a Layer 3 network.

Physical Node Cabling

If you plan to deploy a virtual or cloud form factor cluster, you can skip this section.

The following figure shows the Nexus Dashboard physical node interfaces:

- `eth1-1` and `eth1-2` must be connected to the Management network
- `eth2-1` and `eth2-2` must be connected to the Data network

Figure 1: Node Connectivity

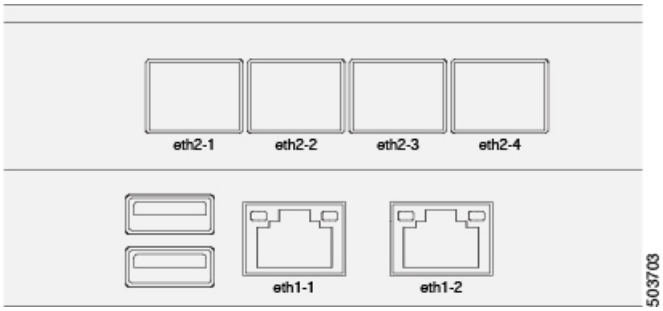

The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode. All interfaces must be connected to individual host ports, PortChannel or vPC are not supported.

When Nexus Dashboard nodes are connected to Cisco Catalyst switches, packets are tagged with `vlan0` if no VLAN is specified. In this case, you must add `switchport voice vlan dot1p` command to the switch interfaces where the nodes are connected to ensure reachability over the data network.

Connecting via External Layer 3 Network

We recommend connecting the Nexus Dashboard cluster to the fabrics via an external Layer 3 network as it does not tie the cluster to any one fabric and the same communication paths can be establish to all sites. Specific connectivity depends on the type of applications deployed in the Nexus Dashboard:

- If you are deploying Multi-Site Orchestrator to manage Cisco ACI fabrics only, you can establish connectivity from the data interface to either the in-band or out-of-band (OOB) interface of each site's APIC or both.
- If you are deploying Multi-Site Orchestrator to manage Cisco DCNM fabrics, you must establish connectivity from the data interface to the in-band interface of each site's DCNM.

- If you are deploying Day-2 Operations applications, such as Nexus Insights, you must establish connectivity from the data interface to the in-band network of each fabric and of the APIC.

If you plan to connect the cluster across a Layer 3 network, keep the following in mind:

- For ACI fabrics, you must configure an L3Out and the external EPG for Cisco Nexus Dashboard data network connectivity in the management tenant.

  Configuring external connectivity in an ACI fabric is described in Cisco APIC Layer 3 Networking Configuration Guide.

- For DCNM fabrics, if the data interface and DCNM's inband interface are in different subnets, you must add a route on DCNM to reach the Nexus Dashboard's data network address.

  You can add the route from the DCNM UI by navigating to **Administration** > **Customization** > **Network Preference** > **In-Band (eth2)** , then adding the route and saving.

- If you specify a VLAN ID for your data interface during setup of the cluster, the host port must be configured as `trunk` allowing that VLAN.

  However, in most common deployments, you can leave the VLAN ID empty and configure the host port in `access` mode.

The following two figures show two distinct network connectivity scenarios when connecting the Nexus Dashboard cluster to the fabrics via a Layer 3 network. The primary purpose of each depends on the type of application you may be running in your Nexus Dashboard.

Note that the "L3 Network" and the "Management Network" can be the same network infrastructure, for example in case the Nexus Dashboard nodes have the management and data network interfaces in the same subnet.

Figure 2: Connecting via Layer 3 Network, Day-2 Operations Applications

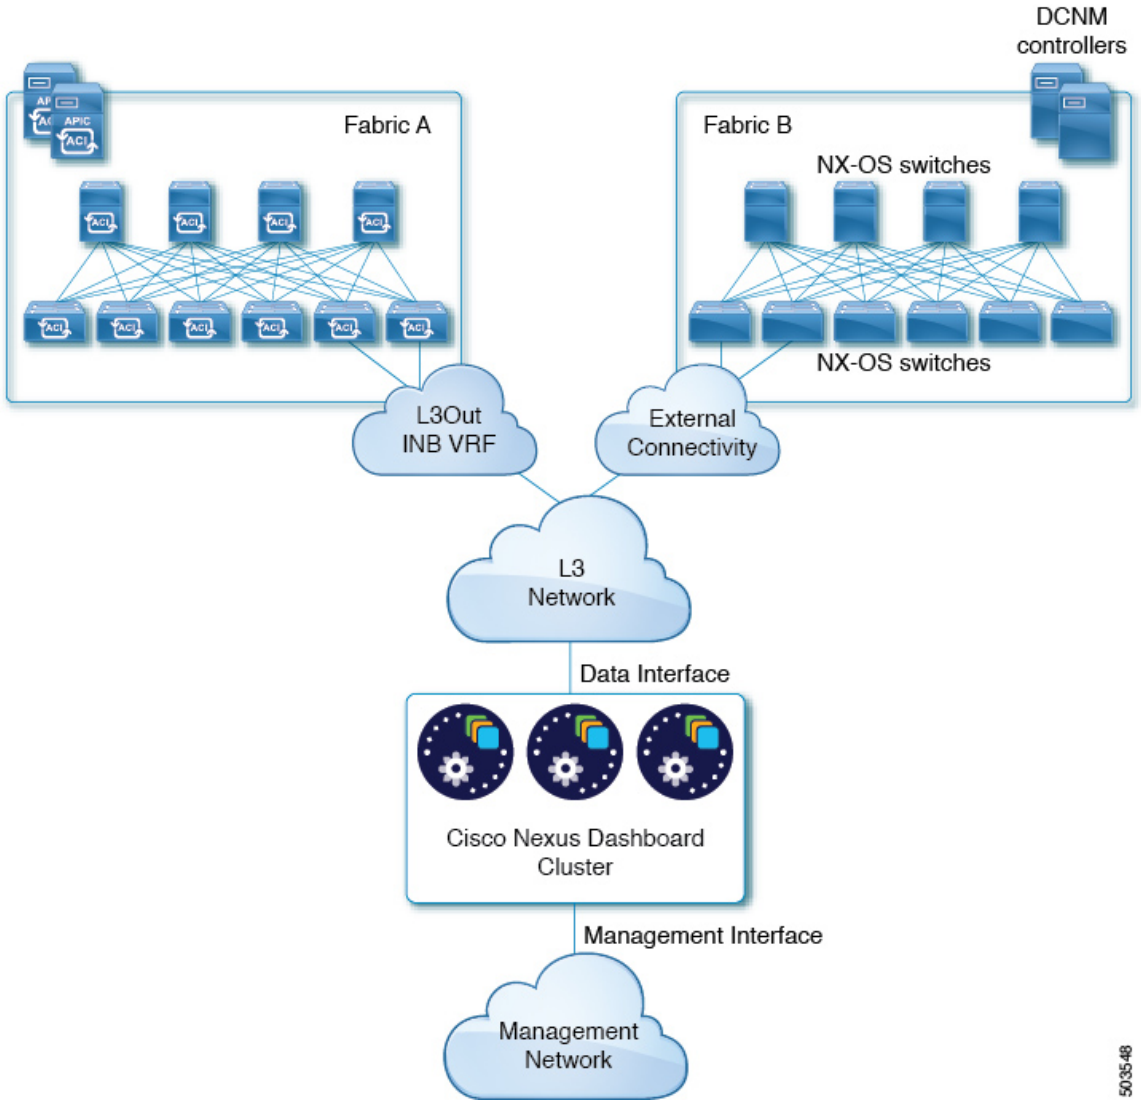

Figure 3: Connecting via Layer 3 Network, Multi-Site Orchestrator

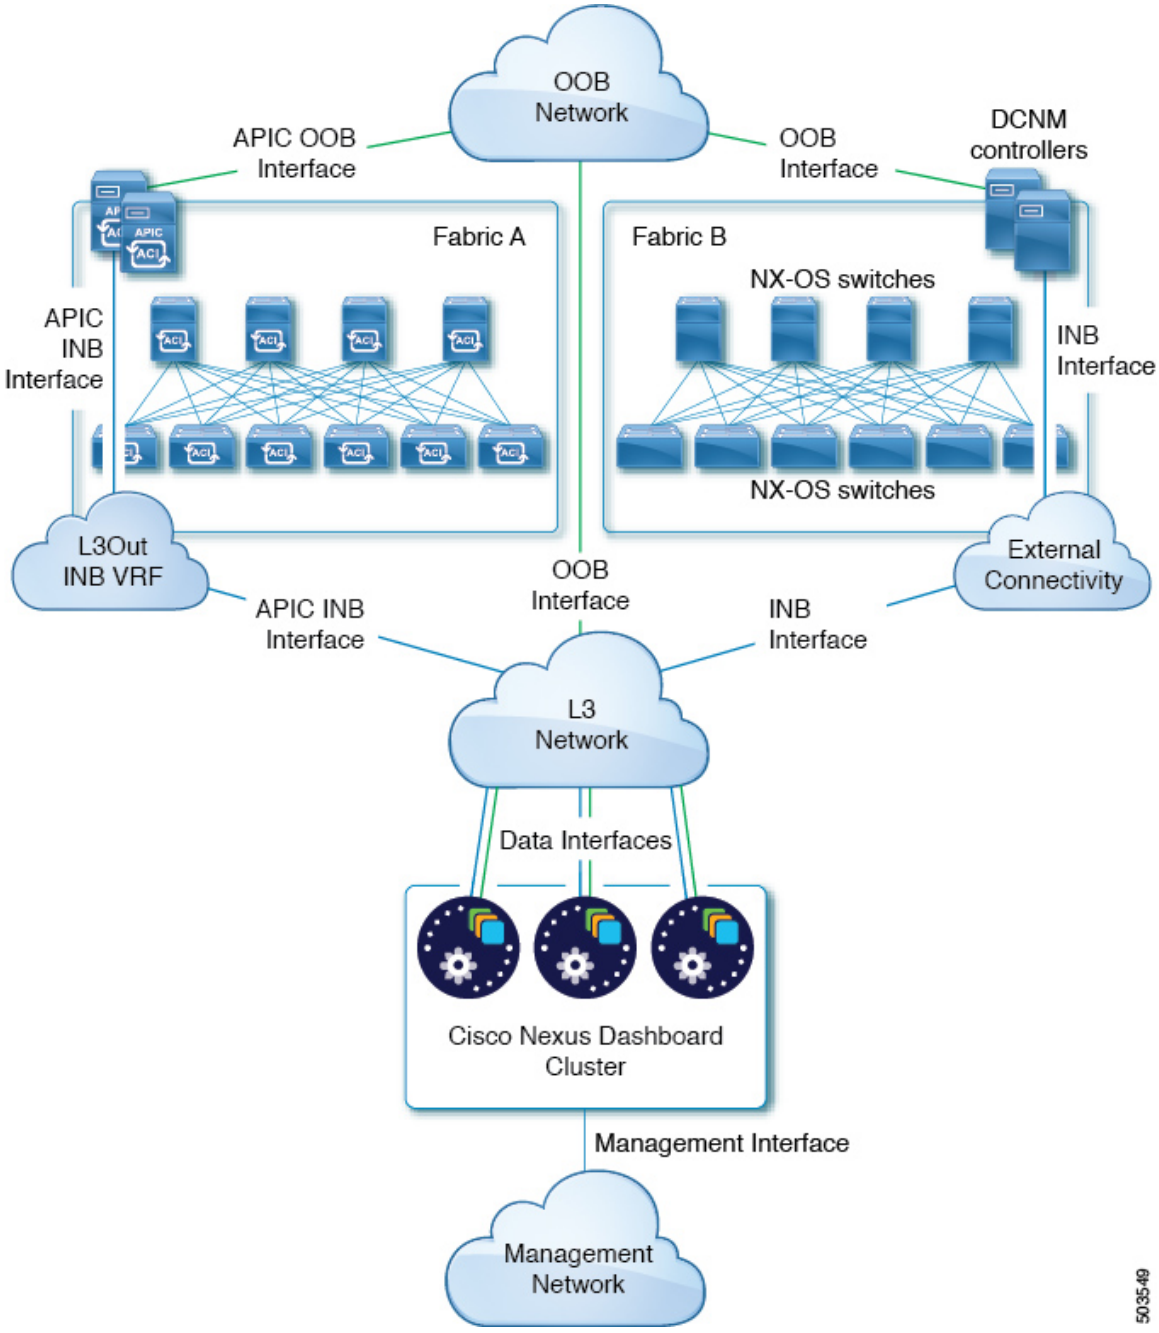

Connecting Nodes Directly to Leaf Switches

You can also connect the Nexus Dashboard cluster directly to one of the fabrics. This provides easy connectivity between the cluster and in-band management of the fabric, but ties the cluster to the specific fabric and requires reachability to other fabrics to be established through external connectivity. This also makes the cluster dependent on the specific fabric so issues within the fabric may impact Nexus Dashboard connectivity. Like in the previous example, connectivity depends on the type of applications deployed in the Nexus Dashboard:

- If you are deploying Multi-Site Orchestrator to manage Cisco ACI fabrics only, you can establish connectivity from the data interface to either the in-band or out-of-band (OOB) interface of each site's APIC
- If you are deploying Nexus Insights or Network Assurance Engine, you must establish connectivity from the data interface to the in-band interface of each fabric.

  For ACI fabrics, the data interface IP subnet connects to an EPG/BD in the fabric and must have a contract established to the local in-band EPG in the management tenant. We recommend deploying the Nexus Dashboard in the management tenant and in-band VRF. Connectivity to other fabrics is established via an L3Out.
- If you are deploying Nexus Insights with ACI fabrics, the data interface IP address and the ACI fabric's in-band IP address must be in different subnets.

If you plan to connect the cluster directly to the leaf switches, keep the following in mind:

- If deploying in VMware ESX or Linux KVM, the host must be connected to the fabric via trunk port.
- For ACI fabrics, we recommend configuring the bridge domain (BD), subnet, and endpoint group (EPG) for Cisco Nexus Dashboard connectivity in management tenant.

  Because the Nexus Dashboard requires connectivity to the in-band EPG in the in-band VRF, creating the EPG in the management tenant means no route leaking is required.
- For ACI fabrics, you must create a contract between the fabric's in-band management EPG and Cisco Nexus Dashboard EPG.
- For ACI fabrics, if several fabrics are monitored with apps on the Nexus Dashboard cluster, L3Out with default route or specific route to other ACI fabric in-band EPG must be provisioned and a contract must be established between the cluster EPG and the L3Out's external EPG.
- If you specify a VLAN ID for your data network during setup of the cluster, the Nexus Dashboard interface and the port on the connected network device must be configured as `trunk`

  However, in most cases we recommend not assigning a VLAN to the data network, in which case you must configure the ports in `access` mode.

The following two figures show two distinct network connectivity scenarios when connecting the Nexus Dashboard cluster directly to the fabrics' leaf switches. The primary purpose of each depends on the type of application you may be running in your Nexus Dashboard.

Note that the "L3 Network" and the "Management Network" can be the same network infrastructure, for example in case the Nexus Dashboard nodes have the management and data network interfaces in the same subnet.

