

Cisco Nexus Dashboard Deployment Guide, Release 3.0.x

First Published: 2023-08-22

Last Modified: 2024-06-24

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: https://www.cisco.com/c/en/us/about/legal/cloud-and-software/software-terms.html. Cisco product warranty information is available at https://www.cisco.com/c/en/us/products/warranty-listing.html. US Federal Communications Commission Notices are found here https://www.cisco.com/c/en/us/products/us-fcc-notice.html.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2023–2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

Trademarks ?

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

CHAPTER 2 **Deployment Overview and Requirements** 3

Deployment Overview 3

Prerequisites and Guidelines 6

Communication Ports 16

Fabric Connectivity 28

Node Distribution Across Sites 36

Services Co-location Use Cases 38

Pre-Installation Checklist 40

CHAPTER 3 **Deploying as Physical Appliance** 45

Prerequisites and Guidelines 45

Deploying Nexus Dashboard as Physical Appliance 48

CHAPTER 4 **Deploying in VMware ESX** 61

Prerequisites and Guidelines 61

Deploying Nexus Dashboard Using VMware vCenter 64

Deploying Nexus Dashboard Directly in VMware ESXi 80

CHAPTER 5 **Deploying in Amazon Web Services** 93

Prerequisites and Guidelines 93

Deploying Nexus Dashboard in AWS 95

CHAPTER 6 **Deploying in Microsoft Azure** 105

Prerequisites and Guidelines 105

Generating an SSH Key Pair in Linux or MacOS 106

Generating an SSH Key Pair in Windows 107

Deploying Nexus Dashboard in Azure 110

CHAPTER 7 **Deploying in Existing Red Hat Enterprise Linux Installation** 119

Prerequisites and Guidelines 119

Deploying Nexus Dashboard in Existing Red Hat Enterprise Linux Installation 121

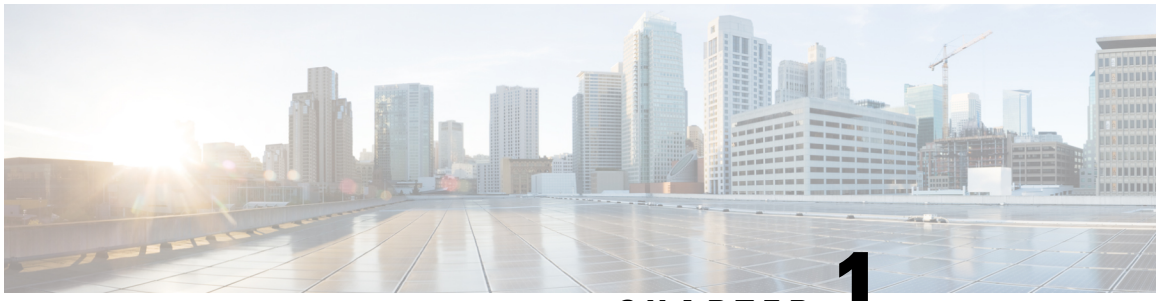
Uninstalling Nexus Dashboard Software 130

Troubleshooting Nexus Dashboard Deployments in RHEL 130

CHAPTER 8 **Upgrading Nexus Dashboard** 133

Prerequisites and Guidelines 133

Upgrading Nexus Dashboard 135



CHAPTER 1

New and Changed Information

• New and Changed Information, on page 1

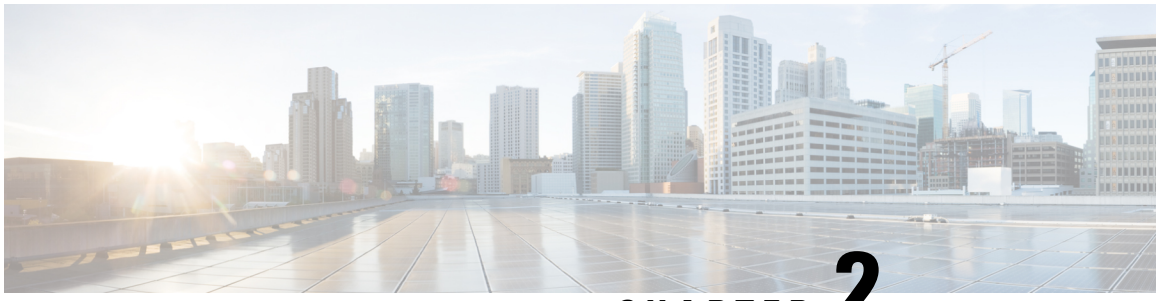
New and Changed Information

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
3.0(1)	First release of this document.	--

New and Changed Information



CHAPTER 2

Deployment Overview and Requirements

- Deployment Overview, on page 3
- Prerequisites and Guidelines, on page 6
- Communication Ports, on page 16
- Fabric Connectivity, on page 28
- Node Distribution Across Sites, on page 36
- Services Co-location Use Cases, on page 38
- Pre-Installation Checklist, on page 40

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 services, such as Nexus Dashboard Insights and Nexus Dashboard Orchestrator. These services are available for all the data center sites and provide real time analytics, visibility, assurance for network policies and operations, as well as policy orchestration for the data center fabrics, such as Cisco ACI or Cisco NDFC.

Nexus Dashboard provides a common platform and modern technology stack for the above-mentioned 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.

Nexus Dashboard cluster typically consists of 1 or 3 `primary` nodes. For 3-node clusters, you can also provision a number of `worker` nodes to enable horizontal scaling and `standby` nodes for easy cluster recovery in case of a primary node failure. For maximum number of `worker` and `standby` nodes supported in this release, see the "Verified Scalability Limits" sections of the *Cisco Nexus Dashboard Release Notes*.

Note This document describes initial configuration of the base 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 hardware" specifically to refer to the hardware and "Nexus Dashboard" to refer to the software stack and the GUI console.



Note Root access to the Nexus Dashboard software is restricted to Cisco TAC only. A special user `rescue-user` is created for all Nexus Dashboard deployments to enable a set of operations and troubleshooting commands. For additional information about the available `rescue-user` commands, see the "Troubleshooting" chapter of the *Nexus Dashboard User Guide*.

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 configuration and operations procedures are described in the *Cisco Nexus Dashboard User Guide*.

Services

Nexus Dashboard is a standard appliance platform to build and deploy services that would allow you to consume all Nexus Dashboard products in a consistent and uniform manner. You can deploy services like Insights, Orchestrator, Fabric Controller, and Data Broker with the Nexus Dashboard platform providing the necessary capacity and life cycle management operations for these services.

Typically, the Nexus Dashboard platform is shipped with only the software required for managing the lifecycle of these services, but no actual services are packaged with the appliance. If you allow public network connectivity from your data centers, you can download and install the services with a few clicks. However, without public network connectivity, you will need to manually download these services' images, upload them to the platform, and perform installation operations before you can use them.

If you are ordering the physical Nexus Dashboard servers, you have the option to choose some services to be pre-installed on the hardware before it is shipped to you. For more information, see the *Nexus Dashboard Ordering Guide*. Note that if you are deploying the virtual or cloud form factors of the Nexus Dashboard, you will need to deploy the services separately after the cluster is ready.

Available Form Factors

This release of Cisco Nexus Dashboard 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. The physical form factor currently supports two different Cisco UCS servers (**SE-NODE-G2 and ND-NODE-L4**) for the cluster nodes, which can be mixed within the same cluster.



Note Not all services are supported on all form factors. When planning your deployment, ensure to check Cisco Nexus Dashboard Cluster Sizing for form factor and cluster size requirements.

- 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.

- Linux KVM (`.qcow2`)

  Virtual form factor that allows you to deploy a Nexus Dashboard cluster using three Linux KVM 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.

- In an existing Red Hat Enterprise Linux (RHEL) system

  Beginning with Release 2.2(1), you can run Nexus Dashboard node in an existing Red Hat Enterprise Linux server.

Cluster Sizing and Availability Guidelines

As mentioned previously, Nexus Dashboard cluster is first deployed using 1 or 3 `primary` nodes. Depending on the type and number of services you choose to run, you may be required to deploy additional worker nodes in your cluster after the initial deployment. For cluster sizing information and recommended number of nodes based on specific use cases, see the Cisco Nexus Dashboard Cluster Sizing tool.

Note

- Single-node clusters are supported for a limited number of services and cannot be extended to a 3-node cluster after the initial deployment.
- Only 3-node clusters support additional `worker` or `standby` nodes.
- If you deploy a single-node cluster and want to extended it to a 3-node cluster or add `worker` nodes, you will need to redeploy it as a base 3-node cluster.
- For 3-node clusters, at least two `primary` nodes are required for the cluster to remain operational. If two `primary` nodes fail, the cluster will go offline and cannot be used until you recover it as described in the *Cisco Nexus Dashboard User Guide*.
- You cannot dynamically add or remove primary or secondary nodes to change the overall size of an existing cluster. To resize a cluster, you must back up the cluster, rebuild it to the desired size, and then restore the configuration. Ensure that the new cluster size meets the supported scale and feature requirements for your deployment.

  An exception applies when a cluster size change is explicitly specified in Cisco documentation. For example, during an upgrade, you may need to change the cluster size for supportability reasons.

After your initial 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 Services

For the full list of supported applications and the associated compatibility and interoperability information, see the Nexus Dashboard and Services Compatibility Matrix.

Prerequisites and Guidelines

Network Time Protocol (NTP) and Domain Name System (DNS)

The Nexus Dashboard nodes require valid DNS and NTP servers for all deployments and upgrades.

Lack of valid DNS connectivity (such as if using an unreachable or a placeholder IP address) can prevent the system from deploying or upgrading successfully.



Note Nexus Dashboard acts as both a DNS client and resolver. It uses an internal Core DNS server which acts as DNS resolver for internal services. It also acts as a DNS client to reach external hosts within the intranet or the Internet, hence it requires an external DNS server to be configured.

Additionally, Nexus Dashboard does not support DNS servers with wildcard records.

Beginning with release 3.0(1), Nexus Dashboard also supports NTP authentication using symmetrical keys. If you want to enable NTP authentication, you need to provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.

The following guidelines apply when enabling NTP authentication:

- For symmetrical authentication, any key you want to use must be configured the same on both your NTP server and Nexus Dashboard.

  The ID, authentication type, and the key/passphrase itself must match and be trusted on both your NTP server and Nexus Dashboard.
- Multiple servers can use the same key.