Figure 4: Connecting Directly to Leaf Switches, Day-2 Operations Applications

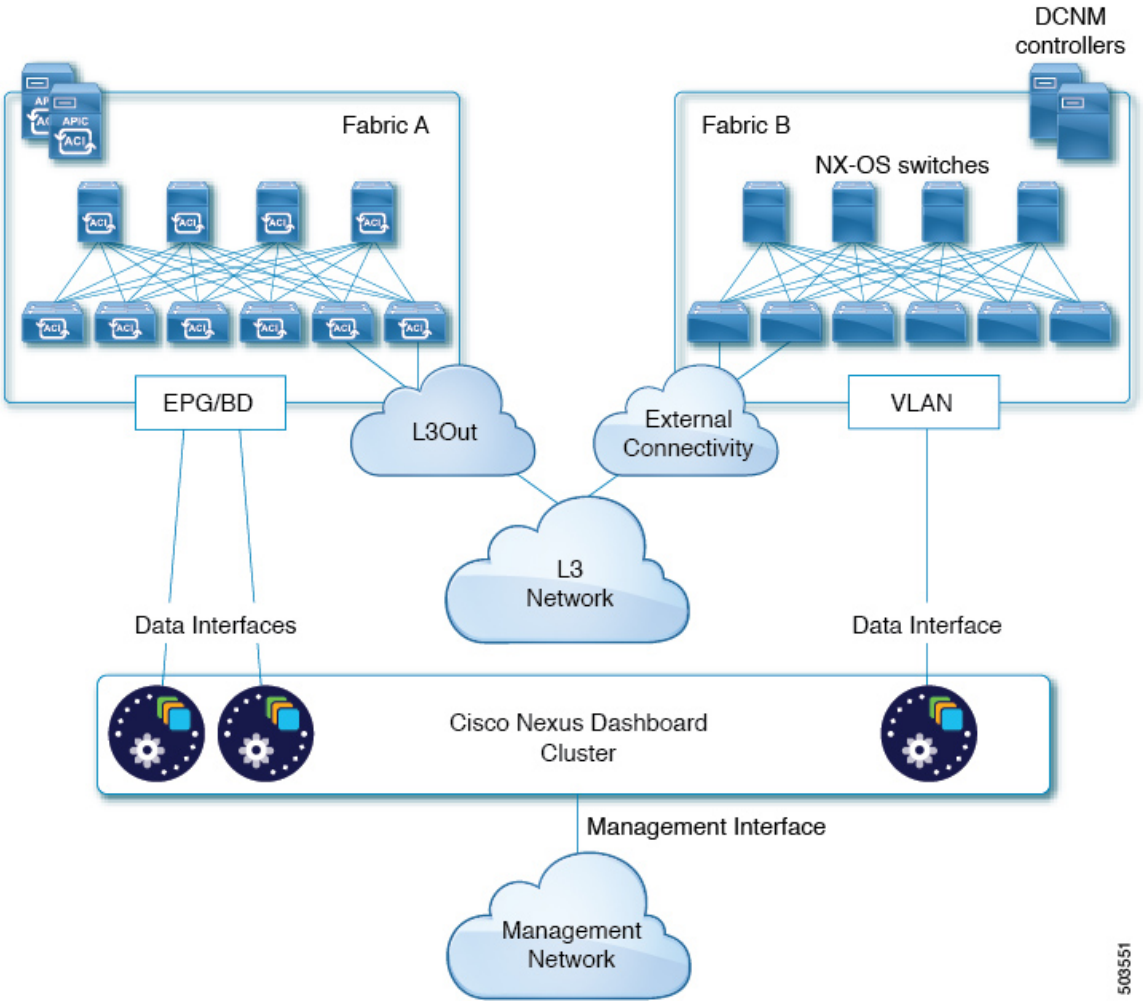

Figure 5: Connecting Directly to Leaf Switches, Multi-Site Orchestrator

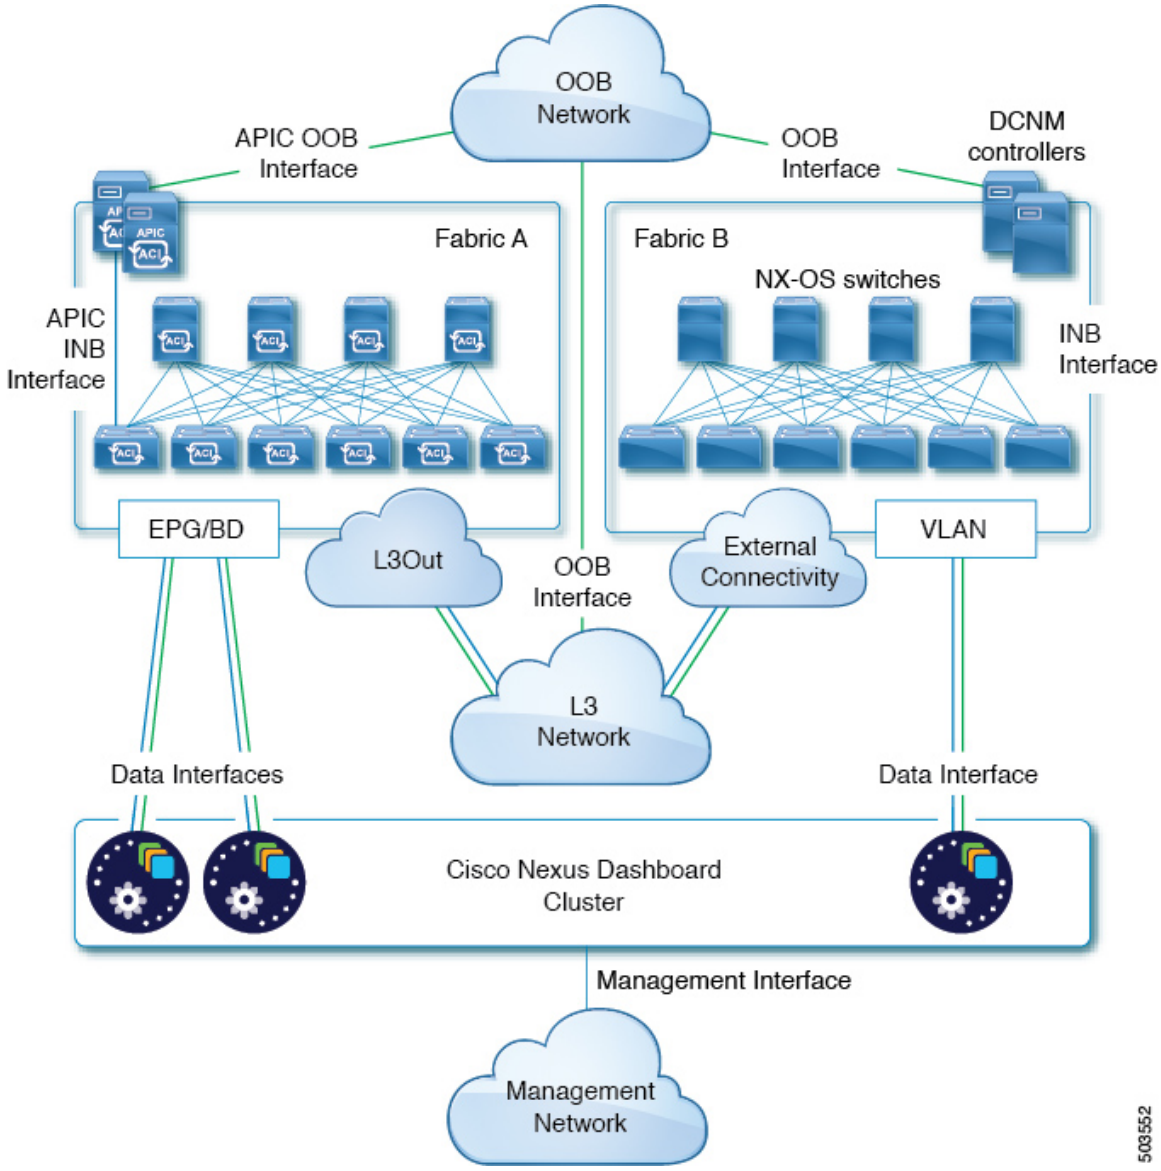

Node Distribution Across Sites

Nexus Dashboard supports distribution of cluster nodes across multiple sites.

We recommend centralized, single-site deployment for Nexus Insights and Network Assurance Engine applications. These applications do not gain redundancy benefits from distributed cluster, which could instead expose the cluster to interconnection failures when nodes are in different sites.

We recommend distributed cluster for Multi-Site Orchestrator deployments. Keep in mind that at least two Nexus Dashboard master nodes are required for the cluster to remain operational, so when deploying a physical Nexus Dashboard cluster across two sites, we recommend deploying a standby node in the site with the single master node as shown in the following figure:

Figure 6: Node Distribution Across Two Sites for Multi-Site Orchestrator

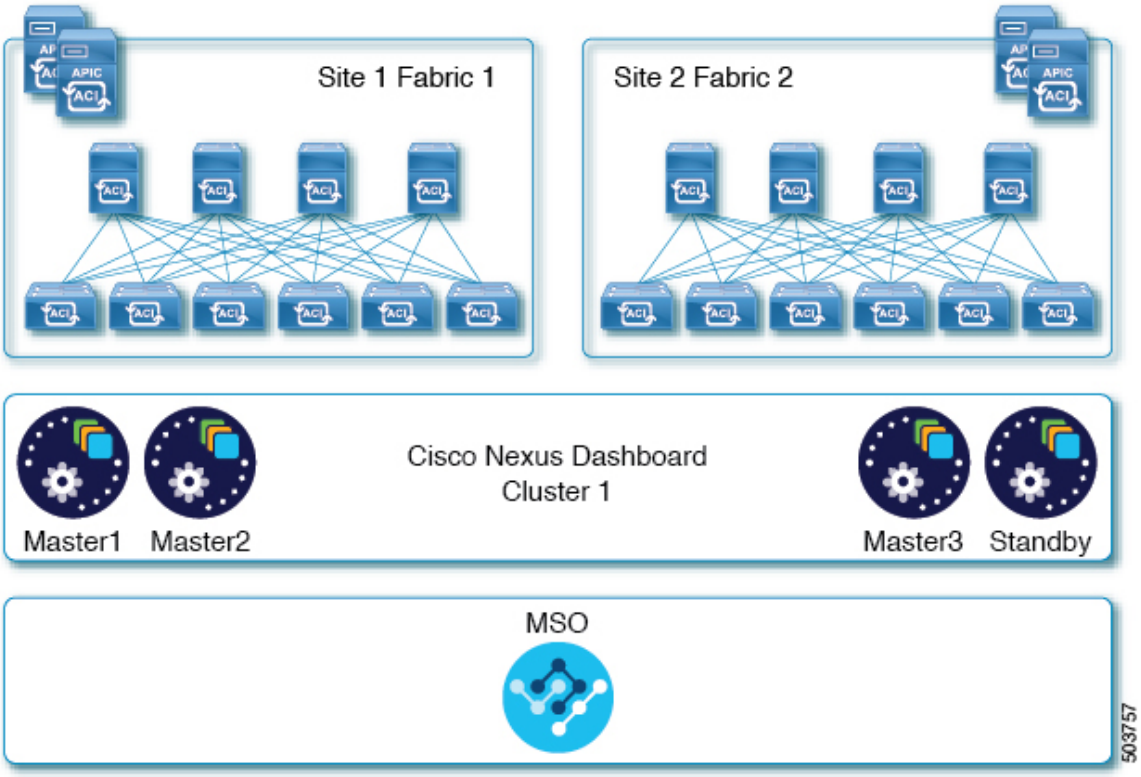

If you are deployed a virtual Nexus Dashboard cluster, standby nodes are not supported and if one of the nodes fails, you will need to bring up a new virtual node to replace it, as described in the "Replacing Virtual Nodes" chapter of the *Cisco Nexus Dashboard User Guide*.

The following table summarizes additional supported scenarios for distribution of physical Nexus Dashboard master (`M1`, `M2`, `M3`) and standby (`S1`) nodes across multiple sites:

Table 5: Nexus Dashboard Node Distribution Across Sites

Number of Sites	Nodes in Site 1	Nodes in Site 2	Nodes in Site 3	Nodes in Site 4
1	`M1`, `M2`, `M3`	--	--	--
2	`M1`, `M2`	`M3`, `S1`	--	--
3	`M1`	`M2`	`M3`	--
4	`M1`	`M2`	`M3`	`S1`

App Co-location Use Cases

This section describes a number of recommended deployment scenarios for specific single-app or multiple apps co-hosting use cases.

Note This release does not support application co-hosting for virtual or cloud form factors. All application co-hosting scenarios below apply for physical Nexus Dashboard clusters only.

Single Site, Nexus Insights and Network Assurance Engine

In a single site scenario with Nexus Insights and Network Assurance Engine applications, a single physical Nexus Dashboard cluster can be deployed with both applications co-hosted on it.

Figure 7: Single Site, Nexus Insights and Network Assurance Engine

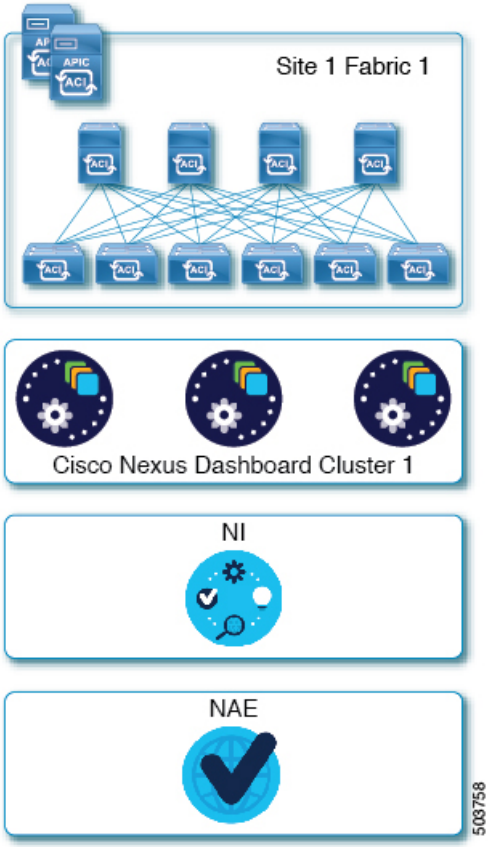

Multiple Sites, Nexus Insights and Network Assurance Engine

In a multiple sites scenario with Nexus Insights and Network Assurance Engine applications, a single Nexus Dashboard cluster can be deployed with both applications co-hosted on it. In this case, the nodes can be distributed between the sites, however since these applications do not gain redundancy benefits from a distributed cluster and could instead be exposed to interconnection failures when nodes are in different sites, we recommend the deployment option on the left:

Figure 8: Single Site, Nexus Insights and Network Assurance Engine

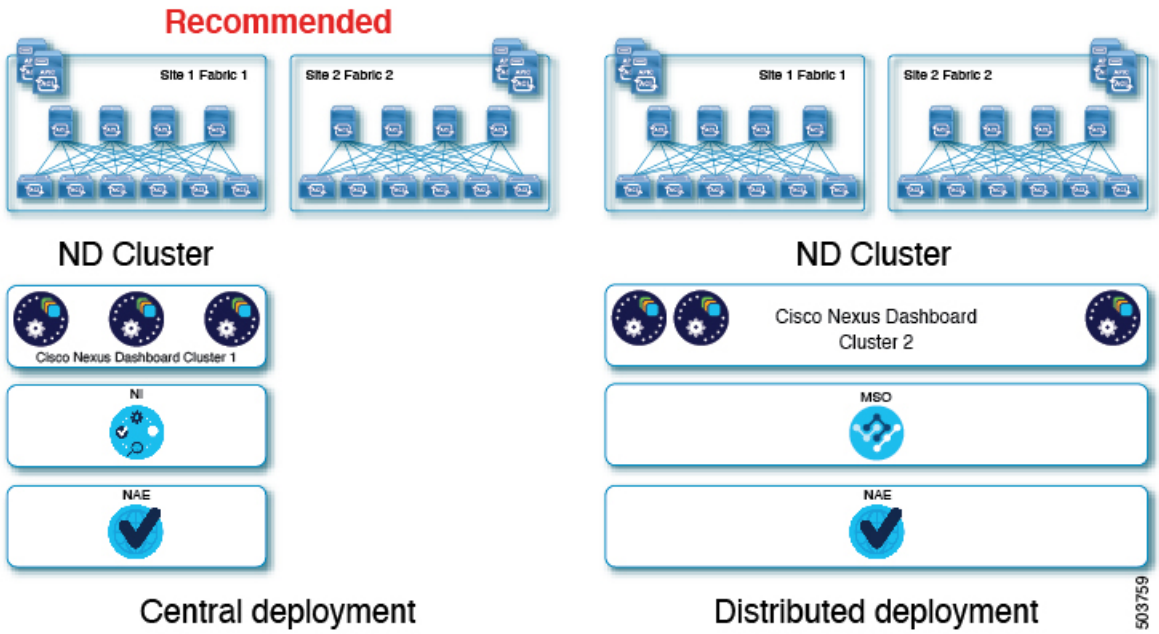

Multiple Sites, Nexus Insights, Network Assurance Engine, and Multi-Site Orchestrator

In this case, we recommend deploying two Nexus Dashboard cluster, with one of them dedicated to the Multi-Site Orchestrator application using the virtual or cloud form factor and the nodes distributed across the sites.

Figure 9: Single Site, Nexus Insights and Network Assurance Engine

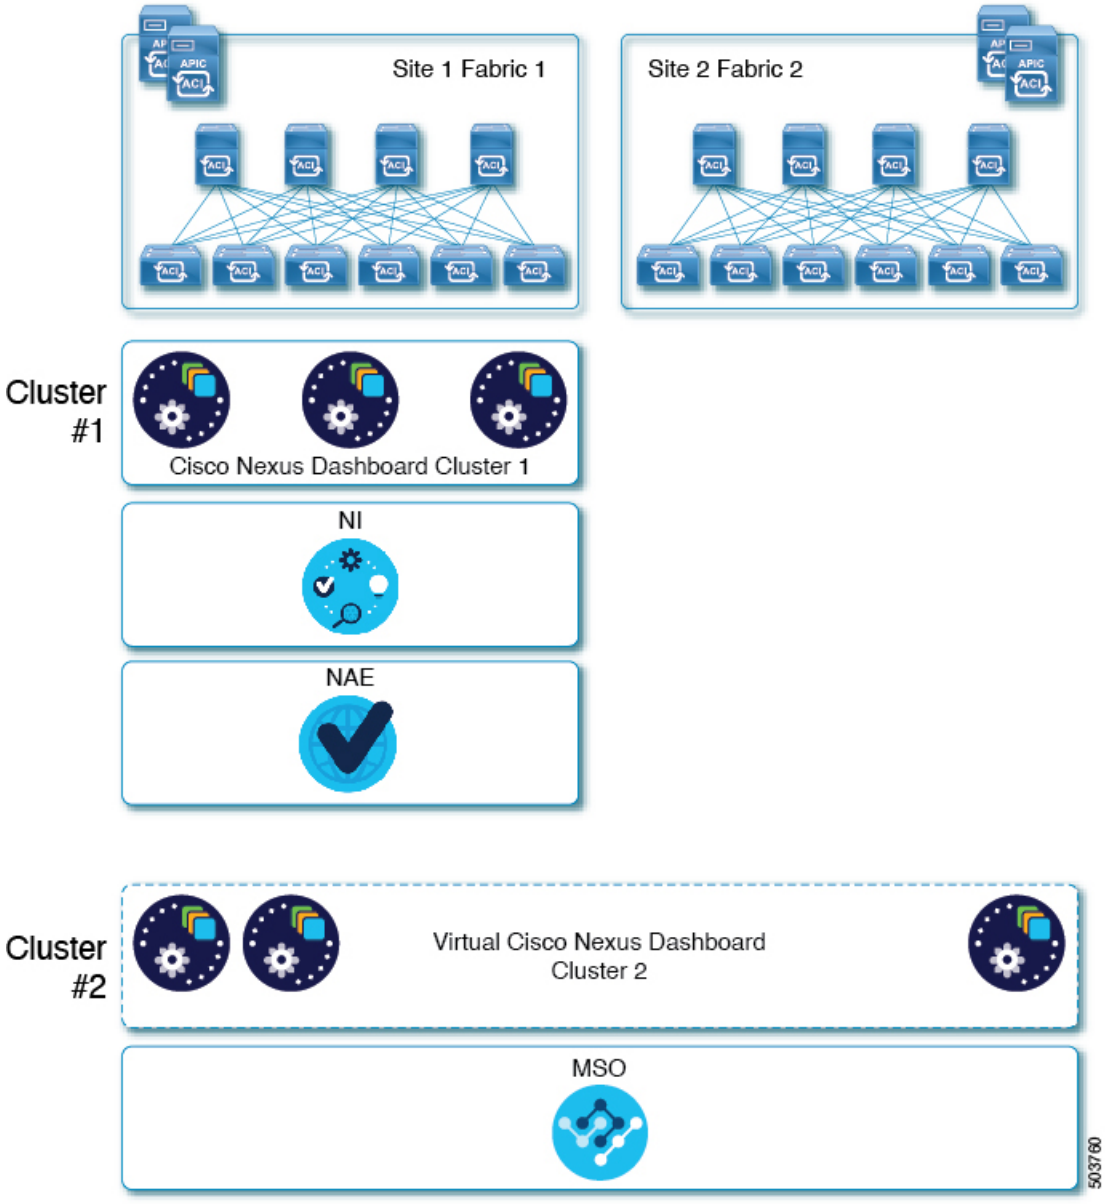

Pre-Installation Checklist

Before you proceed with deploying your Nexus Dashboard cluster, prepare the following information for easy reference during the process:

Table 6: Cluster Details

Parameters	Example	Your Entry
Cluster **Name**	`nd-cluster`	
NTP Server	`171.68.38.65`	
DNS Provider	`64.102.6.247 171.70.168.183`	
DNS Search Domain	`cisco.com`	
App Network	`172.17.0.1/16`	
Service Network	`100.80.0.0/16`	

Table 7: Node Details

Parameters	Example	Your Entry
For physical nodes, **CIMC** address and login information of the first node	`10.195.219.84/24` Username: `admin` Password: `Cisco1234`	
For physical nodes, **CIMC** address and login information of the second node	`10.195.219.85/24` Username: `admin` Password: `Cisco1234`	
For physical nodes, **CIMC** address and login information of the third node	`10.195.219.86/24` Username: `admin` Password: `Cisco1234`	
Password used for each node's `rescue-user` and the initial GUI password. We recommend configuring the same password for all nodes in the cluster.	`Welcome2Cisco!`	
Management IP of the first node	`192.168.9.172/24`	
Management Gateway of the first node.	`192.168.9.1`	
Data Network IP of the first node	`192.168.6.172/24`	
Data Network Gateway of the first node	`192.168.6.1`	
(Optional) **Data Network VLAN** of the first node	`101`	

Parameters	Example	Your Entry
Management IP of the second node	`192.168.9.173/24`	
Management Gateway of the second node.	`192.168.9.1`	
Data Network IP of the second node	`192.168.6.173/24`	
Data Network Gateway of the second node	`192.168.6.1`	
(Optional) **Data Network VLAN** of the second node	`101`	
Management IP of the third node	`192.168.9.174/24`	
Management Gateway of the third node.	`192.168.9.1`	
Data Network IP of the third node	`192.168.6.174/24`	
Data Network Gateway of the third node	`192.168.6.1`	
(Optional) **Data Network VLAN** of the third node	`101`	

CHAPTER 3

Deploying as Physical Appliance

Prerequisites and Guidelines

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

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.

  Note that this section describes how to initially deploy a three-node Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `worker` or `standby`), see the "Deploying Additional Nodes" section of the *Cisco Nexus Dashboard User Guide* instead.

  If you are looking to completely re-image the server, for example in case you cannot log in as the `rescue-user` for manual recovery, see the "Re-Imaging Nodes" section of the *Cisco Nexus Dashboard User Guide*.

  The guide is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

- Ensure you are using the correct hardware and the servers are racked and connected as described in *Cisco Nexus Dashboard Hardware Installation Guide*.

  The physical appliance form factor is supported on the original Nexus Dashboard platform hardware only. The following table lists the PIDs and specifications of the physical appliance servers:

Table 8: Supported Hardware

PID	Hardware
SE-NODE-G2	• UCS C220 M5 Chassis • 2x 10-core 2.2G Intel Xeon Silver CPU • 256 GB of RAM • 4x 2.4TB HDDs 400GB SSD 1.2TB NVME drive • UCS Virtual Interface Card 1455 (4x25G ports) • 1050W power supply
SE-CL-L3	A cluster of 3x `SE-NODE-G2` appliances.

Note The above hardware supports Nexus Dashboard software only. If any other operating system is installed, the node can no longer be used as a Nexus Dashboard node.

- Ensure that you are running a supported version of Cisco Integrated Management Controller (CIMC).

  Recommended version: CIMC, Release 4.1(3b).

  Minimum supported version: CIMC, Release 4.0(1a).

- Ensure that all nodes are running the same release version image.
- If your Nexus Dashboard hardware came with a different release image than the one you would like to deploy, we recommend deploying the cluster with the existing image first and then upgrading it to the desired release.

  For example, if the hardware you received came with Release 2.0.1 image pre-installed, but you want to deploy Release 2.0.2 instead, we recommend:

  - First, bring up the Release 2.0.1 cluster, as described in the following section.
  - Then upgrade to Release 2.0.2, as described in Upgrading Nexus Dashboard, on page 59.

You must have at least a 3-node cluster. Up to four additional worker nodes can be added for horizontal scaling if required by the type and number of applications you will deploy.

Deploying Cisco Nexus Dashboard as Physical Appliance

When you first receive the Nexus Dashboard physical hardware, it comes preloaded with the software image. This section describes how to configure and bring up the initial 3-node Nexus Dashboard cluster.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 23.

Step 1 Configure the first node's basic information.

You only need to complete the following configuration on the first node of the cluster. For the second and third master nodes, simply ensure that they are powered on and their CIMC IP address is reachable from the first node.

a) SSH into the node using CIMC management IP and use the `connect host` command to connect to the node's console.