  In this case the key must only be configured once on Nexus Dashboard, then assigned to multiple servers.
- Both Nexus Dashboard and the NTP servers can have multiple keys as long as key IDs are unique.
- This release supports SHA1, MD5, and AES128CMAC authentication/encoding types for NTP keys.



Note We recommend using AES128CMAC due to its higher security .

- When adding NTP keys in Nexus Dashboard, you must tag them as `trusted`; untrusted keys will fail authentication.

  This option allows you to easily disable a specific key in Nexus Dashboard if the key becomes compromised.

- You can choose to tag some NTP servers as `preferred` in Nexus Dashboard.

  NTP clients can estimate the "quality" of an NTP server over time by taking into account RTT, time response variance, and other variables. Preferred servers will have higher priority when choosing a primary server.

- If you are using an NTP server running `ntpd`, we recommend version 4.2.8p12 at a minimum.

- The following restrictions apply to all NTP keys:
  - The maximum key length for SHA1 and MD5 keys is 40 characters, while the maximum length for AES128 keys is 32 characters.
  - Keys that are shorter than 20 characters can contain any ASCII character excluding '`#`' and spaces. Keys that are over 20 characters in length must be in hexadecimal format.
  - Keys IDs must be in the 1-65535 range.
  - If you configure keys for any one NTP server, you must also configure the keys for all other servers.

Enabling and configuring NTP authentication is described as part of the deployment steps in the later sections.

BGP Configuration and Persistent IPs

Older releases of Nexus Dashboard allowed you to configure one or more persistent IP addresses for services (such as Nexus Dashboard Insights) that require retaining the same IP addresses even in case they are relocated to a different Nexus Dashboard node. However, in those releases, the persistent IPs had to be part of the management and data subnets and the feature could be enabled only if all nodes in the cluster were part of the same Layer 3 network. Here the services used Layer 2 mechanisms like Gratuitous ARP or Neighbor Discovery to advertise the persistent IPs within it's Layer 3 network

Beginning with Release 2.2(1), the Persistent IPs feature is supported even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IPs are advertised out of each node's data links via BGP, which we refer to as "Layer 3 mode". The IPs must also be part of a subnet that is not overlapping with any of the nodes' management or data subnets. If the persistent IPs are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IPs are part of those networks, the feature will operate in Layer 2 mode. BGP can be enabled during cluster deployment or from the Nexus Dashboard GUI after the cluster is up and running.

If you plan to enable BGP and use the persistent IP functionality, you must:

- Ensure that the peer routers exchange the advertised persistent IPs between the nodes' Layer 3 networks.
- Choose to enable BGP at the time of the cluster deployment as described in the subsequent sections or enable it afterwards in the Nexus Dashboard GUI as described in the "Persistent IP Addresses" sections of the *User's Guide*.
- Ensure that the persistent IP addresses you allocate do not be overlap with any of the nodes' management or data subnets.

- Ensure that you fulfill the service-specific persistent IP requirements listed in Table 3: Service-Specific Network Requirements, on page 9 listed in the next section.

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

Table 2: External Network Purpose

Data Network	Management Network
• Nexus Dashboard node clustering • Service to service communication • Nexus Dashboard nodes to Cisco APIC, Cloud Network Controller, and NDFC communication For example, the network traffic for services such as Nexus Dashboard Insights. • Telemetry traffic for switches and on-boarded fabrics	• 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 services, https://dcappcenter.cisco.com must be reachable via the Management Network • Intersight device connector • AAA traffic

The two networks have the following requirements:

- For all new Nexus Dashboard deployments, the management network and data network must be in different subnets.

Note With the exception of the Nexus Dashboard Fabric Controller (SAN Controller), which can be deployed in Nexus Dashboard using the same subnets for the data and management networks.

- When setting up remote authentication, AAA server must not be in the same subnet as the data interface.
- For physical clusters, 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 Dashboard Insights service, the data network must provide IP reachability to the in-band network of each fabric and of the APIC.

- For Nexus Dashboard Insights and AppDynamics integration, the data network must provide IP reachability to the AppDynamics controller.
- For Nexus Dashboard Orchestrator service, 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 NDFC 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.

Note If external VLAN tag is configured for switch ports that are used for data network traffic, you must enable jumbo frames or configure custom MTU equal to or greater than 1504 bytes.

- The table bellow summarize service-specific requirements for the management and data networks.

Note Changing the data subnet requires re-deploying the cluster, so we recommend using a larger subnet than the bare minimum required by the nodes and services to account for any additional services in the future. In addition to the requirements listed in this section, ensure that you consult the *Release Notes* for the specific service you plan to deploy.

Allocating persistent IP addresses is done after the cluster is deployed using the External Service Pools configuration in the UI, as described in the *Cisco Nexus Dashboard User Guide*.

We recommend consulting the specific service's documentation for any additional requirements and caveats related to persistent IP configuration.

Table 3: Service-Specific Network Requirements

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Orchestrator	Layer 3 adjacent	Layer 3 adjacent	N/A
Nexus Dashboard Insights without SFLOW/NetFlow (ACI fabrics)	Layer 3 adjacent	Layer 3 adjacent	N/A
Nexus Dashboard Insights without SFLOW/NetFlow (NDFC fabrics)	Layer 3 adjacent	Layer 2 adjacent	6 IPs in data interface network if using IPv4 7 IPs in data interface network if using IPv6

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Insights with SFLOW/NetFlow (ACI or NDFC fabrics)	Layer 3 adjacent	Layer 2 adjacent	6 IPs in data interface network

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
Nexus Dashboard Fabric Controller, Release 12.1.3	Layer 2 or Layer 3 adjacent	Layer 2 or Layer 3 adjacent	

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
			When operating in Layer 2 mode with LAN deployment type and **LAN Device Management Connectivity** set to `Management` (default) • 2 IPs in the management network for SNMP/Syslog and SCP services • If **EPL** is enabled, 1 additional IP in the data network for each fabric • If **IP Fabric for Media** is enabled, 1 additional IP in the management network for telemetry When operating in Layer 2 mode with LAN deployment type and **LAN Device Management Connectivity** set to `Data`: • 2 IPs in the data network for SNMP/Syslog and SCP services • If **EPL** is enabled, 1 additional IP in the data network for each fabric • If **IP Fabric for Media** is enabled, 1 additional IP in the data network for telemetry When operating in Layer

Nexus Dashboard Service	Management Interface	Data Interface	Total Number of Persistent IPs
			3 mode with LAN deployment type: • **LAN Device Management Connectivity** must be set to `Data` • 2 IPs for SNMP/Syslog and SCP services • If **EPL** is enabled, 1 additional IP in the data network for each fabric • All persistent IPs must be part of a separate pool that must not overlap with the management or data subnets For more information about Layer 3 mode for persistent IPs, see the Persistent IPs section in the User's Guide When operating in Layer 2 or Layer 3 mode with SAN Controller deployment type: • 1 IP for SSH • 1 IP for SNMP/Syslog • 1 IP per Nexus Dashboard cluster node for SAN Insights functionality IP Fabric for Media is not supported in Layer 3 mode

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

Note You must always use the lowest RTT requirement when deploying the Nexus Dashboard cluster and services. For example, if you plan to co-host the Insights and Orchestrator services, site connectivity RTT must not exceed 50ms.

Table 4: RTT Requirements

Service	Connectivity	Maximum RTT
Nexus Dashboard Multi-Cluster connectivity	Between nodes across clusters that are connected via multi-cluster connectivity For more information about multi-cluster connectivity, see *Cisco Nexus Dashboard Infrastructure Management*.	500 ms
Nexus Dashboard Orchestrator	Between nodes	50 ms
	To sites	For APIC sites: 500 ms For NDFC sites: 150 ms
Nexus Dashboard Insights	Between nodes	50 ms
	To switches	150 ms
Nexus Dashboard Fabric Controller	Between nodes	50 ms
	To switches	200 ms*

* POAP (PowerOn Auto Provisioning) is supported with a max RTT of 50 ms between Nexus Dashboard Fabric Controller and the switches.

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 pre-populated during deployment.

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

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

If you are planning to deploy multiple Nexus Dashboard clusters, they can use the same Application and Service subnets.



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.

For the same reason, we recommend not using `169.254.0.0/16` (the Kubernetes `br1` subnet) for the App or Service subnets.

IPv4 and IPv6 Support

Prior releases of Nexus Dashboard supported either pure IPv4 or dual stack IPv4/IPv6 (for management network only) configurations for the cluster nodes. Beginning with release 3.0(1), Nexus Dashboard supports pure IPv4, pure IPv6, or dual stack IPv4/IPv6 configurations for the cluster nodes and services.

When defining IP configuration, the following guidelines apply:

- All nodes and networks in the cluster must have a uniform IP configuration – either pure IPv4, or pure IPv6, or dual stack IPv4/IPv6.
- If you deploy the cluster in pure IPv4 mode and want to switch to dual stack IPv4/IPv6 or pure IPv6, you must redeploy the cluster.
- For dual stack configurations:
    - Both external (data and management) and internal (app and services) networks must be in dual stack mode.

      Partial configurations, such as IPv4 data network and dual stack management network, are not supported.
    - IPv6 addresses are also required for physical servers' CIMCs.
    - You can configure either IPv4 or IPv6 addresses for the nodes' management network during initial node bring up, but you must provide both types of IPs during the cluster bootstrap workflow.

      Management IPs are used to log in to the nodes for the first time to initiate cluster bootstrap process.
    - All internal certificates will be generated to include both IPv4 and IPv6 Subject Alternative Names (SANs).
    - Kubernetes internal core services will start in IPv4 mode.
    - DNS will serve and forward to both IPv4 and IPv6 and server both types of records.
    - VxLAN overlay for peer connectivity will use data network's IPv4 addresses.

      Both IPv4 and IPv6 packets are encapsulated within the VxLAN's IPv4 packets.
    - The UI will be accessible on both IPv4 and IPv6 management network addresses.

- For pure IPv6 configurations:
  - Pure IPv6 mode is supported for physical and virtual form factors only.

    Clusters deployed in AWS, Azure, or an existing Red Hat Enterprise Linux (RHEL) system do not support pure IPv6 mode.
  - You must provide IPv6 management network addresses when initially configuring the nodes.

    After the nodes (physical, virtual, or cloud) are up, these IPs are used to log in to the UI and continue cluster bootstrap process.
  - You must provide IPv6 CIDRs for the internal App and Service networks described above.
  - You must provide IPv6 addresses and gateways for the data and management networks described above.
  - All internal certificates will be generated to include IPv6 Subject Alternative Names (SANs).
  - All internal services will start in IPv6 mode.
  - VxLAN overlay for peer connectivity will use data network's IPv6 addresses.