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` CLI login as well as the initial GUI password.

```
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) 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

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

Step 2 Wait for the initial bootstrap process to complete.

After you provide and confirm management network information, the initial setup 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 3 Browse to the node's management IP address to open the GUI.

The rest of the configuration workflow takes place from the first node's GUI. 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**

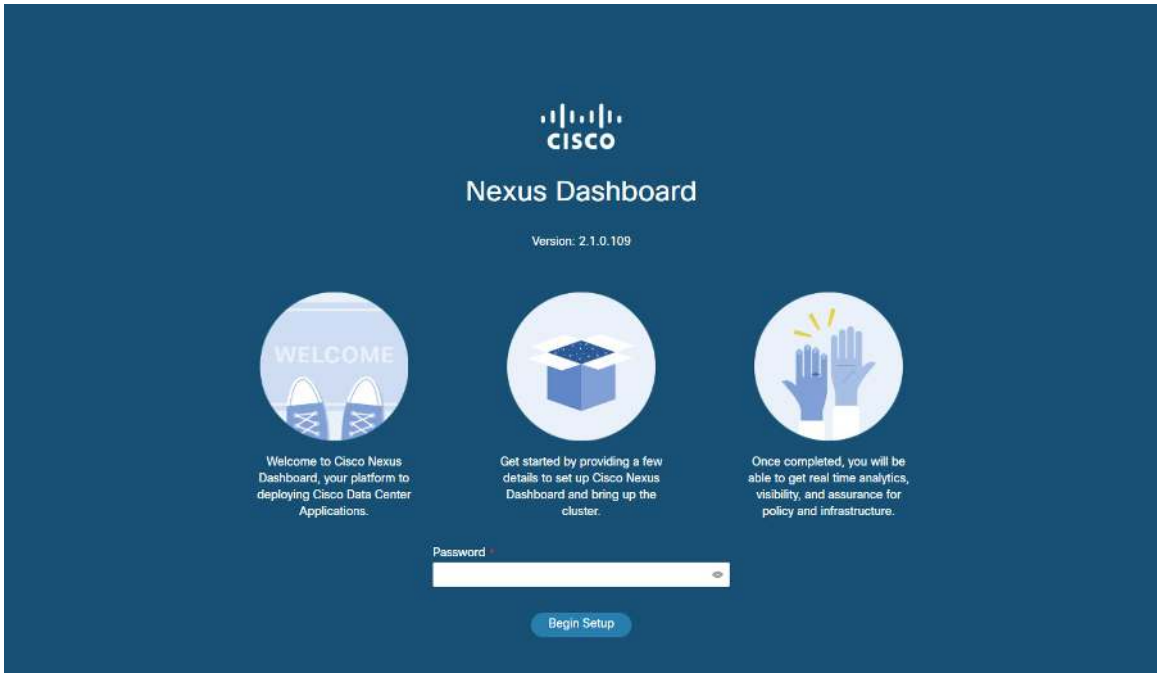

Step 4 In the **Cluster Details** screen, provide cluster information.

a) Provide the cluster **Name**.
b) Click **Add NTP Host** and provide the NTP server information.
c) Click **Add DNS Provider** and provide the DNS server information
d) (Optional) Expand **View Advanced Settings** menu and configure the DNS search domain and the internal networks (Application and Services).

Application and Services networks are described in the Prerequisites and Guidelines, on page 6 section earlier in this document.

e) Click **Next** to continue.

Step 5 In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.
b) Provide the node's **Data Network** IP address and gateway.

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.

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

Step 6 Click **Add Node** to add another node to the cluster.

The **Node Details** window opens.

a) Provide the node's CIMC details and click **Verify**.

The IP address and login information of the node's CIMC is used to pull that node's information, such as the serial number.

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

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.

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

Step 7 Repeat the previous step to add the 3rd node.

Step 8 Click **Next** to continue.

Step 9 In the **Confirmation** screen, review the entered 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.

It may take up to 20 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 10 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.

Step 11 If you deployed Release 2.0.2 and plan to host multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you deployed Release 2.0.1 or you are hosting only a single application in your Nexus Dashboard cluster, skip this step.

If you are co-hosting multiple applications in the same cluster, you must configure the App Infra Services with deployment profiles appropriate for your combination of applications and fabric sizes.

After the cluster upgrade is completed, follow the instructions described in the "App Infra Services" section of the *Cisco Nexus Dashboard User Guide*, which is also available in the products GUI.

CHAPTER 4

Deploying in VMware ESX

Prerequisites and Guidelines

Virtual deployments are supported starting with Nexus Dashboard, Release 2.0.2h. Earlier releases support only the physical form factor described in Deploying as Physical Appliance, on page 23.

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

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.

  Note that this document describes how to initially deploy a three-node Nexus Dashboard cluster. If you want to expand an existing cluster with additional nodes (such as `worker` or `standby`), see the "Deploying Additional Nodes" section of the *Cisco Nexus Dashboard User Guide* instead.

  The guide is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*

- Ensure that the ESX form factor supports your scale and application requirements.

  Scale and application 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.

- Ensure you have enough system resources:

Table 9: Deployment Requirements

Nexus Dashboard Version	Requirements

Release 2.0.2h
Earlier releases are not supported.
	• VMware vCenter 6.x
	• VMware ESXi 6.5 or 6.7
	• Each VM requires:
	  • 16 vCPUs
	  • 64 GB of RAM
	  • 500 GB disk
	• We recommend that each Nexus Dashboard node is deployed in a different ESXi server.

- 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.

ESX Host Network Connectivity

If you plan to install Nexus Dashboard Insights or Fabric Controller service and use the Persistent IPs feature, you must ensure that the ESX host where the cluster nodes are deployed has a single logical uplink. In other words, it is connected via a single link, PC, or vPC and not a dual Active/Active (`A/A`) or Active/Standby (`A/S`) link without PC/vPC.

The following diagrams summarize the supported and unsupported network connectivity configurations for the ESX host where the nodes are deployed:

- In case the ESX host is connected directly, the following configurations are supported:
  - `A/A` uplinks of Port-Group or virtual switch with PC or vPC
  - Single uplink of Port-Group or virtual switch
  - Port-Channel used for the uplink of Port-Group or virtual switch.

`A/A` or `A/S` uplinks of Port-Group or virtual switch without PC or vPC are not supported

Figure 10: ESX Host Connectivity (Direct)

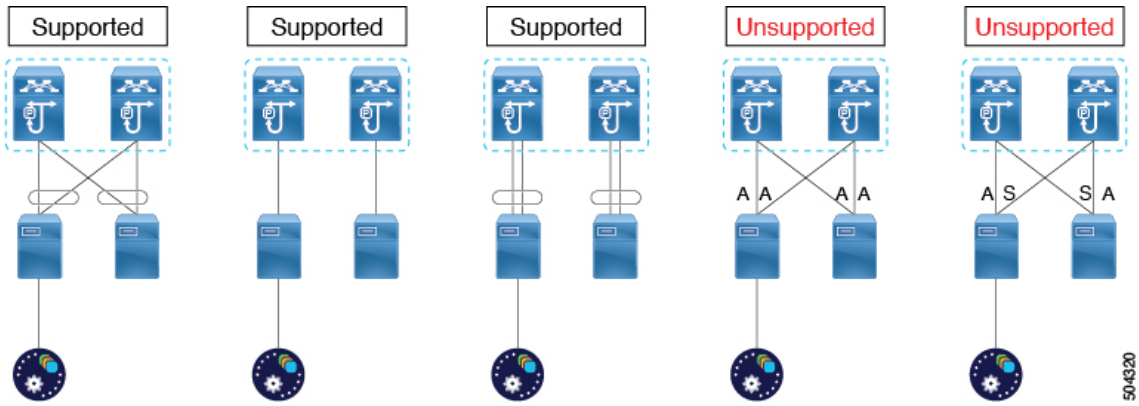

- In case the ESX host is connected via a UCS Fabric Interconnect (or equivalent), the following configurations are supported:
  - `A/S` uplinks of Port-Group or virtual switch at UCS Fabric Interconnect level without PC or vPC

    In this case, the `Active/Standby` links are based on the server technology, such as Fabric Failover for Cisco UCS and not at the ESXi hypervisor level.
  - Single uplink of Port-Group or virtual switch

`A/A` or `A/S` uplinks of Port-Group or virtual switch at the hypervisor level without PC or vPC are not supported

Figure 11: ESX Host Connectivity (with Fabric Interconnect)

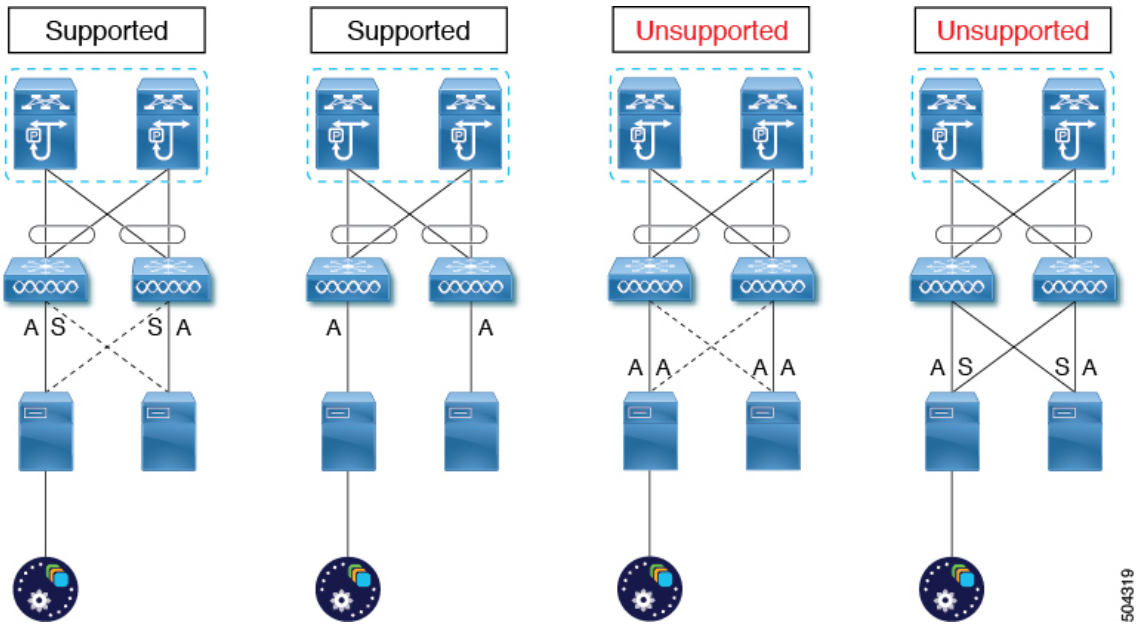

Deploying Cisco Nexus Dashboard in VMware ESX

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 29.

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html

b) Click the **Downloads** tab.

c) Choose the Nexus Dashboard version you want to download.

d) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.ova`).

Step 2 Log in to your VMware vCenter.

You cannot deploy the OVA directly in the ESX host, you must deploy it using the vCenter.

Note 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 6.7.

Step 3 Start the new VM deployment.

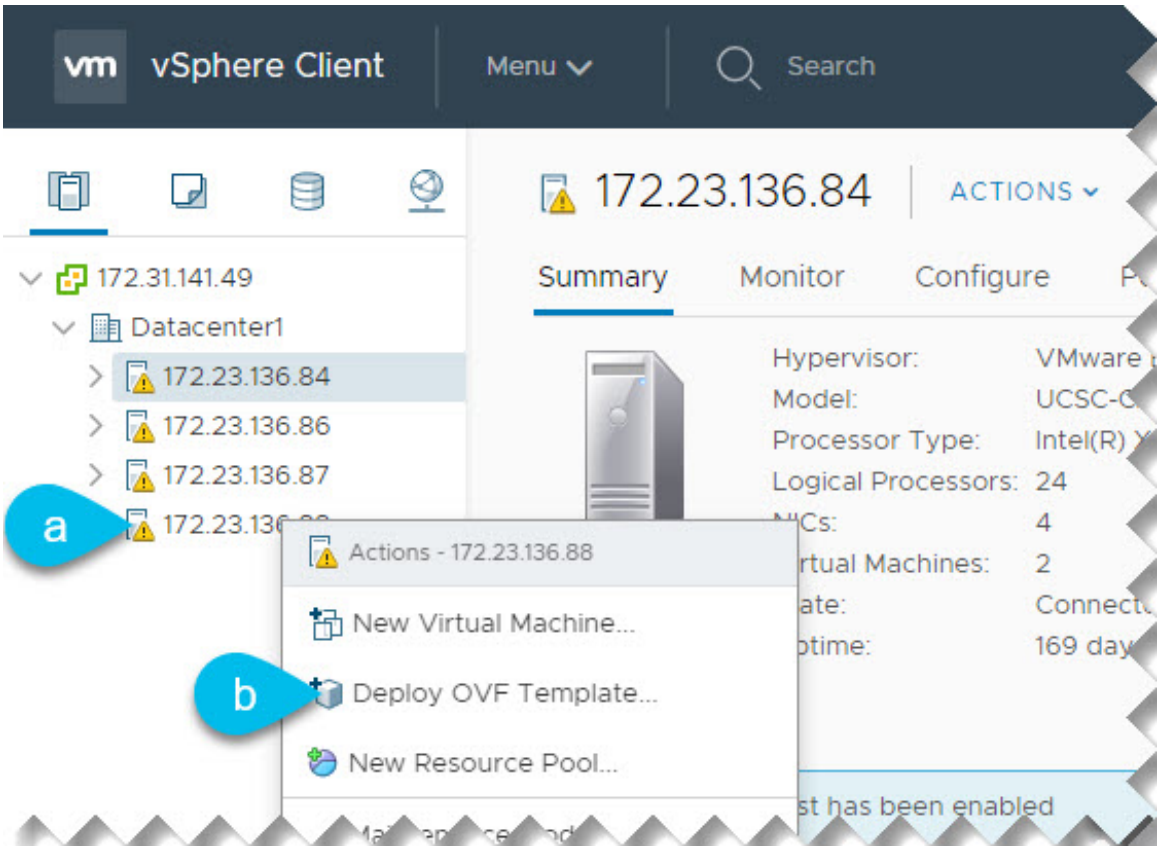

a) Right-click the ESX host where you want to deploy.
b) Then select **Deploy OVF Template...**

The **Deploy OVF Template** wizard appears.

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

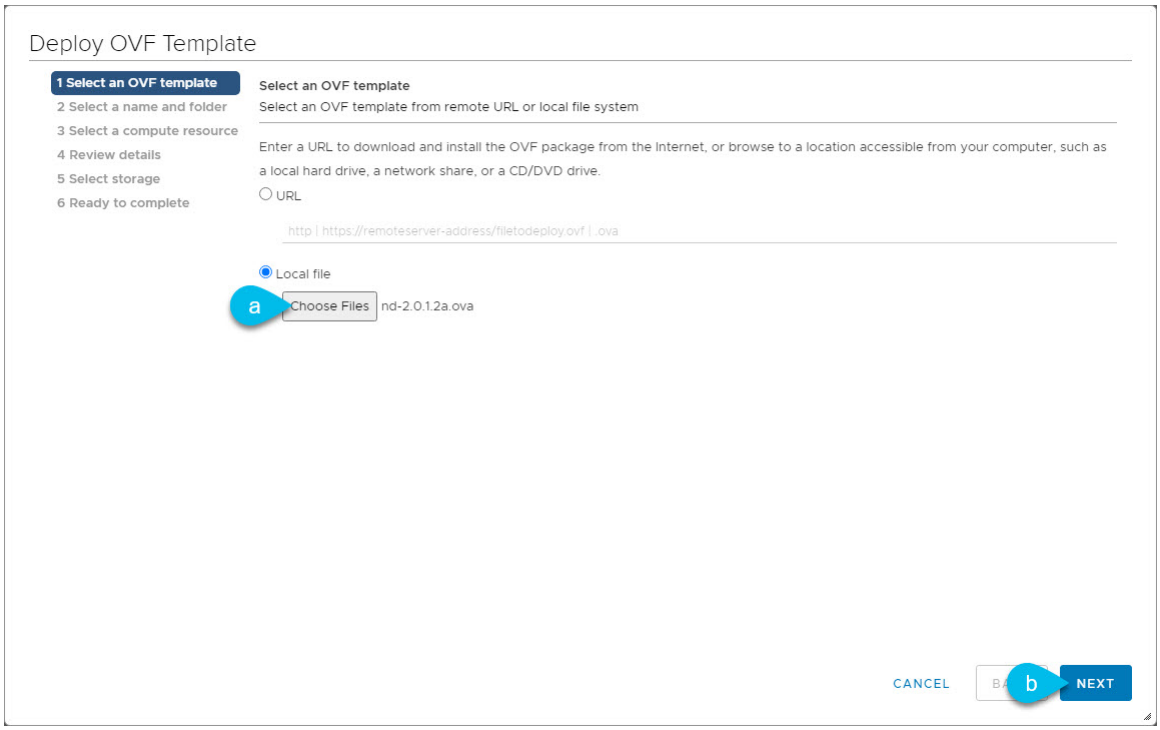

a) Select **Local file** and click **Choose Files** to select the OVA file you downloaded..
b) Click **Next** to continue.

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

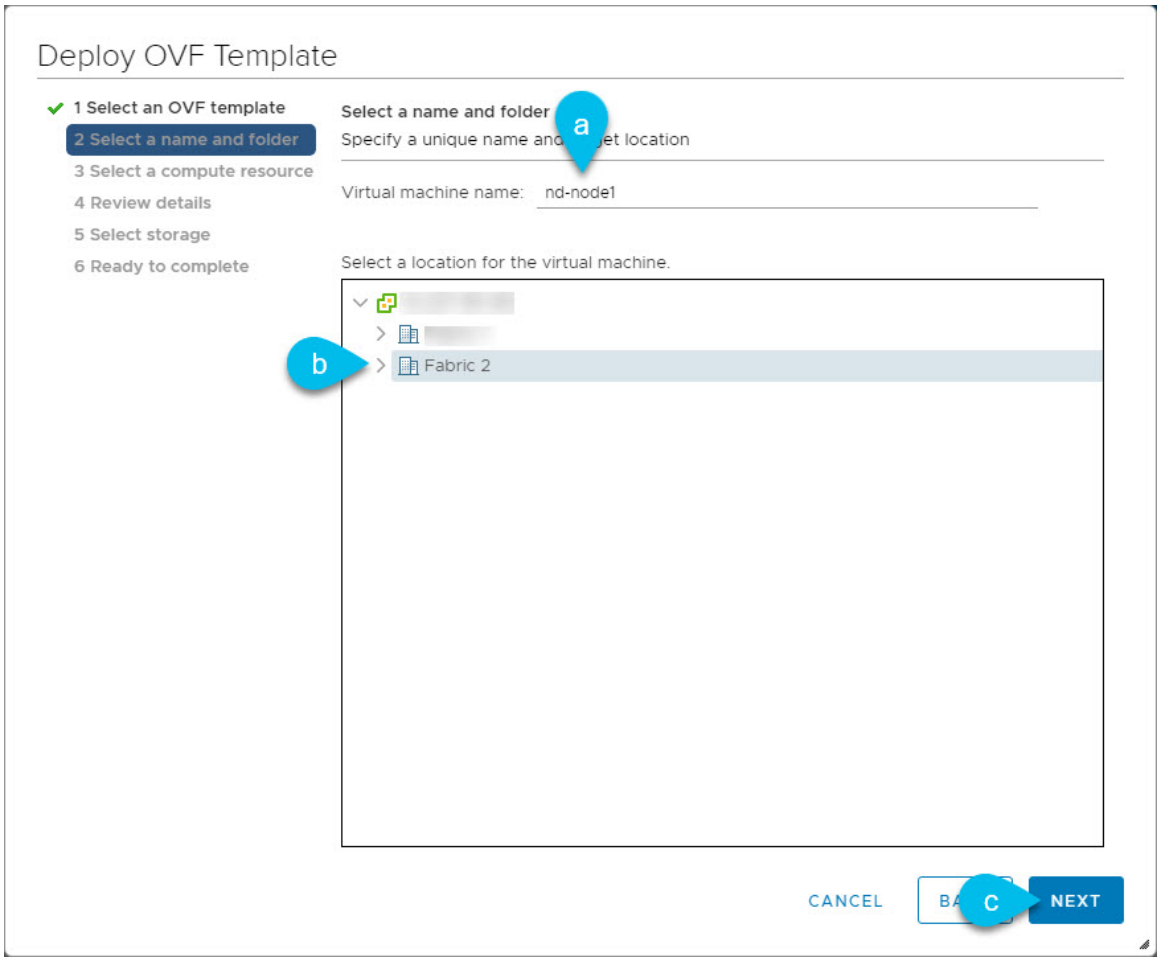

a) Provide the name for your virtual machine.
b) Select the location for the virtual machine.
c) Click **Next** to continue

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

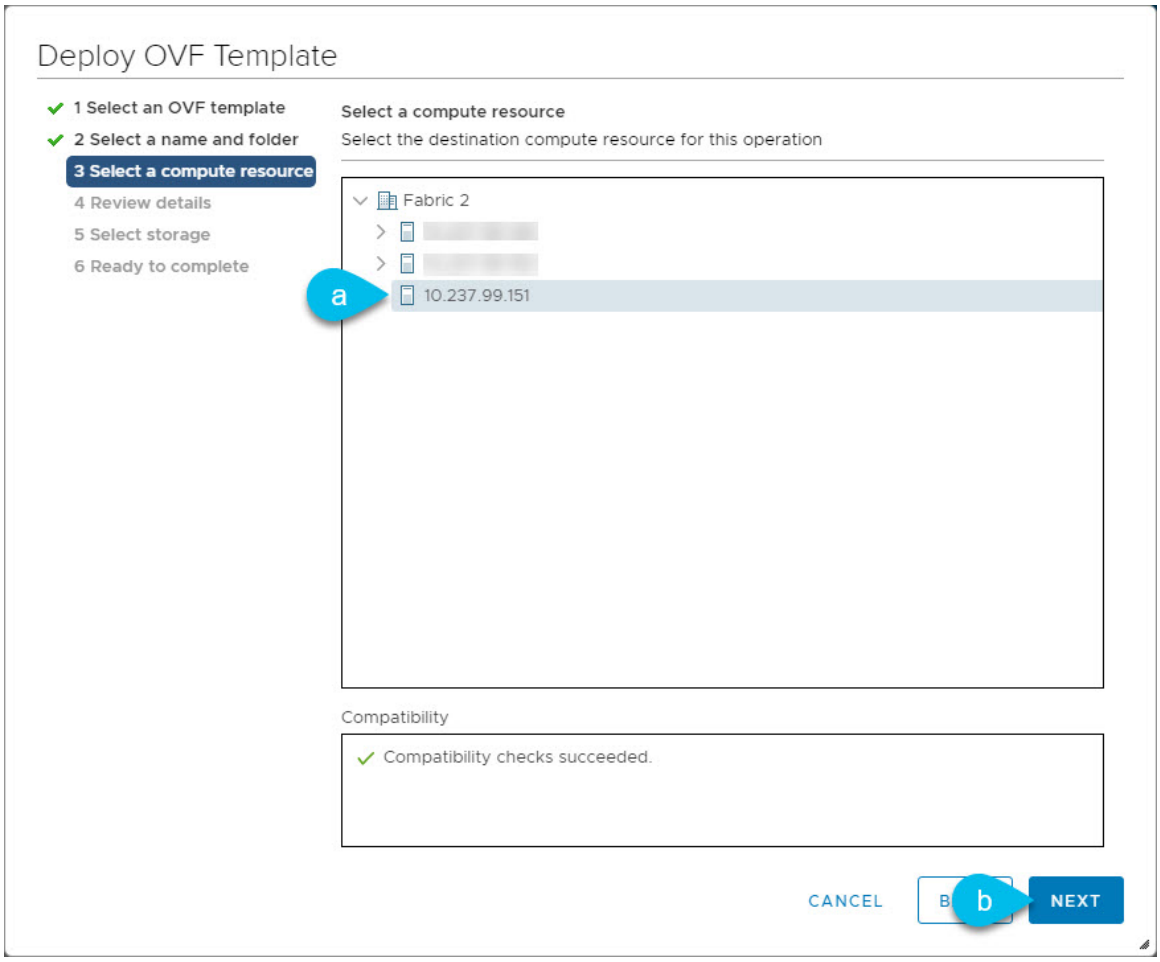

a) Select the vCenter datacenter and the ESX host for the virtual machine.
b) Click **Next** to continue

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

Step 8 In the **Select storage** screen, provide the storage information.

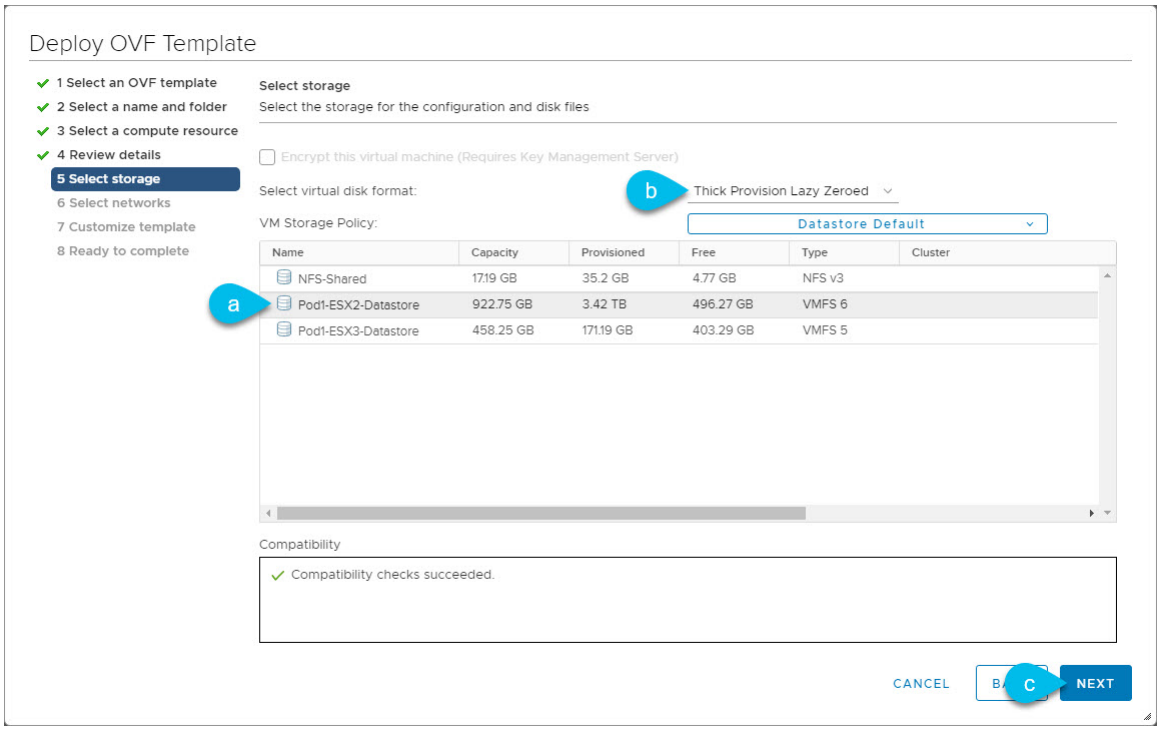

a) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

b) From the **Select virtual disk format** dropdown, select **Thick Provision Lazy Zeroed**.

c) Click **Next** to continue

Step 9 In the **Select networks** screen, accept default values and click **Next** to continue.

There are two networks, **fabric0** is used for the data network and **mgmt0** is used for the management network.

Step 10 In the **Customize template** screen, provide the required information.

Note The following few steps may be listed in different order depending on the version of the vSphere client you are using. The provided order and examples are using VMware vSphere 6.7.

In the **Resource Configuration** and **Node Configuration** categories, provide the following details:

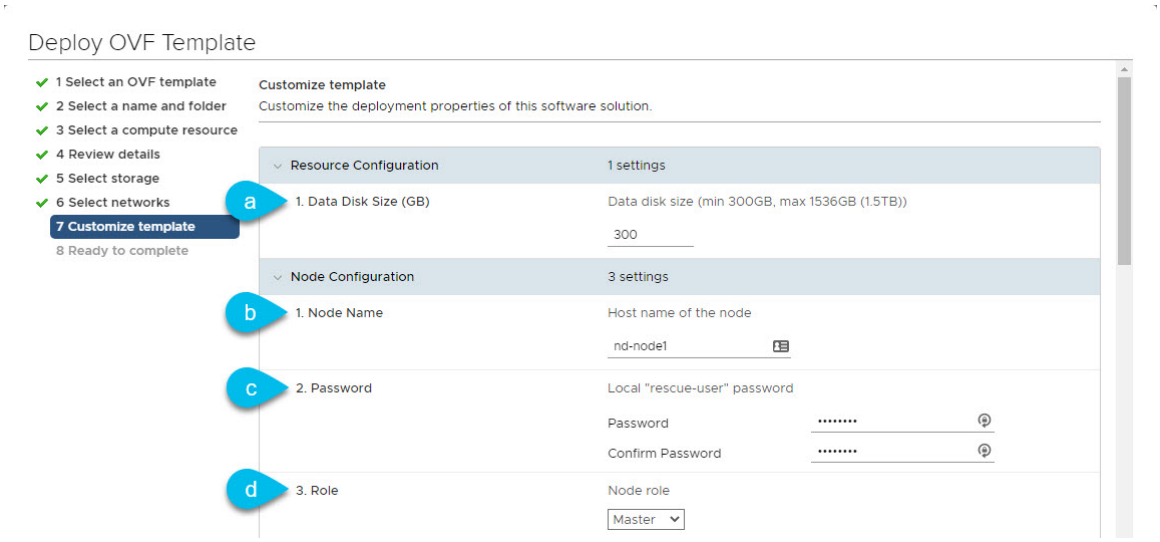

a) Provide the sizes for the node's data disks.

   We recommend using the default values for the required data disks.

b) Provide the **Node Name**.

   This will be the hostname for node, do not use the fully qualified domain name (FQDN).

   For example, `nd-node1`

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

   We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node. If you provide different passwords, the first node's password will be used as the initial password of the `admin` user in the GUI.

d) From the **Role** dropdown, select `Master`.

   When first deploying the cluster, all 3 nodes must be `Master`. Adding `Worker` and `Standby` nodes is described in the *Cisco Nexus Dashboard User Guide*.

In the **Network Configuration** category, provide the following details:

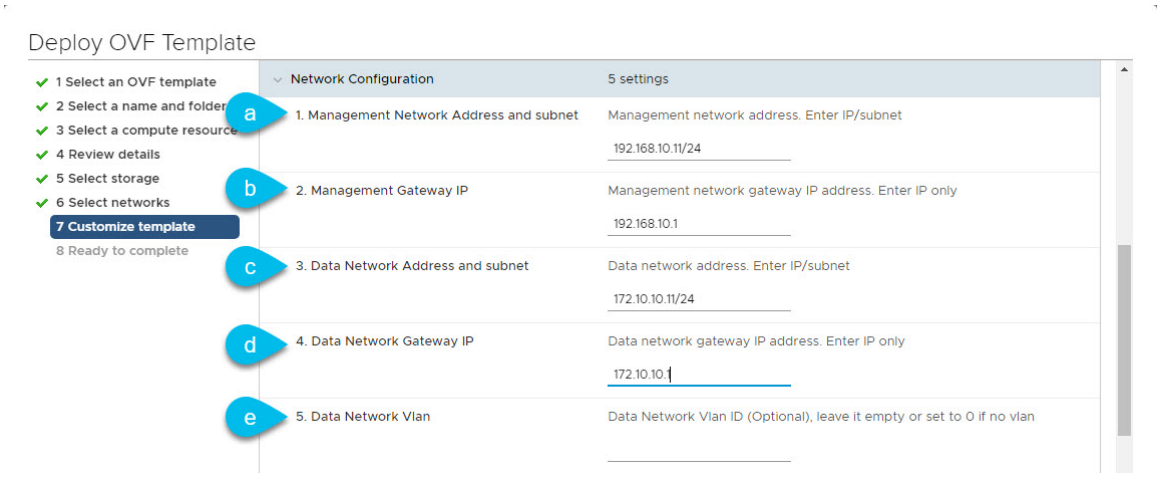

a) Provide the **Management Address and Subnet** for the node.

The management IP address can be in the same or different subnet as the data network IP address.

For example, `192.168.10.11/24`.

b) Provide the **Management Gateway IP**.

For example, `192.168.10.1`.

c) Provide the **Data Network Address and subnet**.

The data network IP address can be in the same or different subnet as the management IP address.

For example, `172.10.10.11/24`.

d) Provide the **Data Network Gateway**.

For example, `172.10.10.1`.

e) (Optional) If the data traffic is on a VLAN, provide the **Data Network Vlan**.

For most deployments, you can leave this field blank. If you do want to provide a VLAN ID for your data network, you can enter it in this field, for example `100`.

In the **Cluster Configuration Mandatory** and **Cluster Configuration Optional** categories, provide the following details:

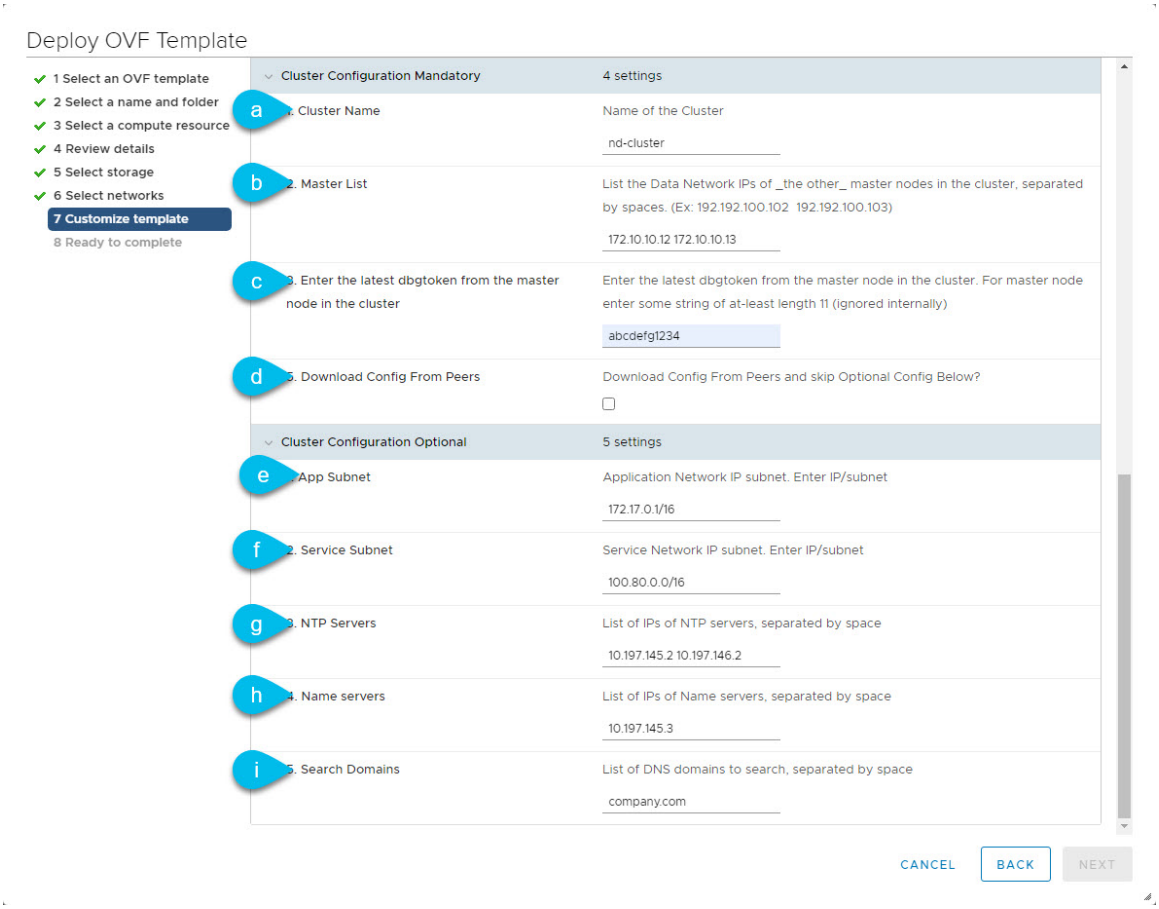

a) Provide the **Cluster Name** for the Nexus Dashboard cluster.

This name must be the same for all nodes.

For example, `nd-cluster`.

b) In the **Master List** field, provide the data network IP addresses of the other 2 nodes you will configure for your cluster.

Each IP address in the list must be separated by a space.

For example, if the data network IP addresses of all 3 nodes are `172.10.10.11`, `172.10.10.12`, and `172.10.10.13`, the value of this field for the first node would be `172.10.10.12 172.10.10.13`

c) Provide a value for the **dbgtoken** field.

Since this is the first node you are deploying, provide any 11-character value for this field (for example, `abcdef12345`). When you deploy the other two nodes, you will use this field to provide a token from the first node to simplify configuration.

d) Leave the **Download Config From Peers** checkbox unchecked.

You will use this option when configuring the other two nodes.

e) Provide the **App Subnet**.

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.

f) Provide the **Services Subnet**.

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.

g) Provide the **NTP Servers** information.

For example, `10.197.145.2 10.197.146.2`.

h) Provide the **Name servers** information.

For example, `10.197.145.3`.

i) (Optional) Provide the **Search Domains** information.

For example, `company.com`.

Step 11 Verify that all information is valid and click **Next** to continue.

After you complete the **Customize template** screen, a verification banner is shown at the top.

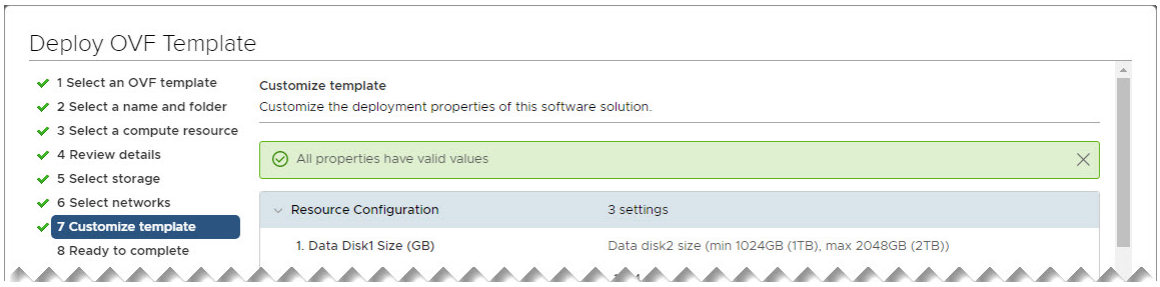

Step 12 In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the first node.

Step 13 Wait for the VM deployment to complete, ensure that the VMware Tools periodic time synchronization is disabled, then start the VM.

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 14 Log in to the first node's console as the `rescue-user`.

Use the password you specified in the OVF template when deploying the VM.

Step 15 Retrieve the `dbgtoken`.

Run the following command:

```
$ acs debug-token
09GZ1PMB8CML
```

Make a note of this token, you will use it to deploy the other two nodes.

Keep in mind, the token expires and is refreshed every 30 minutes, so ensure to retrieve it when ready to deploy the second and third nodes.

Step 16 Deploy the second node.

The steps to deploy the second and third nodes are similar, with the exception that you can now use the `dbgtoken` from the first node to skip some of the configuration.

a) Repeat Steps 2 through 9 to start deploying the 2nd node.

We recommend using a different ESX host for each node.

b) In the **Cluster Configuration** screen, provide the following information:

- **Node Name**

  Do not use the fully qualified domain name (FQDN).

- **Password**

  We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node. If you provide different passwords, the first node's password will be used as the initial password of the `admin` user in the GUI.

- **Role**

  When first deploying the cluster, all 3 nodes must be `Master`.

- **Management Network Address and subnet**
- **Management Gateway IP**
- **Data Network Address and subnet**
- **Data Network Gateway**
- (Optional) If the data traffic is on a VLAN, provide the **Data Network Vlan**.

- **Cluster Name**

  This name must be the same for all nodes. For example, `nd-cluster`.

- **Master List**

  Provide the data network IP addresses of the other 2 nodes in your cluster separated by a space.

  For example, if the data network IP addresses of all 3 nodes are `172.10.10.11`, `172.10.10.12`, and `172.10.10.13`, the value of this field for the second node would be `172.10.10.11 172.10.10.13`

- Provide the **dbgtoken** you obtained from the first node.

  The token expires and is refreshed every 30 minutes, ensure to obtain the latest valid token from the first node before continuing. For example, `09GZ1PMB8CML`.

- Check the **Download Config From Peers**

  The second and third nodes will download common configuration parameters from the first node using the `dbgtoken`.

c) Skip **Cluster Configuration Optional** fields and click **Next** to continue.

d) In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the second node.

Step 17 Repeat the previous step to deploy the third node.

Step 18 Wait for the second and third node VMs deployment to complete, then start the VMs.

Step 19 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.

CHAPTER 5

Deploying in Amazon Web Services

Prerequisites and Guidelines

Cloud deployments are supported starting with Nexus Dashboard, Release 2.0.2b. Earlier releases support only the physical form factor described in Deploying as Physical Appliance, on page 23.

Before you proceed with deploying the Nexus Dashboard cluster in Amazon Web Services (AWS), you must:

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.
- Ensure that the AWS form factor supports your scale and application requirements.

  Scale and application co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.
- Have appropriate access privileges for your AWS account.

  You must be able to launch multiple instances of Elastic Compute Cloud (`m5.2xlarge`) to host the Nexus Dashboard cluster.
- Have at least 6 AWS Elastic IP addresses.

  A typical Nexus Dashboard deployment consists of 3 nodes with each node requiring 2 AWS Elastic IP addresses for the management and data networks.

  By default, your AWS account has lower elastic IP limit, so you may need to request an increase. To request IP limit increase:

  1. In your AWS console, navigate to **Computer** > **EC2**.
  2. In the EC2 Dashboard, click **Network & Security** > **Elastic IPs** and note how many Elastic IPs are already being used.
  3. In the EC2 Dashboard, click **Limits** and note the maximum number of **EC2-VPC Elastic IPs** allowed.

     Subtract the number of IPs already being used from the limit to get. Then if necessary, click **Request limit increase** to request additional Elastic IPs.
- Create a Virtual Private Cloud (VPC).

A VPC is an isolated portion of the AWS cloud for AWS objects, such as Amazon EC2 instances. To create a VPC:

1. In your AWS console, navigate to **Networking & Content Delivery Tools** > **VPC**.
2. In the VPC Dashboard, click **Your VPCs** and choose **Create VPC**. Then provide the **Name Tag** and **IPv4 CIDR block**.

   The CIDR block is a range of IPv4 addresses for your VPC and must be in the `/16` to `/24` range. For example, `10.9.0.0/16`.

- Create an Internet Gateway and attach it to the VPC.

  Internet Gateway is a virtual router that allows a VPC to connect to the Internet. To create an Internet Gateway:

  - In the VPC Dashboard, click **Internet Gateways** and choose **Create internet gateway**. Then provide the **Name Tag**.
  - In the **Internet Gateways** screen, select the Internet Gateway you created, then choose **Actions** > **Attach to VPC**. Finally, from the **Available VPCs** dropdown, select the VPC you created and click **Attach internet gateway**.

- Create a routes table.

  Routes table is used for connecting the subnets within your VPC and Internet Gateway to your Nexus Dashboard cluster. To create a routes table:

  - In the VPC Dashboard, click **Route Tables**, choose the **Routes** tab, and click **Edit routes**.
  - In the **Edit routes** screen, click **Add route** and create a `0.0.0.0/0` destination. From the **Target** dropdown, select `Internet Gateway` and choose the gateway you created. Finally, click **Save routes**.

- Create a key pair.

  A key pair consists of a private key and a public key, which are used as security credentials to verify your identity when connecting to an EC2 instance. To create a key pair:

  - Navigate to **All services** > **Compute** > **EC2**.
  - In the EC2 Dashboard, click **Network & Security** > **Key Pairs**. Then click **Create Key Pairs**.
  - Provide a name for your key pair, select the **pem** file format, and click **Create key pair**.

    This will download the `.pem` private key file to your system. Move the file to a safe location, you will need to use it the first time you log in to an EC2 instance's console.

  By default only PEM-based login is enabled for each node. If you'd like to be able to SSH into the nodes using a password, you will need to explicitly enable password-based logins. You can do that by logging into each node separately using the PEM file the first time, then executing the following command:

  ```
  # acs login prompt-enable
  ```

Deploying the Cisco Nexus Dashboard in AWS

This section describes how to deploy Cisco Nexus Dashboard cluster in Amazon Web Services (AWS).

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 45.

Step 1 Subscribe to Cisco Nexus Dashboard product in AWS Marketplace.

a) Log into your AWS account and navigate to the AWS Management Console

The Management Console is available at https://console.aws.amazon.com/.

b) Navigate to **Services** > **AWS Marketplace Subscriptions**.
c) Click **Manage Subscriptions**.
d) Click **Discover products**.
e) Search for **Cisco Nexus Dashboard** and click the result.
f) In the product page, click **Continue to Subscribe**.
g) Click **Accept Terms**.

It may take a couple of minutes for the subscription to be processed.

h) Finally click **Continue to Configuration**.

Step 2 Select software options and region.

a) From the **Delivery Method** dropdown, select `Cisco Nexus Dashboard for Cloud`.
b) From the **Software Version** dropdown, select the version you want to deploy.
c) From the **Region** dropdown, select the regions where the template will be deployed.

This must be the same region where you created your VPC.

d) Click **Continue to Launch**.

The product page appears, which shows a summary of your configuration and enables you to launch the cloud formation template.

Step 3 From the **Choose Action**, select `Launch CloudFormation` and click **Launch**.

The **Create stack** page appears.

Step 4 Create stack.

a) In the **Prerequisite - Prepare template** area, select `Template is ready`.
b) In the **Specify Template** area, select `Amazon S3 URL` for the template source.

The template will be populated automatically.

c) Click **Next** to continue.

The **Specify stack details** page appears.

Step 5 Specify stack details.

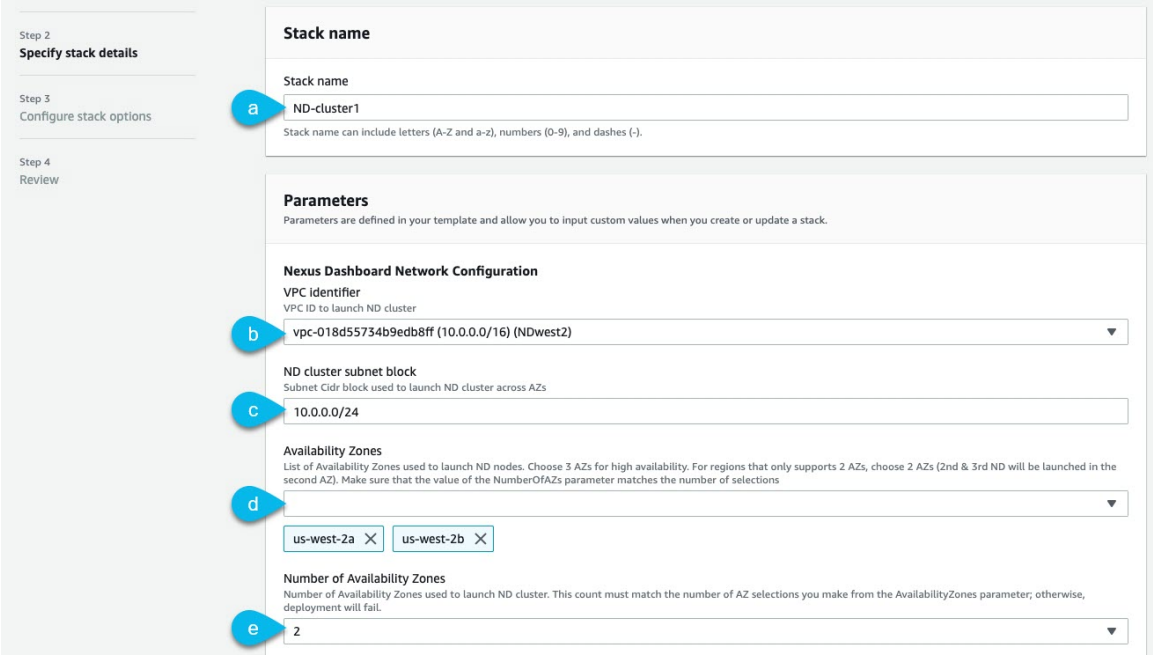

a) Provide the **Stack name**.

b) From the **VPC identifier** dropdown, select the VPC you created.

For example, `vpc-038f83026b6a48e98(10.176.176.0/24).`

c) In the **ND cluster Subnet block**, provide the VPC subnet CIDR block.

Choose a subnet from the VPC CIDR that you defined. You can provide a smaller subnet or use the whole CIDR.

For example, `10.176.176.0/24`.

d) From the **Availability Zones** dropdown, select one or more available zones.

We recommend you choose 3 availability zones. For regions that support only 2 availability zones, 2nd and 3rd nodes of the cluster will launch in the second availability zone.

e) From the **Number of Availability Zones** dropdown, select the number of zones you added in the previous substep.

Ensure that the number matches the number of availability zones you selected in the previous substep.

Provide the rest of the node information.

a) Enable **Data Interface EIP support**.

   This field enables external connectivity for the node. External connectivity is required for communication with Cisco ACI fabrics outside AWS.

b) From the **Instance type**, select `m5.2xlarge`
c) Provide the **Cluster name**.

   The cluster name must be the same across all nodes you deploy.

d) Provide the **Host name** prefix.

   The template will deploy a 3-node cluster with each node using the **Host name** prefix and appending `-1`, `-2`, and `-3` to create unique host names for each node.

e) Provide the **NTP servers** information.
f) Provide the **Name servers** information.
g) (Optional) Provide the **DNS search domains list**.
h) Provide the **Application IP subnet**.

   For example, `10.101.0.0/16`.

i) Provide the **Service IP subnet**.

   The services network is an internal network used by the Nexus Dashboard and its processes.

   For example, `10.102.0.0/16`.

Finally, provide the login and access information.

a) In the **Password** fields, provide the password.

This password will be used for the Nexus Dashboard's `rescue-user` login, as well as the initial password for the GUI's `admin` user.

b) From the **SSH key pair** dropdown, select the key pair you created.

c) In the **Access control** field, provide the external network allowed to access the cluster.

For example, `0.0.0.0/0` to be able to access the cluster from anywhere.

d) Click **Next** to continue.

Step 6 In the **Advanced options** screen, simply click **Next**.

Step 7 In the **Review** screen, verify template configuration and click **Create stack**.

Step 8 Wait for the instance deployment to complete, then start the instance.

You can view the status of the instance deployment in the **CloudFormation** page, for example `CREATE_IN_PROGRESS`. You can click the refresh button in the top right corner of the page to update the status.

When the status changes to `CREATE_COMPLETE`, you can proceed to the next step.

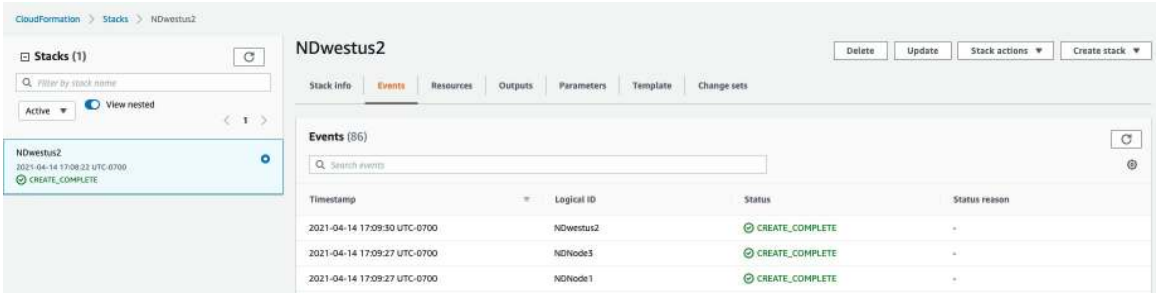

Step 9 Verify that the cluster is healthy.

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

After all three nodes' status is `CREATE_COMPLETE`, proceed with the following substeps to verify cluster health.

a) Verify that the AWS EC2 instances are up and running.

Navigate to **Services** > **EC2**. Then confirm that the **Status Checks** tab displays `2/2` checks.

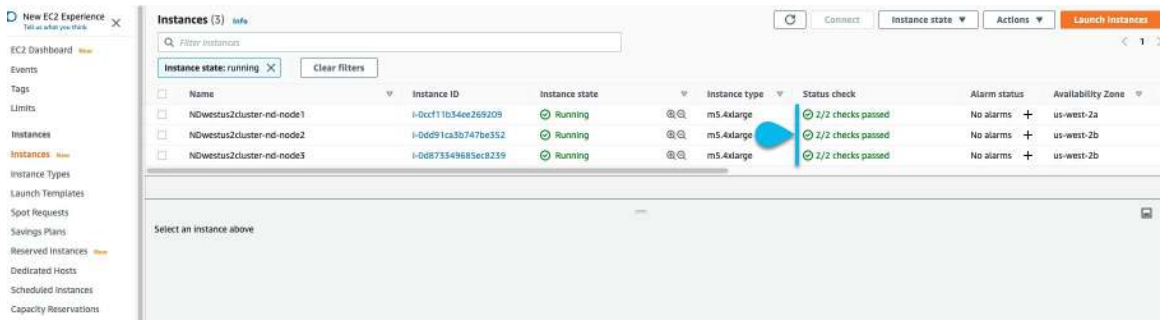

b) Login in to one of the nodes.

You will need to use the private key `.pem` file you downloaded when creating a key pair in the following command:

```
$ ssh -i <pem-file-name>.pem rescue-user@<node-ip-address>
```

c) 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
```

d) 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.

When you first log in, you will be prompted to change the password.

Step 10 (Optional) Enable password-based login.

By default only PEM-based login is enabled for each node. If you'd like to be able to SSH into the nodes using a password, you will need to explicitly enable password-based logins. You can do that by logging into each node separately using the PEM file the first time, then executing the following command:

```
# acs login-prompt enable
```

CHAPTER 6

Deploying in Microsoft Azure

Prerequisites and Guidelines

Cloud deployments are supported starting with Nexus Dashboard, Release 2.0.2b. Earlier releases support only the physical form factor described in Deploying as Physical Appliance, on page 23.

Before you proceed with deploying the Nexus Dashboard cluster in Microsoft Azure, you must:

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.
- Ensure that the Azure form factor supports your scale and application requirements.

  Scale and application co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.
- Have appropriate access privileges for your Azure account and subscription.
- Create an SSH key pair.

  A key pair consists of a private key and a public key, which are used as security credentials to verify your identity when connecting to Nexus Dashboard VMs. You will be asked to provide the public key when creating the Nexus Dashboard nodes.

  You can use an external utility, such as `putty`, to generate a key pair for your cluster.

Deploying the Cisco Nexus Dashboard in Azure

This section describes how to deploy Cisco Nexus Dashboard cluster in Microsoft Azure.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 53.

Step 1 Subscribe to Cisco Nexus Dashboard product in Azure Marketplace.

a) Log into your Azure account and browse to https://azuremarketplace.microsoft.com
b) In the search field, type `Cisco Nexus Dashboard` and select the option that is presented.

You will be re-directed to the Nexus Dashboard Azure Marketplace page.

c) Click **Get it now**.
d) In the **Select a plan** dropdown, select the version and click **Create**.

Step 2 In the **Basics** tab, provide the subscription details, region, and password.

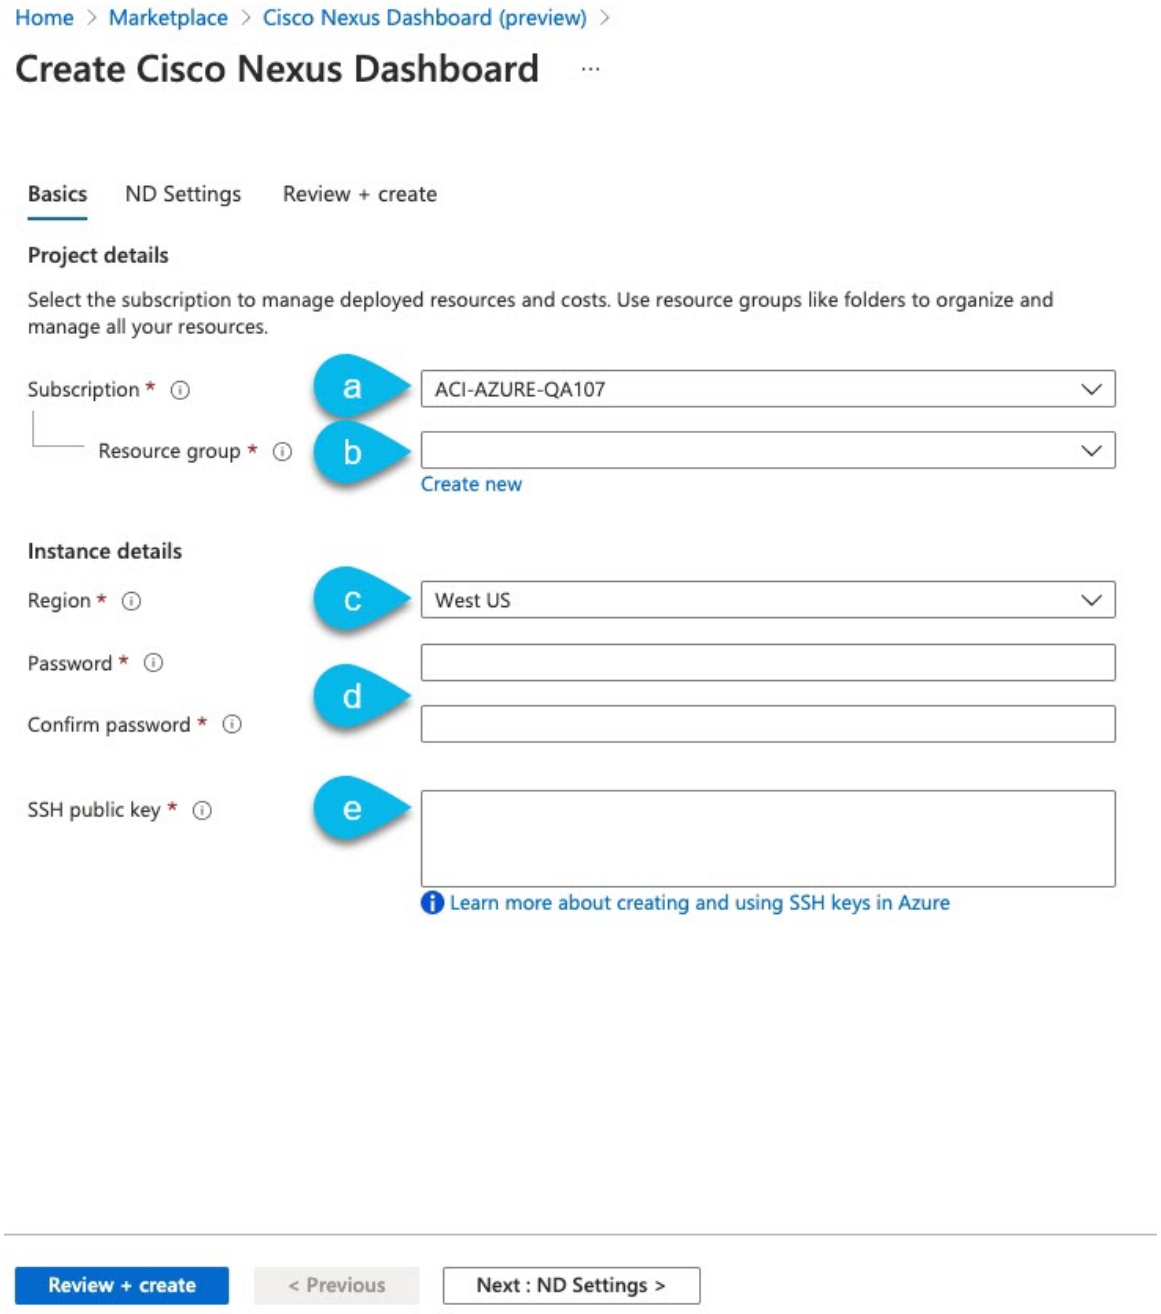

a) From the **Subscription** dropdown, select the subscription you want to use for this.

b) From the **Resource group** dropdown, select an existing resource group for the cluster or click **Create new** to create one.

c) From the **Region** dropdown, select the regions where the template will be deployed.

This must be the same region where you created your resource group and VNET.

d) Provide and confirm the node password.

This is the same password that will be used for the `rescue-user` on each node.

e) In the **SSH public key** field, paste the public key from the key pair you generate as part of the Prerequisites and Guidelines, on page 53 section.

Step 3 Provide **ND Settings** cluster details.

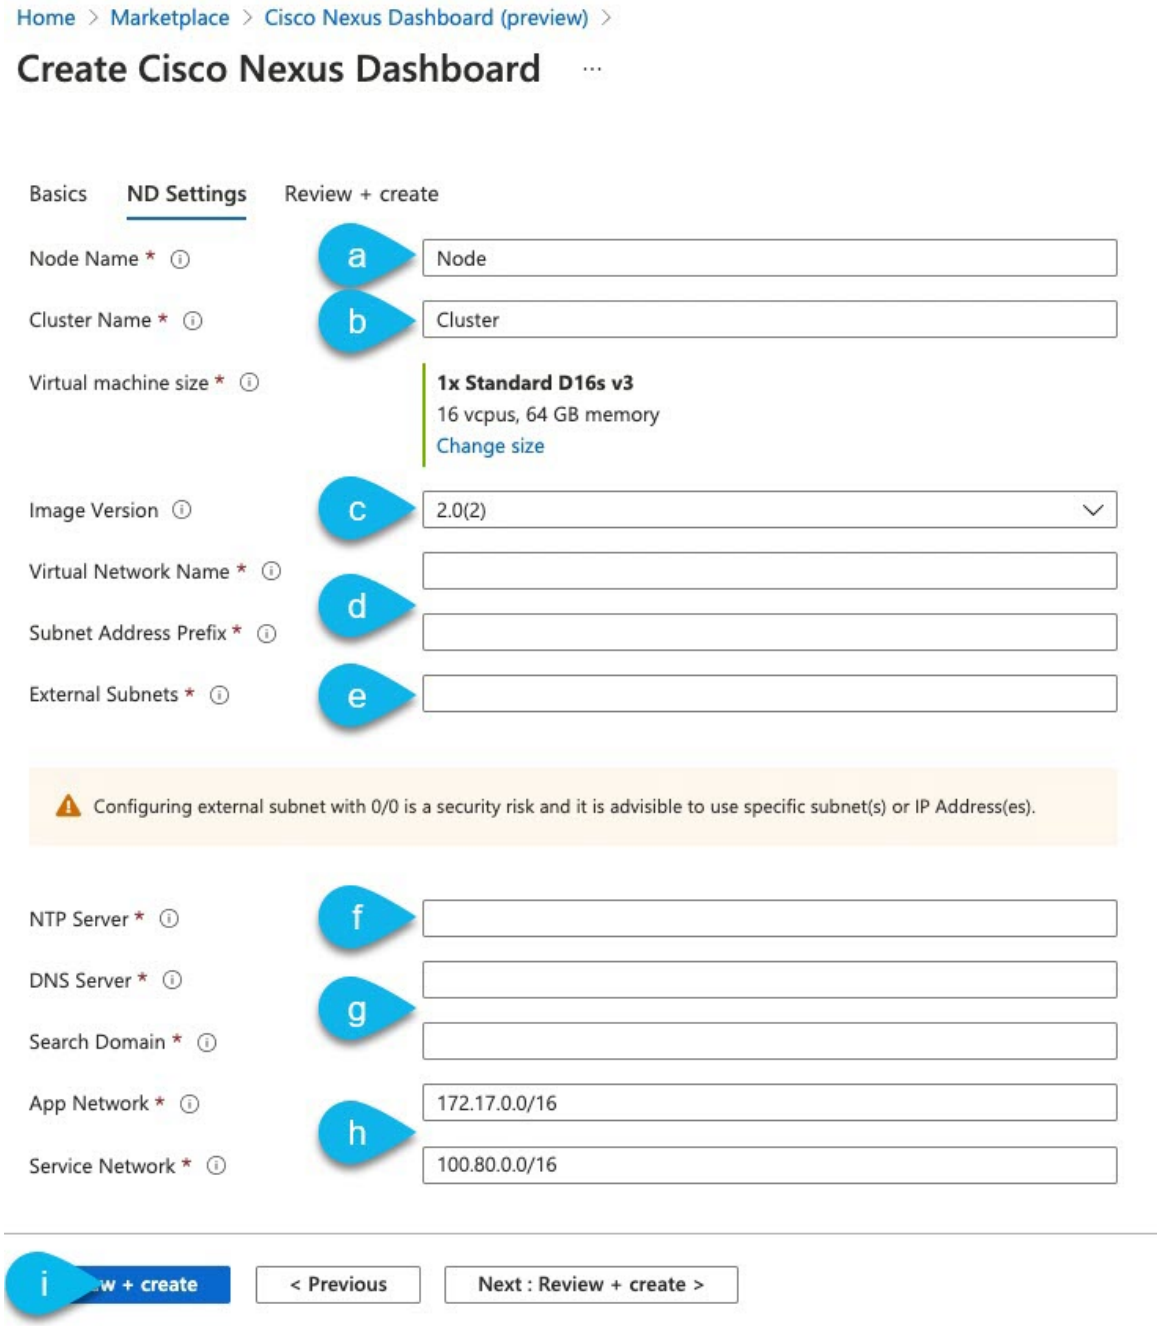

a) Provide the **Node Name**.

The template will deploy a 3-node cluster with each node using the **Node Name** prefix and appending `1`, `2`, and `3` to create unique host names for each node.

b) Provide the **Cluster Name**.

The cluster name must be the same across all nodes you deploy.

c) In the **Image Version** dropdown, ensure the latest release is selected.

d) In the **Virtual Network Name** and **Subnet Address Prefix** fields, provide the name of the VNET and choose a subnet within that VNET.

If the VNET with the name you enter does not exist, it will be created for you.

e) In the **External Subnets** field, provide the external network allowed to access the cluster.

For example, `0.0.0.0/0` to be able to access the cluster from anywhere.

f) Provide the **NTP Servers** information.

g) Provide the **DNS Servers** and **Search Domains** information.

h) Provide the **Application Network** and **Service Network**.

These are internal networks used by the Nexus Dashboard and its processes.

For example, `172.17.0.1/16` and `100.80.0.0/16`.

i) Click **Review + create**.

The product page appears, which shows a summary of your configuration and enables you to launch the cloud formation template.

Step 4 Wait for the VMs deployment to complete, then start the VMs.

Step 5 (Optional) Enable password-based SSH login.

By default only key-based SSH login is enabled for each node. If you'd like to be able to SSH into the nodes using a password, you will need to explicitly enable password-based logins. You can do that by connecting to the VM console of each node from Azure, logging in as `rescue-user` using the password you provided during cluster deployment, and then executing the following command:

```
# acs login-prompt enable
```

Step 6 Verify that the cluster is healthy.