    IPv6 packets are encapsulated within the VxLAN's IPv6 packets.
  - All internal services will use IPv6 addresses.

Communication Ports

The following sections provide a reference for ports required by the Nexus Dashboard cluster and services.

Note All services use TLS or mTLS with encryption to protect data privacy and integrity while in transit.

Nexus Dashboard Ports

The following ports are required by the Nexus Dashboard cluster.

Table 5: Nexus Dashboard Ports (Management Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, CIMC, default gateway
SSH	22	TCP	In/Out	CLI and CIMC of the cluster nodes

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
TACACS	49	TCP	Out	TACACS server
DNS	53	TCP/UDP	Out	DNS server
HTTP	80	TCP	Out	Internet/proxy
NTP	123	UDP	Out	NTP server
HTTPS	443	TCP	In/Out	UI, other clusters (for multi-cluster connectivity), fabrics, Internet/proxy
LDAP	389 636	TCP	Out	LDAP server
Radius	1812	TCP	Out	Radius server
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
Infra-Service	30012 30021 30500-30600	TCP/UDP	In/Out	Other cluster nodes

Table 6: Nexus Dashboard Ports (Data Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, CIMC, default gateway
SSH	22	TCP	Out	In-band of switches and APIC
DNS	53	TCP/UDP	In/Out	Other cluster nodes and DNS server
NFSv3	111	TCP/UDP	In/Out	Remote NFS server

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
HTTPS	443	TCP	Out	In-band of switches and APIC/NDFC
NFSv3	608	UDP	In/Out	Remote NFS server
SSH	1022	TCP/UDP	In/Out	Other cluster nodes
NFSv3	2049	TCP	In/Out	Remote NFS server
VXLAN	4789	UDP	In/Out	Other cluster nodes
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
Infra-Service	3379 3380 8989 9090 9969 9979 9989 15223 30002-30006 30009-30010 30012 30014-30015 30018-30019 30025 30027	TCP	In/Out	Other cluster nodes
Infra-Service	30016 30017	TCP/UDP	In/Out	Other cluster nodes
Infra-Service	30019	UDP	In/Out	Other cluster nodes
Infra-Service	30500-30600	TCP/UDP	In/Out	Other cluster nodes

Nexus Dashboard Insights Ports

In addition to the ports required by the Nexus Dashboard cluster nodes, which are listed above, the following ports are required by the Nexus Dashboard Insights service.

Table 7: Nexus Dashboard Insights Ports (Data Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
Show Techcollection	2022	TCP	In/Out	In-band of switches and APIC/NDFC
Flow Telemetry	5640-5671	UDP	In	In-band of switches
TAC Assist	8884	TCP	In/Out	Other cluster nodes
KMS	9989	TCP	In/Out	Other cluster nodes and ACI fabrics
Kafka	30001	TCP	In/Out	In-band IP of switches and APIC/NDFC
SW Telemetry	5695 30000 57500 30570	TCP	In/Out	Other cluster nodes

Nexus Dashboard Fabric Controller Ports

In addition to the ports required by the Nexus Dashboard (ND) cluster nodes, the following ports are required by the Nexus Dashboard Fabric Controller (NDFC) service.

Note The following ports apply to the Nexus Dashboard management network and/or data network interfaces depending on which interface provides IP reachability from the NDFC service to the switches.

Table 8: Nexus Dashboard Fabric Controller Ports

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
SCP	22	TCP	Out	SCP clients archiving NDFC backup files to remote server.
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's **Server Settings** menu. This is an optional feature.
DHCP	67	UDP	In	If NDFC local DHCP server is configured for Bootstrap/POAP purposes. This applies to LAN deployments only. **Note** When using NDFC as a local DHCP server for POAP purposes, all ND master node IPs must be configured as DHCP relays. Whether the ND nodes' management or data IPs are bound to the DHCP server is determined by the LAN Device Management Connectivity in the NDFC Server Settings.
DHCP	68	UDP	Out	
SNMP	161	TCP/UDP	Out	SNMP traffic from NDFC to devices.
HTTPS/HTTP (NX-API)	443/80	TCP	Out	NX-API HTTPS/HTTP client connects to device NX-API server on port 443/80, which is also configurable. NX-API is an optional feature, used by limited set of NDFC functions. This applies to LAN deployments only.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	NDFC provides an integrated host and physical network topology view by correlating the information obtained from registered VMM domains, such as VMware vCenter or OpenStack, as well as container orchestrators, such as Kubernetes. This is an optional feature

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.

Note The following ports apply to the External Service IPs, also known as persistent IPs, used by some of the NDFC services. These External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 9: Nexus Dashboard Fabric Controller Persistent IP Ports

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SCP	22	TCP	In	SCP is used by various features to transfer files between devices and the NDFC service. The NDFC SCP service serves as the SCP server for both downloads and uploads. SCP is also used by the POAP client on the devices to download POAP-related files. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
TFTP (POAP)	69	TCP	In	Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings. This applies to LAN deployments only.
HTTP (POAP)	80	TCP	In	Only used for device zero-touch provisioning via POAP, where devices can send (limited jailed write-only access to NDFC) basic inventory information to NDFC to start secure POAP communication. NDFC Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings. This applies to LAN deployments only.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
BGP	179	TCP	In/Out	For Endpoint Locator, per fabric where it is enabled, an EPL service is spawned with its own persistent IP. This service is always associated with the Nexus Dashboard data interface. NDFC EPL service peers with the appropriate BGP entity (typically BGP Route-Reflectors) on the fabric to get BGP updates needed to track endpoint information. This feature is only applicable for VXLAN BGP EVPN fabric deployments. This applies to LAN deployments only.
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished via the NDFC HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP assigned to that pod. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings. This applies to LAN deployments only.
Syslog	514	UDP	In	When NDFC is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SCP	2022	TCP	Out	Transport tech-support file from persistent IP of NDFC POAP-SCP pod to a separate ND cluster running Nexus Dashboard Insights. The SCP-POAP service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings
SNMP Trap	2162	UDP	In	SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings
HTTP (PnP)	9666	TCP	In	Cisco Plug and Play (PnP) for Catalyst devices is accomplished via NDFC HTTP port 9666 and HTTPS port 9667. HTTP on port 9666 is used to send CA certificate bundle to devices to prime the device for HTTPS mode and actual PnP happens over HTTPS on port 9667 afterwards. PnP service, like POAP, runs on persistent IP that is associated with either the management or data subnet. Persistent IP subnet is controlled by the **LAN Device Management Connectivity** setting in the NDFC Server Settings. This applies to LAN deployments only.
HTTPS (PnP)	9667	TCP	In	

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
GRPC (Telemetry)	33000	TCP	In	SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP. This is enabled on SAN deployments only.
GRPC (Telemetry)	50051	TCP	In	Information related to multicast flows for IP Fabric for Media deployments as well as PTP for general LAN deployments is streamed out via software telemetry to a persistent IP associated with a NDFC GRPC receiver service pod. This is enabled on LAN and Media deployments only.

Nexus Dashboard Fabric Controller Ports for SAN Deployments

Nexus Dashboard Fabric Controller can be deployed on a single-node or 3-node Nexus Dashboard cluster. The following ports are required for NDFC SAN deployments on single-node clusters.

Table 10: Nexus Dashboard Fabric Controller Ports for SAN Deployments on Single-Node Clusters

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
SCP	22	TCP	Out	SCP clients archiving NDFC backup files to remote server.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's Server Settings menu. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from NDFC to devices.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	NDFC provides an integrated host and physical network topology view by correlating the information obtained from registered VMM domains, such as VMware vCenter or OpenStack, as well as container orchestrators, such as Kubernetes. This is an optional feature.

Note The following ports apply to the External Service IPs, also known as Persistent IPs, used by some of the NDFC services. These External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 11: Nexus Dashboard Fabric Controller Persistent IP Ports for SAN Deployments on Single-Node Clusters

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SCP	22	TCP	In	SCP is used by various features to transfer files between devices and the NDFC service. The NDFC SCP service functions for both downloads and uploads.
Syslog	514	UDP	In	When NDFC is configured as a Syslog server, syslogs from the devices are sent out towards the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet. This is controlled by the LAN Device Management Connectivity setting in the NDFC Server Settings.

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SNMP Trap	2162	UDP	In	SNMP traps from devices to NDFC are sent out toward the persistent IP associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in NDFC has a persistent IP that is associated with either the management or data subnet.
GRPC (Telemetry)	33000	TCP	In	SAN Insights Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to NDFC Persistent IP. This is enabled on SAN deployments only.

Fabric Connectivity

The following sections describe how to connect your Nexus Dashboard cluster nodes to the management and data networks and how to connect the cluster to your fabrics.

For on-premises APIC or NDFC 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.

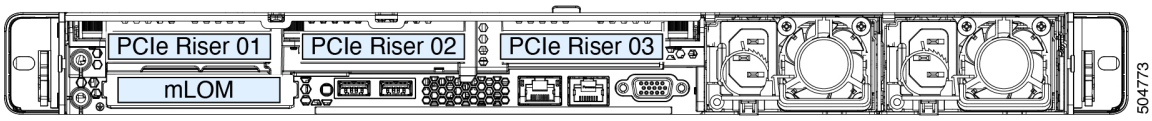
For Cisco Cloud Network Controller fabrics, you must connect via a Layer 3 network.

Physical Node Cabling

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

Physical nodes can be deployed in UCS-C220-M5 (`SE-NODE-G2`) and UCS-C225-M6 (`ND-NODE-L4`) physical servers with the following guidelines:

Figure 1: mLOM and PCIe Riser 01 Card Used for Node Connectivity



- Both servers come with a Modular LAN on Motherboard (mLOM) card, which you use to connect to the Nexus Dashboard management network.
- The UCS-C220-M5 server includes a 4-port VIC1455 card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity
- The UCS-C225-M6 server includes either a 2x10GbE NIC (`APIC-P-ID10GC`), or 2x25/10GbE SFP28 NIC (`APIC-P-I8D25GF`), or the VIC1455 card in the "`PCIe-Riser-01`" slot (shown in the above diagram), which you use for Nexus Dashboard data network connectivity.

When connecting the node to your management and data networks:

- The interfaces are configured as Linux bonds (one for the data interfaces and one for the management interfaces) running in active-standby mode.
- For management network:
  - You must use the `mgmt0` and `mgmt1` on the mLOM card.
  - All ports must have the same speed, either 1G or 10G.
- For data network:
  - On the UCS-C220-M5 server, you must use the VIC1455 card.
  - On the UCS-C225-M6 server, you can use the 2x10GbE NIC (`APIC-P-ID10GC`), or 2x25/10GbE SFP28 NIC (`APIC-P-I8D25GF`), or the VIC1455 card.

Note If you connect using the 25G Intel NIC, you must disable the FEC setting on the switch port to match the setting on the NIC:

```
(config-if)# fec off
# show interface ethernet 1/34
Ethernet1/34 is up
admin state is up, Dedicated Interface
  [...]
  FEC mode is off
```

- All interfaces must be connected to individual host-facing switch ports; Fabric Extenders (FEX), PortChannel (PC), and Virtual PortChannel (vPC) are not supported.
- Configure Cisco VIC-connected data uplinks as trunk ports on the connected switches. Add a native VLAN on the upstream trunk port if VLAN tagging is not used in Nexus Dashboard for the data network.



Note When connecting Nexus Dashboard nodes to Arista switches on the data interfaces (fabric interfaces), configure the switch interfaces as access ports rather than trunk ports, with the access VLAN set to the data subnet VLAN.

- Verify that the port-channel feature is not enabled on the VIC:
  1. Log into the CIMC.
  2. Click the hamburger icon, then navigate to:

     Networking > **Adapter Card 1** > **General**
  3. In the **Adapter Card Properties** area, locate the **Port Channel** field and verify that the option is not enabled (check box is unchecked).

     If the **Port Channel** option is enabled, remove the check from the check box before beginning the Nexus Dashboard deployment process.
- All ports must have the same speed, either 10G or 25G.
- Port-1 corresponds to `fabric0` in Nexus Dashboard and Port-2 corresponding to `fabric1`.

  You can use both `fabric0` and `fabric1` for data network connectivity.



Note When using a 4-port card, the order of ports depends on the model of the server you are using:

- On the UCS-C220-M5 server, the order from left to right is Port-1, Port-2, Port-3, Port-4.
- On the UCS-C225-M6 server, the order from left to right is Port-4, Port-3, Port-2, Port-1.
- Port-1 and Port-2 OR Port-3 and Port-4 are used in pairs. We recommend that you use the Port-1/Port-2 pair.

- If you connect the nodes to Cisco Catalyst switches, you must add `switchport voice vlan dot1p` command to the switch interfaces.

  On the 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.

Note This command is not supported on Cisco Catalyst 9500X or 9600X models. For those models, you must configure the interface as a trunk interface and specify the VLAN in the Nexus Dashboard data network configuration.

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 Nexus Dashboard 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 Nexus Dashboard Orchestrator to manage Cisco NDFC fabrics, you must establish connectivity from the data interface to the in-band interface of each site's NDFC.
- If you are deploying Day-2 Operations applications, such as Nexus Dashboard 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 NDFC fabrics, if the data interface and NDFC's inband interface are in different subnets, you must add a route on NDFC to reach the Nexus Dashboard's data network address.

  You can add the route from the NDFC 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.

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

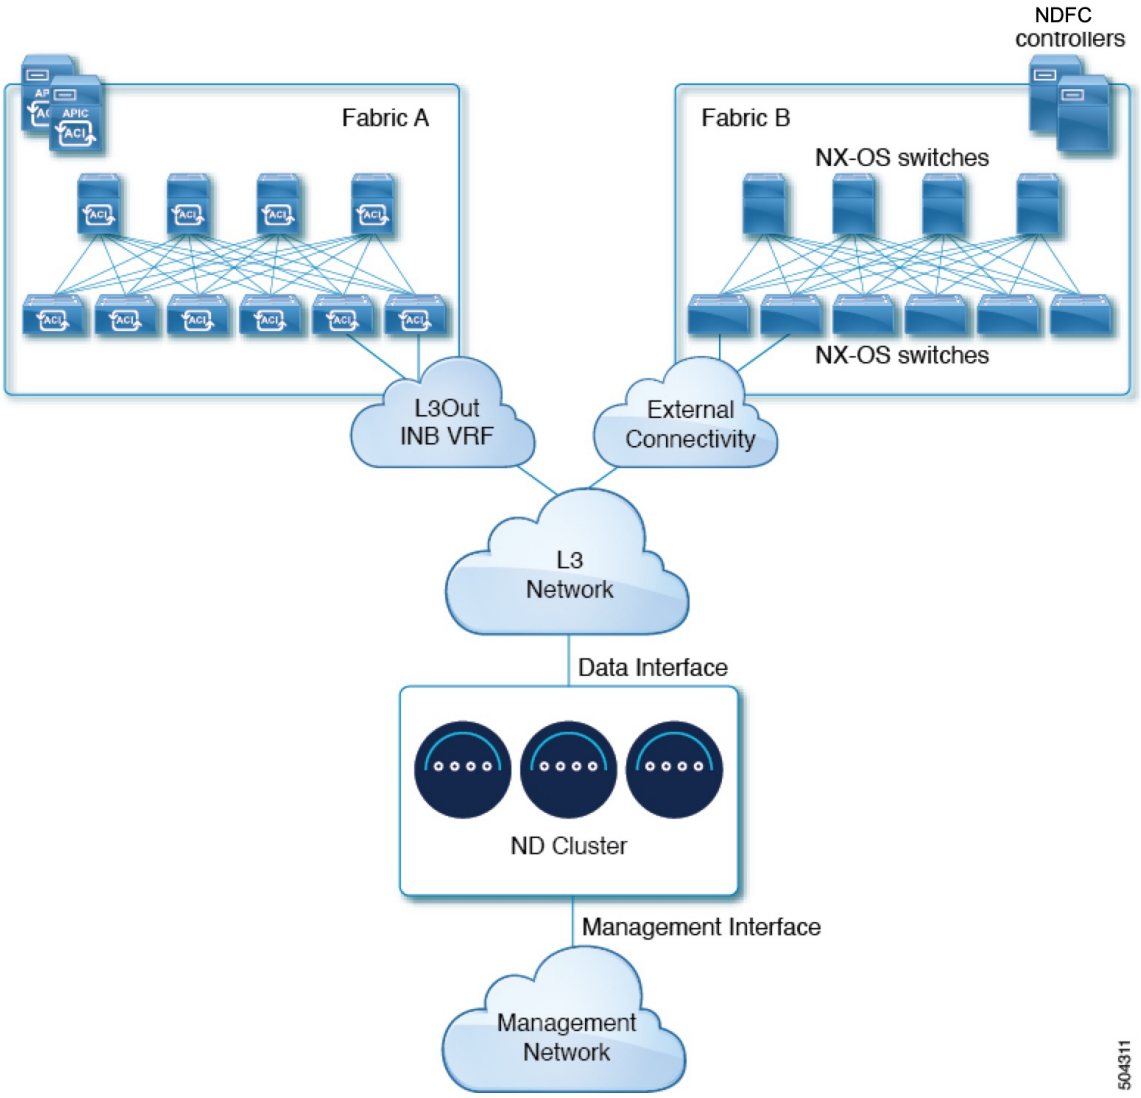
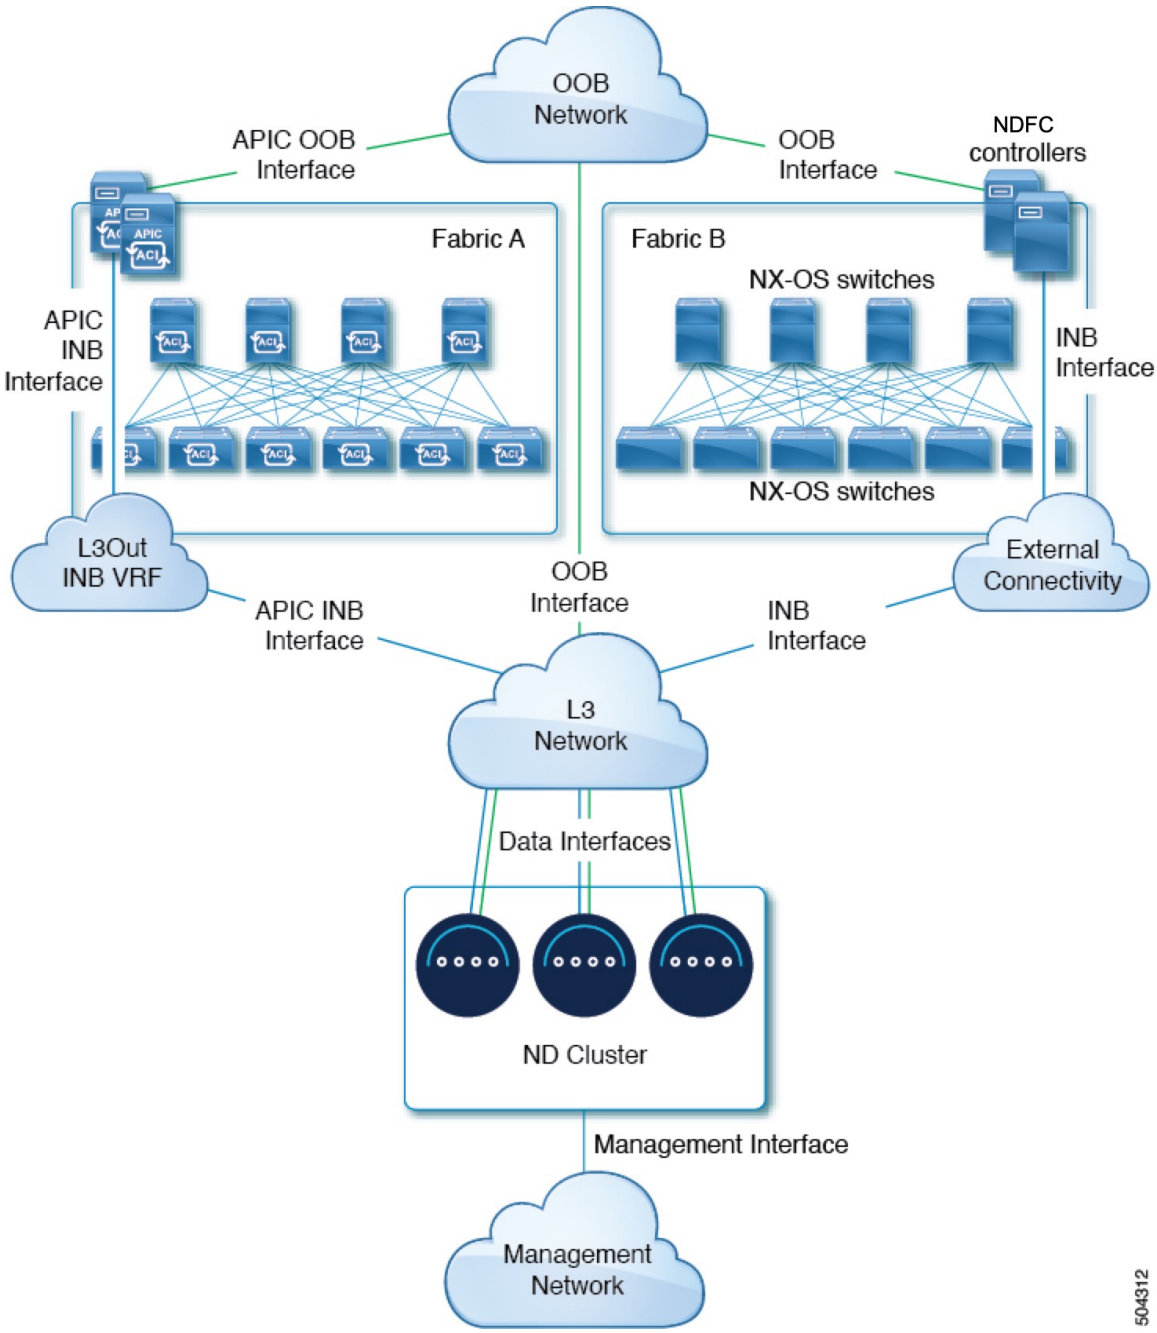