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

a) Login in to one of the nodes.

```
$ ssh rescue-user@<node-ip-address>
```

b) 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
```

c) 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.

When you first log in, you will be prompted to change the password.

CHAPTER 7

Upgrading Nexus Dashboard

Prerequisites and Guidelines

Before you upgrade your existing Nexus Dashboard cluster:

- Ensure that you have read the target release's *Release Notes* for any changes in behavior, guidelines, and issues that may affect your upgrade.
- You must be running Cisco Nexus Dashboard, Release 2.0.1 or later.

  If you are running Cisco Application Services Engine, follow the steps described in Upgrading From Application Services Engine, on page 63 instead.
- The upgrade process is the same for all Nexus Dashboard form factors.

  Regardless of whether you deployed your cluster using physical servers, VMware ESX OVA, or in Azure or AWS cloud, you will use the target release's ISO image to upgrade.
- Ensure that your current Nexus Dashboard cluster is healthy.

  You can check the system status on the **System Overview** page of the Nexus Dashboard GUI or by logging in to one of the nodes as `rescue-user` and executing the `acs health` command
- You must disable any application installed in your cluster before the upgrade and re-enable them after the upgrade completes successfully.

  Before you re-enable the applications, you must configure the App Infra Services deployment profiles as described in the "App Infra Services" section of the *Cisco Nexus Dashboard User Guide*.
- After upgrading to Release 2.0.2, we recommend upgrading all the applications to their latest versions.
- Downgrading from Release 2.0.2 is not supported.

Upgrading Nexus Dashboard

This section describes how to upgrade an existing Nexus Dashboard cluster.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 59

Step 1 Download the Nexus Dashboard image.

a) Browse to the Software Download page.

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

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

c) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.iso`).

Note You must download the `.iso` image for all upgrades, even if you used the VMware ESX `.ova` image or a cloud provider's marketplace for initial cluster deployment.

d) (Optional) Host the image on a web server in your environment.

When you upload the image to your Nexus Dashboard cluster, you will have an option to provide a direct URL to the image.

Step 2 Log in to your current Nexus Dashboard GUI as an `Administrator` user.

Step 3 Upload the new image to the cluster.

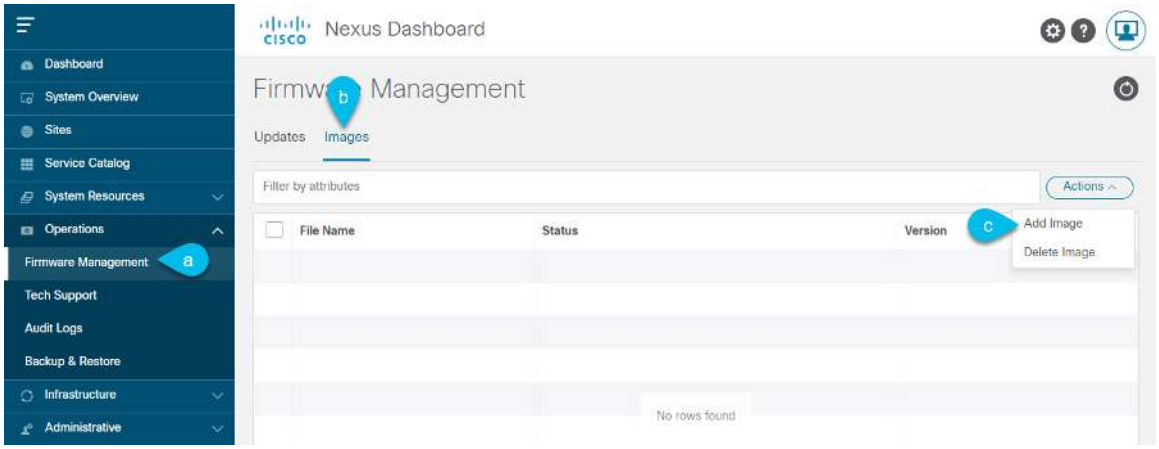

a) Navigate to **Operations** > **Firmware Management**.

b) Select the **Images** tab.

c) From the **Actions** menu, select **Add Image**.

Step 4 Select the new image.

a) In the **Add Firmware Image** window, select **Local**.

Alternatively, if you hosted the image on a web server, choose **Remote** instead.

b) Click **Select file** and select the ISO image you downloaded in the first step.

If you chose to upload a remote image, provide the file path for the image on the remote server.

c) Click **Upload** to add the image.

The image will be uploaded to the Nexus Dashboard cluster, unpacked, processed, and made available for the upgrade. The whole process may take several minutes and you will be able to see the status of the process in the **Images** tab.

Step 5 Wait for the image status to change to `Downloaded`.

You can check the status of the image download progress in the **Images**.

Step 6 Set up the update.

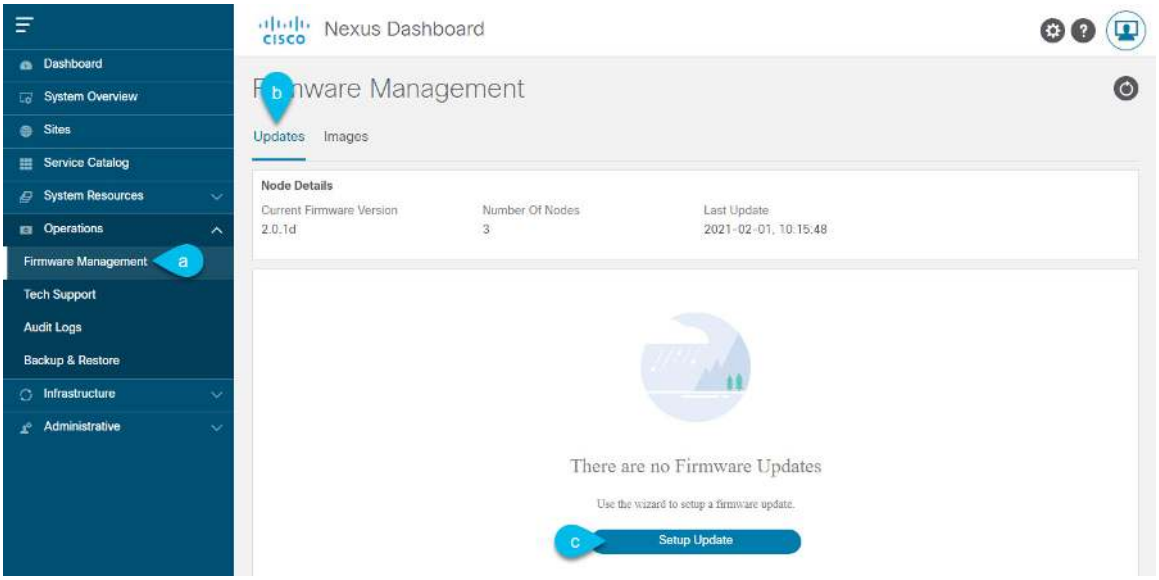

a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Updates** tab.
c) Click **Setup Update**.

The **Firmware Update** screen opens.

Step 7 Choose the upgrade image.

a) In the **Firmware Update** > **Version selection** screen, select the firmware version you uploaded and click **Next**.
b) In the **Firmware Update** > **Confirmation** screen, verify the details and click **Begin Install**.

The installation progress window is displayed. You can navigate away from this screen while the update is in progress. To check on the update status at a later time, navigate to the **Firmware Management** screen and click **View Details** in the **Last Update Status** tile.

This will set up the required Kubernetes images and services but will not switch the cluster to the new version. The cluster will continue to run the existing version until you activate the new image in the next step. The entire process may take up to 20 minutes.

Step 8 Activate the new image.

a) Navigate back to the **Operations** > **Firmware Management** screen
b) In the **Last Update Status** tile, click **View Details**.
c) Click **Activate**.
d) In the **Activation Confirmation** window, click **Continue**.

It may take up to 20 additional minutes for all the cluster services to start and the GUI to become available. The page will automatically reload when the process is completed.

Step 9 If you are hosting multiple applications in the same cluster, configure deployment profiles for the App Infra Services.

If you are hosting only a single application in your Nexus Dashboard cluster, skip this step.

If you are co-hosting multiple applications in the same cluster, you must configure the App Infra Services with deployment profiles appropriate for your combination of applications and fabric sizes.

After the cluster upgrade is completed, follow the instructions described in the "App Infra Services" section of the *Cisco Nexus Dashboard User Guide*, which is also available in the products GUI.

CHAPTER 8

Upgrading From Application Services Engine

Prerequisites and Guidelines

Before you upgrade an existing Cisco Application Services Engine, Release 1.1.3 cluster to Cisco Nexus Dashboard:

- Ensure that you have read the target release's *Release Notes* for any changes in behavior, guidelines, and issues that may affect your upgrade.
- You must be running Cisco Application Services Engine, Release 1.1.3d as a physical appliance.

  If you are already running Cisco Nexus Dashboard, follow the steps described in Upgrading Nexus Dashboard, on page 59 instead.

  Upgrading from earlier releases of Application Services Engine is not supported and you would need to deploy a new cluster as described in earlier chapters in this document.

  If your Application Services Engine is deployed in VMware ESX, Linux KVM, or Amazon Web Services, you cannot upgrade to Nexus Dashboard.
- The upgrade process is the same for all Nexus Dashboard form factors.

  Regardless of whether you deployed your cluster using physical servers, VMware ESX OVA, or in Azure or AWS cloud, you will use the target release's ISO image to upgrade.
- Ensure that your current Application Services Engine is healthy.
- If you have any disabled applications on the existing Application Services Engine cluster, we recommend deleting them before upgrading to Nexus Dashboard.
- If you have Multi-Site Orchestrator application running in your Application Services Engine cluster, you must uninstall it before upgrading the cluster to Nexus Dashboard.

  Migrating to Multi-Site Orchestrator running on Nexus Dashboard consists of multiple steps, which include platform upgrade, application installation, configuration restore, and cloud site upgrades. We strongly recommend following the MSO migration procedure detailed in the "Migrating Existing Cluster to Nexus Dashboard" chapter of the *Multi-Site Deployment Guide*.

- After upgrading to Nexus Dashboard release 2.0.2, we recommend upgrading all the applications to their latest versions.
- Downgrading from Nexus Dashboard release 2.0.2 is not supported.

Upgrading From Application Services Engine

This section describes how to upgrade an existing Application Services Engine, Release 1.1.3d cluster to Nexus Dashboard.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 59

Step 1 Download the Nexus Dashboard image.

a) Browse to the Software Download page.

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

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

c) Download the Cisco Nexus Dashboard image (`nd-dk9.<version>.iso`).

Note You must download the `.iso` image for all upgrades, even if you used the VMware ESX `.ova` image or a cloud provider's marketplace for initial cluster deployment.

d) (Optional) Host the image on a web server in your environment.

When you upload the image to your Nexus Dashboard cluster, you will have an option to provide a direct URL to the image.

Step 2 Log in to your current Application Services Engine GUI as an `Administrator` user.

Step 3 Upload the new image to the cluster.

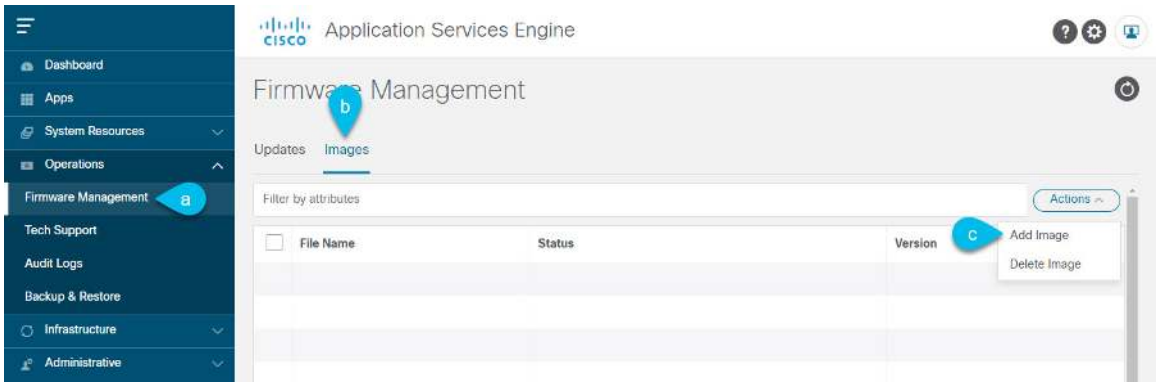

a) Navigate to **Operations** > **Firmware Management**.

b) Select the **Images** tab.

c) From the **Actions** menu, click **Add Image**.

Step 4 Select the new image.

a) In the **Add Firmware Image** window, select **Local**.

Alternatively, if you hosted the image on a web server, choose **Remote** instead.

b) Click **Select file** and select the ISO image you downloaded in the first step.

If you chose to upload a remote image, provide the file path for the image on the remote server.

c) Click **Upload** to add the image.

The image will be uploaded to the Application Services Engine cluster, unpacked, processed, and made available for the upgrade. The whole process may take several minutes and you will be able to see the status of the process in the **Images** tab.

Step 5 Wait for the image status to change to `Downloaded`.

You can check the status of the image download progress in the **Images**.

Step 6 Set up the update.

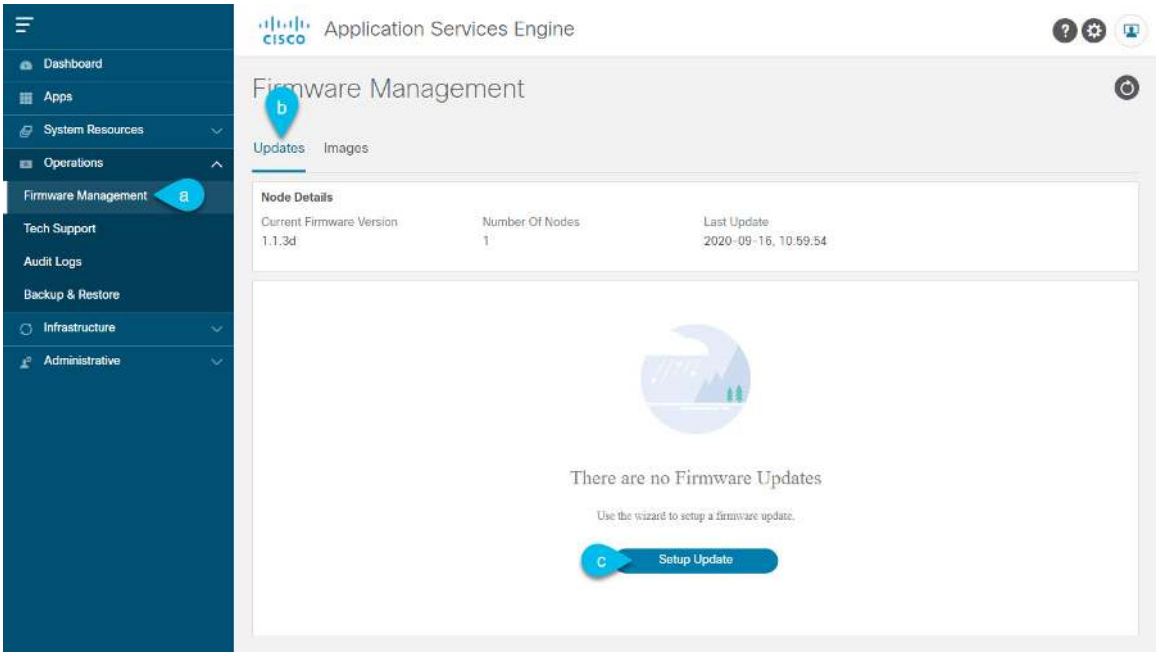

a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Updates** tab.
c) Click **Setup Update**.

Step 7 Provide update details.

a) In the **Version Selection** screen, select the firmware version you uploaded, then click **Next**.
b) In the **Confirmation** screen, verify the details, then click **Begin Install**.

The installation progress window is displayed. You can navigate away from this screen while the update is in progress. To check on the update status at a later time, navigate to the **Firmware Management** screen and click **View Details** in the **Last Update Status** tile.

Step 8 Activate the new image.

a) Navigate back to the **Operations** > **Firmware Management** screen

b) In the **Last Update Status** tile, click **View Details**.
c) Click **Activate**.
d) In the **Activation Confirmation** window, click **Continue**.

It may take up to 20 additional minutes for all the cluster services to start and the GUI to become available. The page will automatically reload when the process is completed.