Figure 3: Connecting via Layer 3 Network, Nexus Dashboard Orchestrator



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 Nexus Dashboard 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 Dashboard Insights, 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 Dashboard 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.
- 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.
- 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.
  - You must create a contract between the fabric's in-band management EPG and Cisco Nexus Dashboard EPG.
  - 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.

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.

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

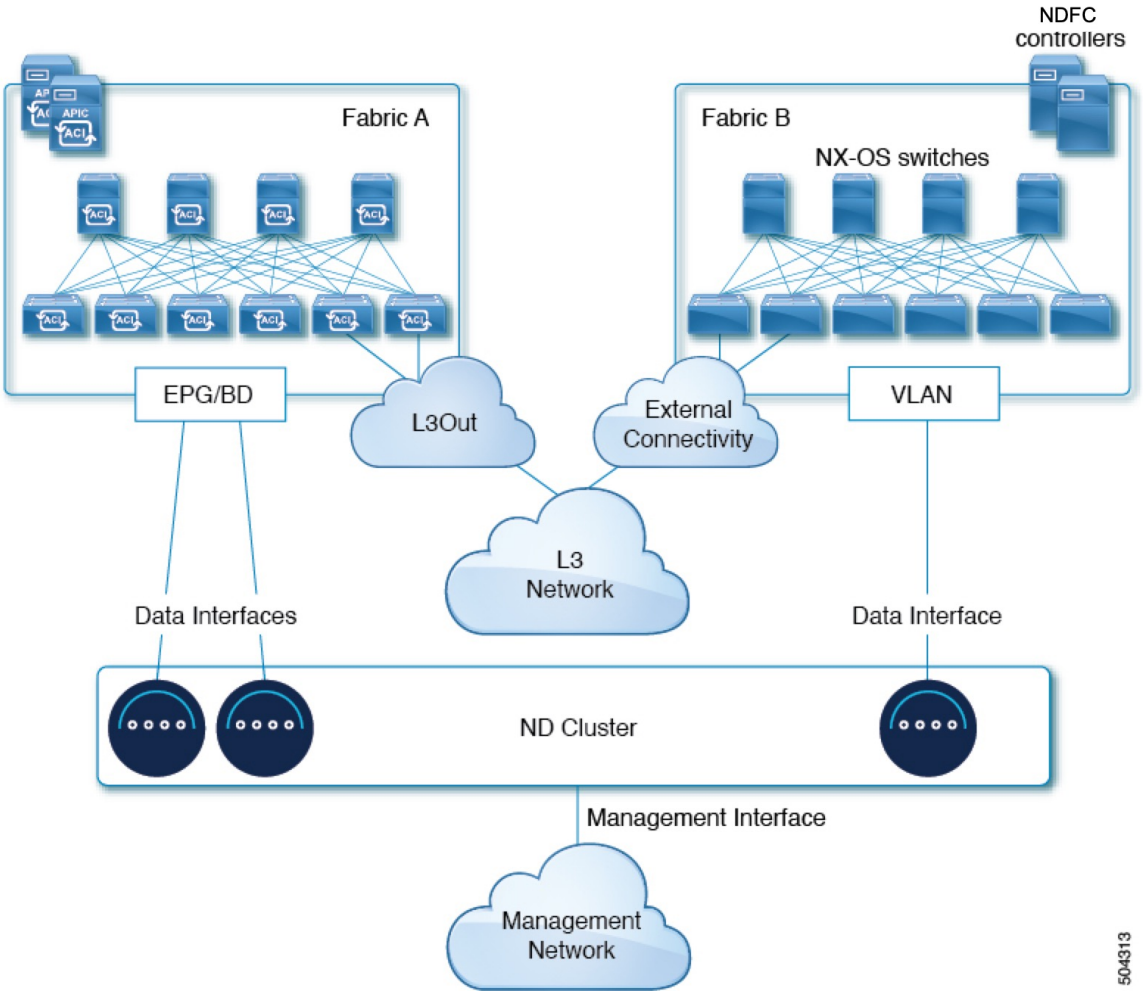
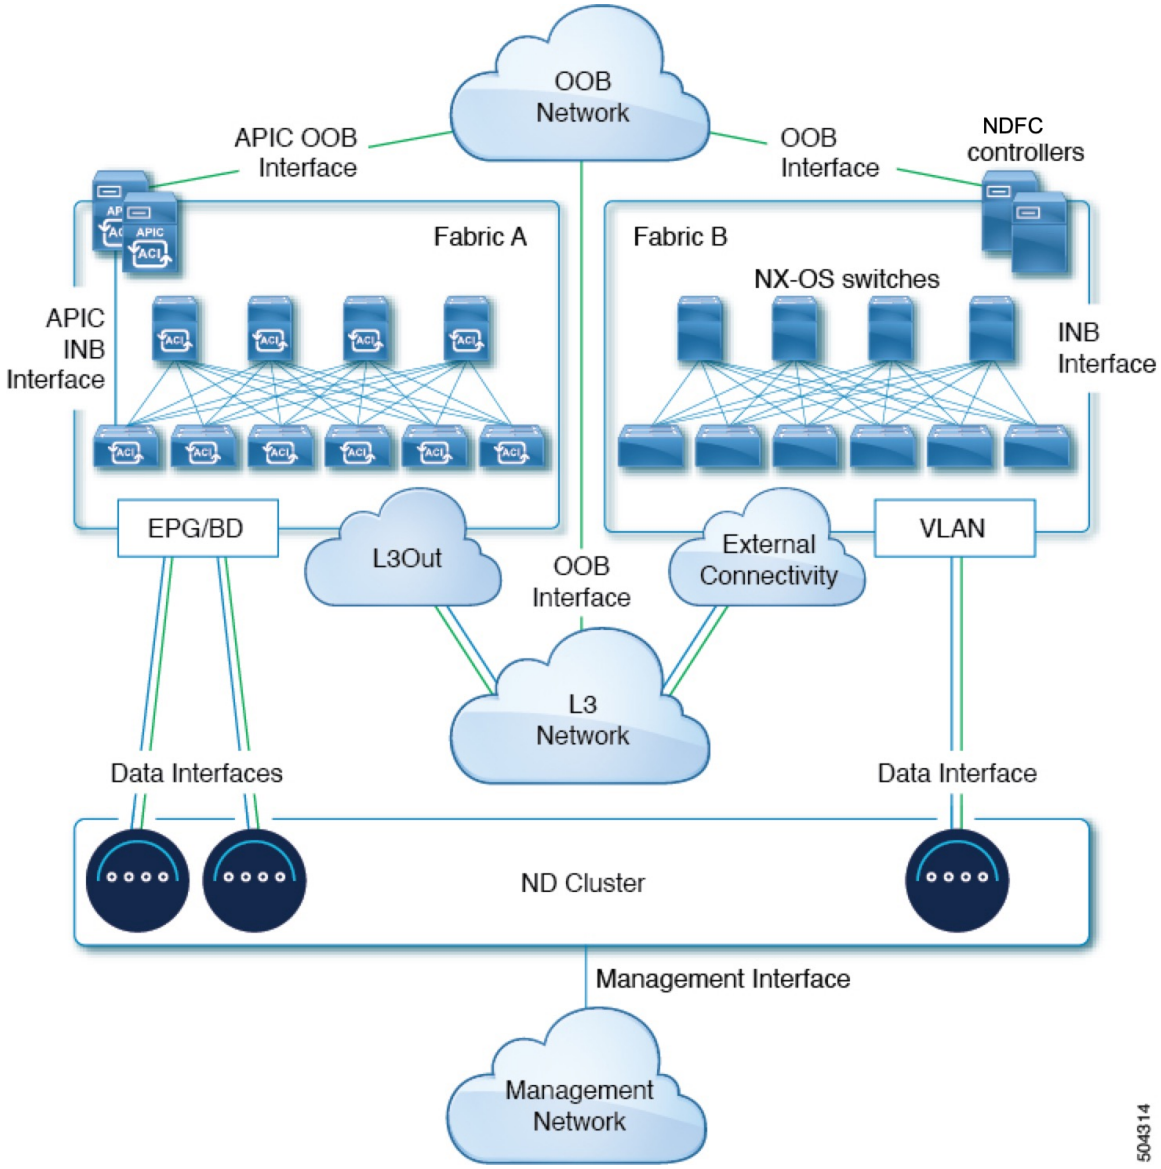


Figure 5: Connecting Directly to Leaf Switches, Nexus Dashboard Orchestrator



Node Distribution Across Sites

Nexus Dashboard supports distribution of cluster nodes across multiple sites. The following node distribution recommendations apply to both physical and virtual clusters.

Note The diagrams in the following sections provide some examples of possible deployment scenarios for physical or virtual Nexus Dashboard cluster nodes. For details on the exact number of nodes required for your specific use case, see the Nexus Dashboard Capacity Planning tool.

Node Distribution for Nexus Dashboard Insights

For Nexus Dashboard Insights, we recommend a centralized, single-site deployment. This service does not support recovery in case if two `primary` nodes are not available and so it gains no redundancy benefits from a distributed cluster, which could instead expose the cluster to interconnection failures when nodes are in different sites.

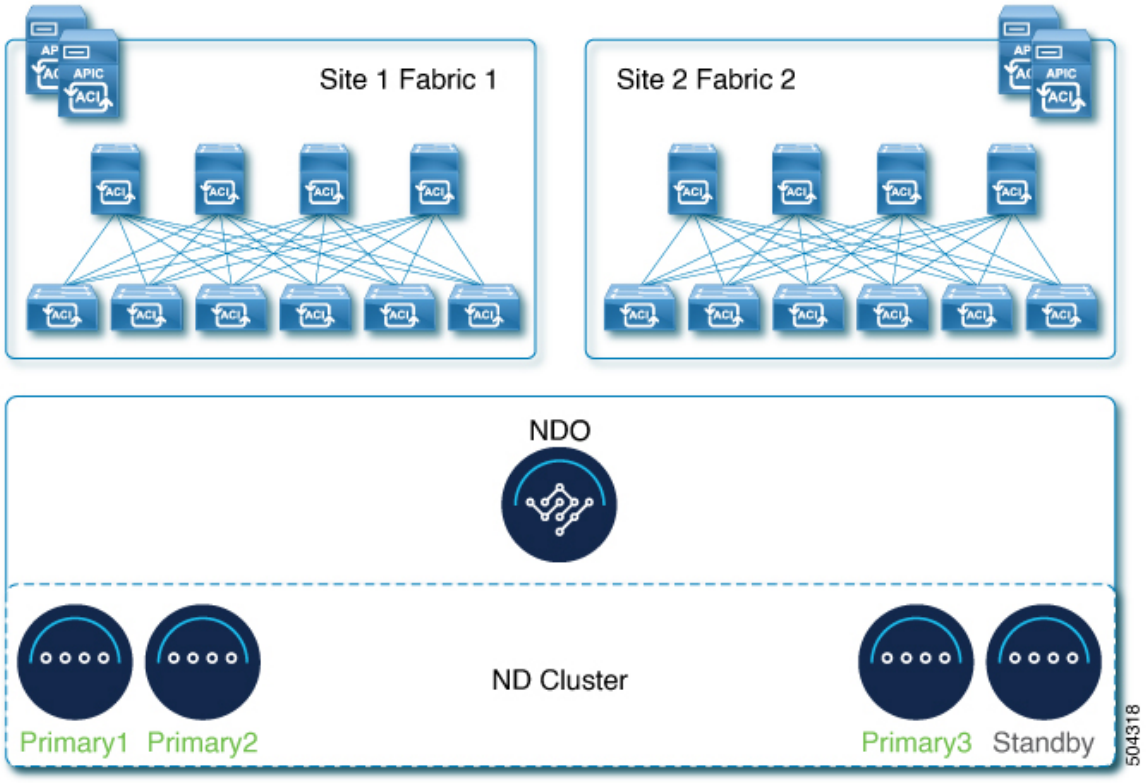
Node Distribution for Fabric Controller

For Nexus Dashboard Fabric Controller, we recommend a centralized, single-site deployment. This service does not support recovery in case if two `primary` nodes are not available and so it gains no redundancy benefits from a distributed cluster, which could instead expose the cluster to interconnection failures when nodes are in different sites.

Node Distribution for Nexus Dashboard Orchestrator

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

Figure 6: Node Distribution Across Two Sites for Nexus Dashboard Orchestrator



Services Co-location Use Cases

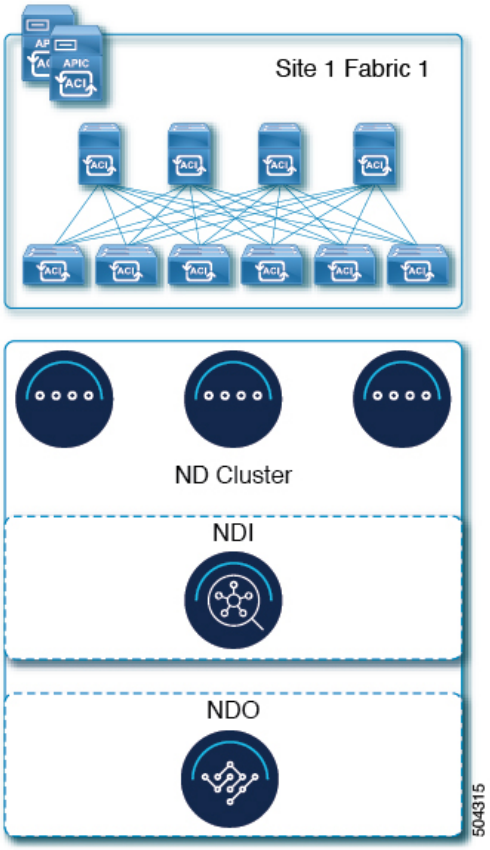
This section describes a number of recommended deployment scenarios for specific single-service or multiple services co-hosting use cases.

Note This release does not support co-hosting services in Nexus Dashboard clusters that are deployed in Linux KVM, AWS, Azure, or RHEL. All services co-hosting scenarios below apply for physical or VMware ESX cluster form factors only. For additional cluster sizing and deployment planning reference information, see the Cisco Nexus Dashboard Cluster Sizing tool.

Single Site, Nexus Dashboard Insights and Orchestrator

In a single site scenario with Nexus Dashboard Insights and Orchestrator services, a single Nexus Dashboard cluster can be deployed with both services co-hosted on it.

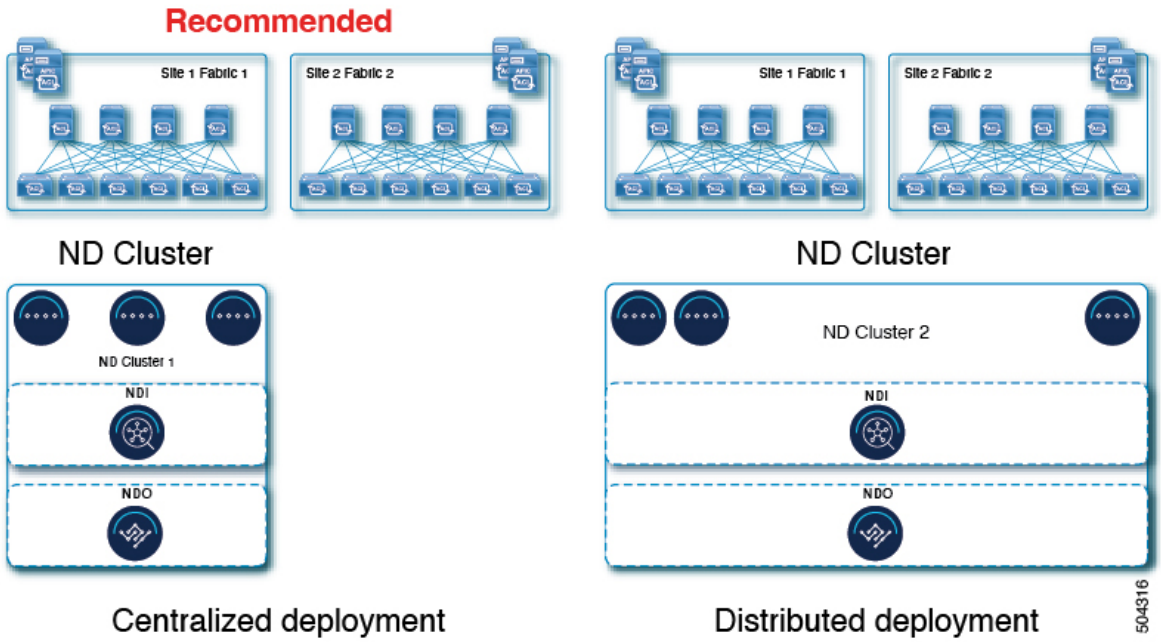
Figure 7: Single Site, Nexus Dashboard Insights and Orchestrator



Multiple Sites, Single Cluster for Nexus Dashboard Insights and Orchestrator

In a multiple sites scenario with Nexus Dashboard Insights and Orchestrator services, a single Nexus Dashboard cluster can be deployed with both services co-hosted on it. In this case, the nodes can be distributed between the sites, however since the Insights service does 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:

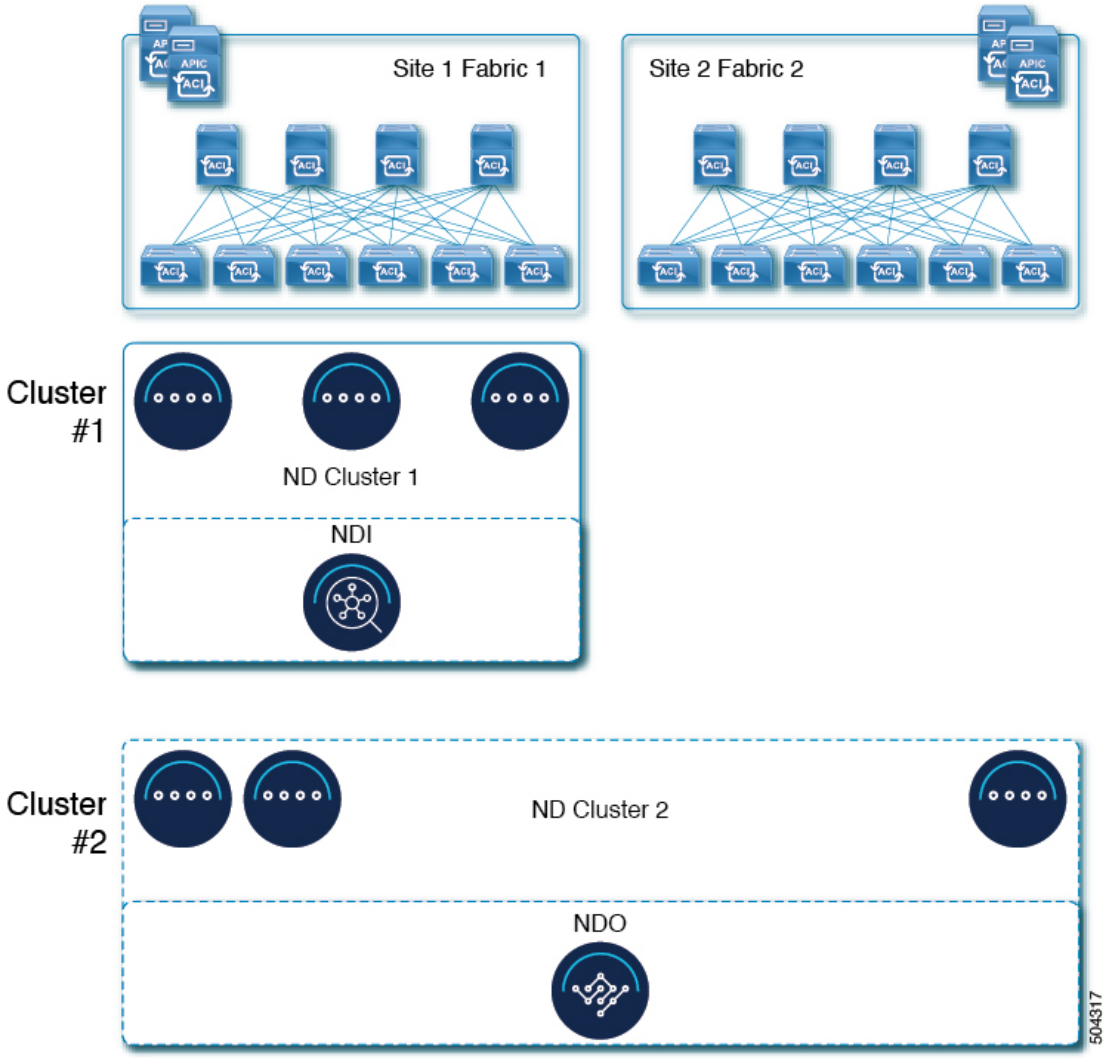
Figure 8: Multiple Sites, Single Cluster for Nexus Dashboard Insights and Orchestrator



Multiple Sites, Multiple Clusters for Nexus Dashboard Insights and Orchestrator

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

Figure 9: Multiple Sites, Multiple Clusters for Nexus Dashboard Insights and Orchestrator



Pre-Installation Checklist

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

Table 12: 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`	

Parameters	Example	Your Entry
DNS Search Domain	`cisco.com`	
App Network	`172.17.0.0/16`	
Service Network	`100.80.0.0/16`	

Table 13: 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`	
If you enable BGP, **ASN** of the first node	`63331`	

Parameters	Example	Your Entry
If you enable BGP and use pure IPv6 deployment, **Router ID** for the first node in the form of an IPv4 address	`1.1.1.1`	

If you enable BGP, IP addresses of the first node's **BGP Peer**(s)	`200.11.11.2`
	or
	`200:11:11::2`	

If you enable BGP, ASNs of the first node's **BGP Peer**(s)	`55555`	
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`	
If you enable BGP, **ASN** of the second node	`63331`	
If you enable BGP and use pure IPv6 deployment, **Router ID** for the second node in the form of an IPv4 address	`2.2.2.2`	

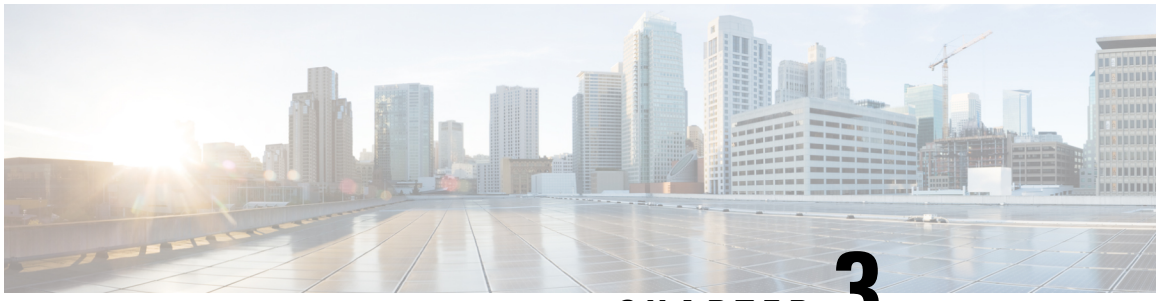
If you enable BGP, IP addresses of the second node's **BGP Peer**(s)	`200.12.12.2`
	or
	`200:12:12::2`	

If you enable BGP, ASNs of the second node's **BGP Peer**(s)	`55555`	
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`	

Parameters	Example	Your Entry
(Optional) **Data Network VLAN** of the third node	`101`	
If you enable BGP, **ASN** of the third node	`63331`	
If you enable BGP and use pure IPv6 deployment, **Router ID** for the third node in the form of an IPv4 address	`3.3.3.3`	

If you enable BGP, IP addresses of the third node's **BGP Peer**(s)	`200.13.13.2`
	or
	`200:13:13::2`	

If you enable BGP, ASNs of the third node's **BGP Peer**(s)	`55555`	



CHAPTER 3

Deploying as Physical Appliance

- Prerequisites and Guidelines, on page 45
- Deploying Nexus Dashboard as Physical Appliance, on page 48

Prerequisites and Guidelines

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

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

  If you are looking to completely re-image the server, for example in case you cannot sign in as the `rescue-user` for manual recovery, see the *Troubleshooting*.

- Ensure that you have enough physical nodes to support your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor and the specific services you plan to deploy. You can use the Nexus Dashboard Capacity Planning tool to verify that the virtual form factor satisfies your deployment requirements.

Note The capacity planning tool lists the total number of primary and worker nodes required. In addition, this form factor supports up to 2 standby nodes.

This document describes how to initially deploy the base Cisco Nexus Dashboard cluster. If you want to expand an existing cluster with another nodes (such as `worker` or `standby`), see the *Infrastructure Management* article instead, which is also available from the Cisco Nexus Dashboard UI.

- Review and complete any additional prerequisites that is described in the *Release Notes* for the services you plan to deploy.

  You can find the service-specific documents at the following links:

  - Nexus Dashboard Fabric Controller Release Notes
  - Nexus Dashboard Insights Release Notes
  - Nexus Dashboard Orchestrator Release Notes

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

  The physical appliance form factor is supported on the UCS-C220-M5 and UCS-C225-M6 original Cisco Nexus Dashboard platform hardware only. The following table lists the PIDs and specifications of the physical appliance servers:

Table 14: Supported UCS-C220-M5 Hardware

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

Table 15: Supported UCS-C225-M6 Hardware

Process ID	Hardware
ND-NODE-L4=	• Cisco UCS C225 M6 Chassis • 2.8-GHz AMD CPU • 256 GB of RAM • 4x 2.4-TB HDDs 960-GB SSD 1.6-TB NVME drive • Intel X710T2LG 2x10 GbE (Copper) Intel E810XXVDA2 2x25/10 GbE (Fiber Optic) Cisco UCS Virtual Interface Card 1455 (4x25G Ports) • 1050-W power supply
ND-CLUSTER-L4	A cluster of 3x `ND-NODE-L4=` appliances.

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

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

  The minimum that is supported and recommended versions of CIMC are listed in the "Compatibility" section of the *Release Notes* for your Cisco Nexus Dashboard release.

- Ensure that you have configured an IP address for the server's CIMC.

  To configure a CIMC IP address:

  1. Power on the server.

     After the hardware diagnostic is complete, you will be prompted with different options controlled by the function (Fn) keys.

  2. Press the **F8** key to enter the **Cisco IMC configuration Utility**.

  3. Provide the following information:

     - Set **NIC mode** to `Dedicated`.
     - Choose between the **IPv4** and **IPv6** IP modes.

       You can choose to enable or disable DHCP. If you disable DHCP, provide the static IP address, subnet, and gateway information.

     - Ensure that **NIC Redundancy** is set to `None`.
     - Press **F1** for more options such as hostname, DNS, default user passwords, port properties, and reset port profiles.

  4. Press **F10** to save the configuration and then restart the server.

- Ensure that Serial over LAN (SoL) is enabled in CIMC.

  SoL is required for the `connect host` command, which you use to connect to the node to provide basic configuration information. To use the SoL, you must first enable it on your CIMC. SSH into the node using the CIMC IP address and enter the sign-in credentials. Run the following commands:

```
Server# scope sol
Server /sol # set enabled yes
Server /sol *# set baud-rate 115200
Server /sol *# commit
Server /sol *#
Server /sol # show

C220-WZP23150D4C# scope sol
C220-WZP23150D4C /sol # show

Enabled Baud Rate(bps)  Com Port SOL SSH Port
------- --------------- -------- -------------
yes     115200          com0     2400


C220-WZP23150D4C /sol #
```

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

  For example, if the hardware you received came with Release 2.2.1 image pre-installed, but you want to deploy Release 3.0.1 instead, we recommend:

  1. First, bring up the Release 2.2.1 cluster, as described in the deployment guide for that release.
  2. Then upgrade to Release 3.0.1, as described in Upgrading Nexus Dashboard, on page 133.

Note For brand new deployments, you can also choose to simply re-image the nodes with the latest version of the Cisco Nexus Dashboard (for example, if the hardware came with an image which does not support a direct upgrade to this release through the GUI workflow) before returning to this document for deploying the cluster. This process is described in the "Re-Imaging Nodes" section of the *Troubleshooting* article for this release.

- You must have at least a 3-node cluster. Extra worker nodes can be added for horizontal scaling if required by the enter and number of services you deploy. For the maximum number of `worker` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

Deploying 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 45.

Procedure

Step 1 Configure the first node's basic information.

You must configure only a single ("first") node as described in this step. Other nodes will be configured during the GUI-based cluster deployment process described in the following steps and will accept settings from the first `primary` node. The other two `primary` nodes do not require any additional configuration besides ensuring that their CIMC IP addresses are reachable from the first `primary` node and login credentials are set.

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
```

Note
If you want to configure pure IPv6 mode, provide the IPv6 in the above example instead.

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.

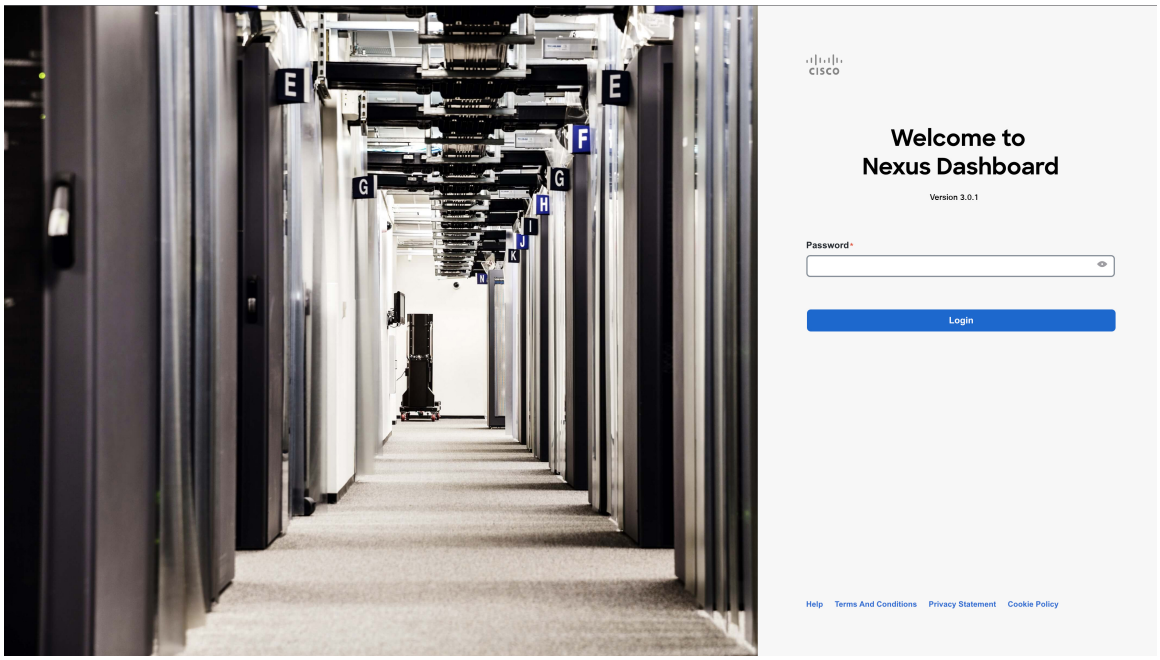
```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

System UI online, please login to https://192.168.9.172 to continue.
```

Step 3 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

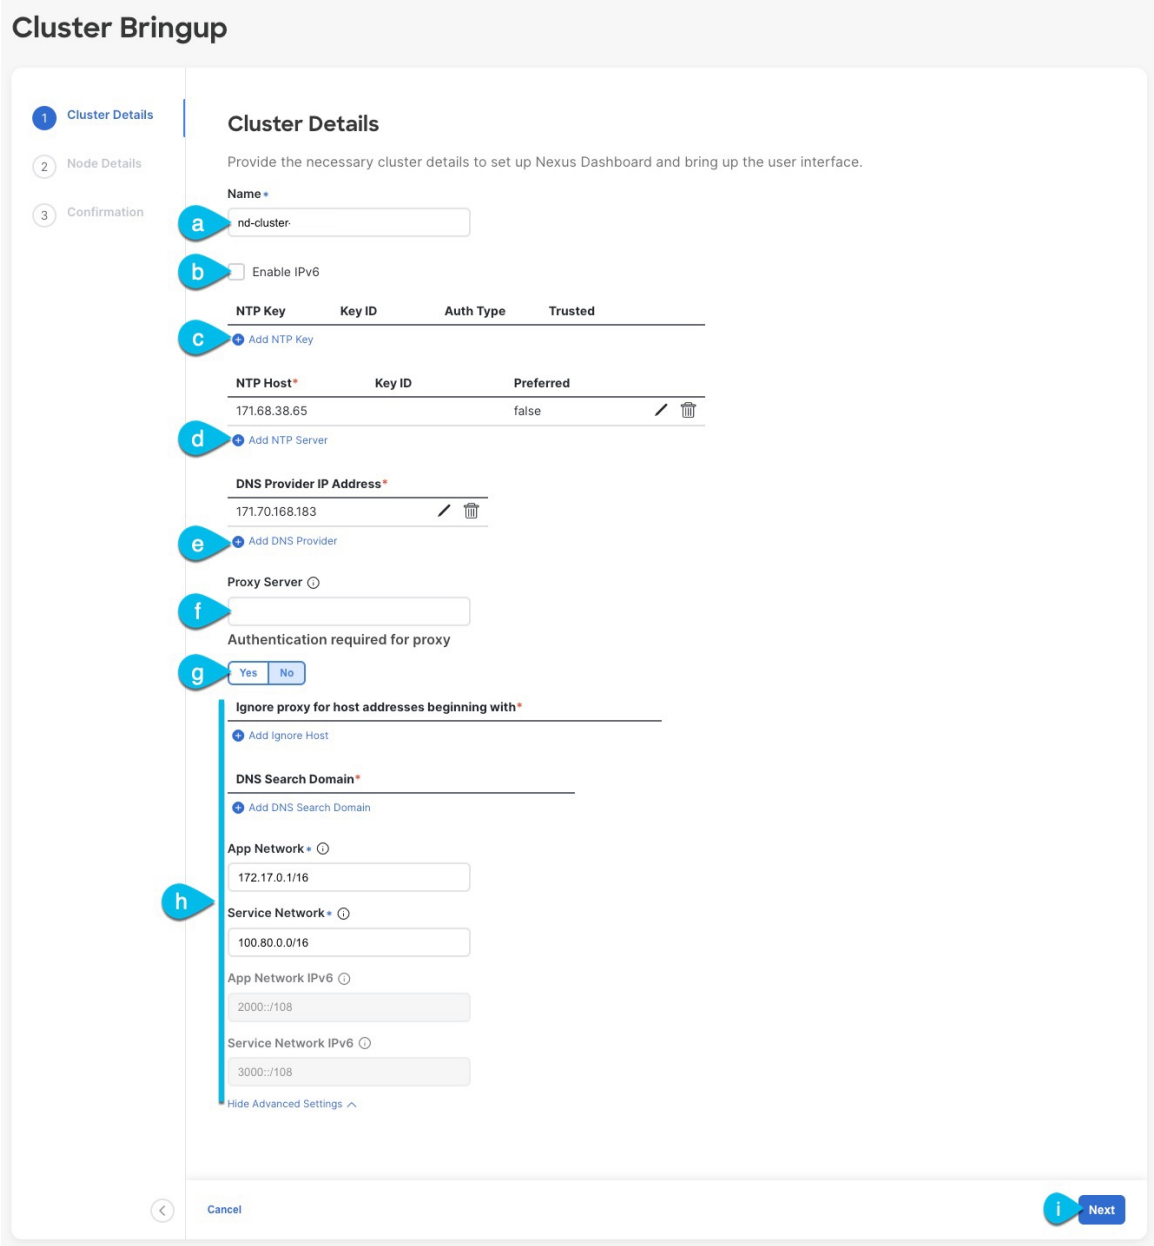
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Login**



Step 4 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

b) (Optional) If you want to enable IPv6 functionality for the cluster, check the **Enable IPv6** checkbox.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

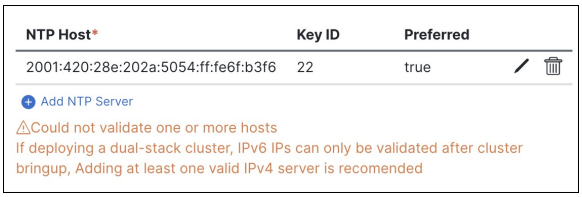
d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

Note
If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and provided an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

The proxy server must have the following URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
```

```
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

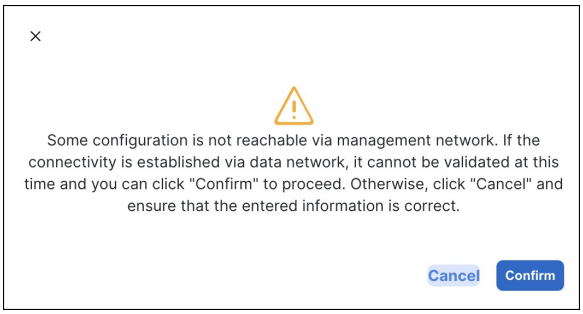
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

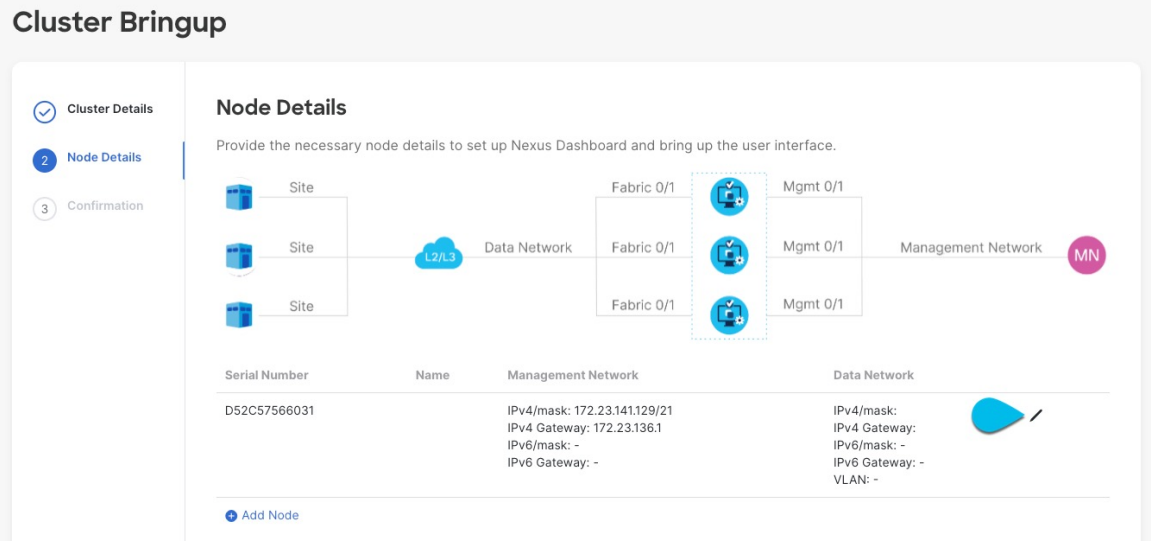
i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



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

You have defined the Management network and IP address for the node into which you are currently logged in during the initial node configuration in earlier steps, but you must also provide the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.



a) Click the **Edit** button next to the first node.

b) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated.

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

c) 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, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

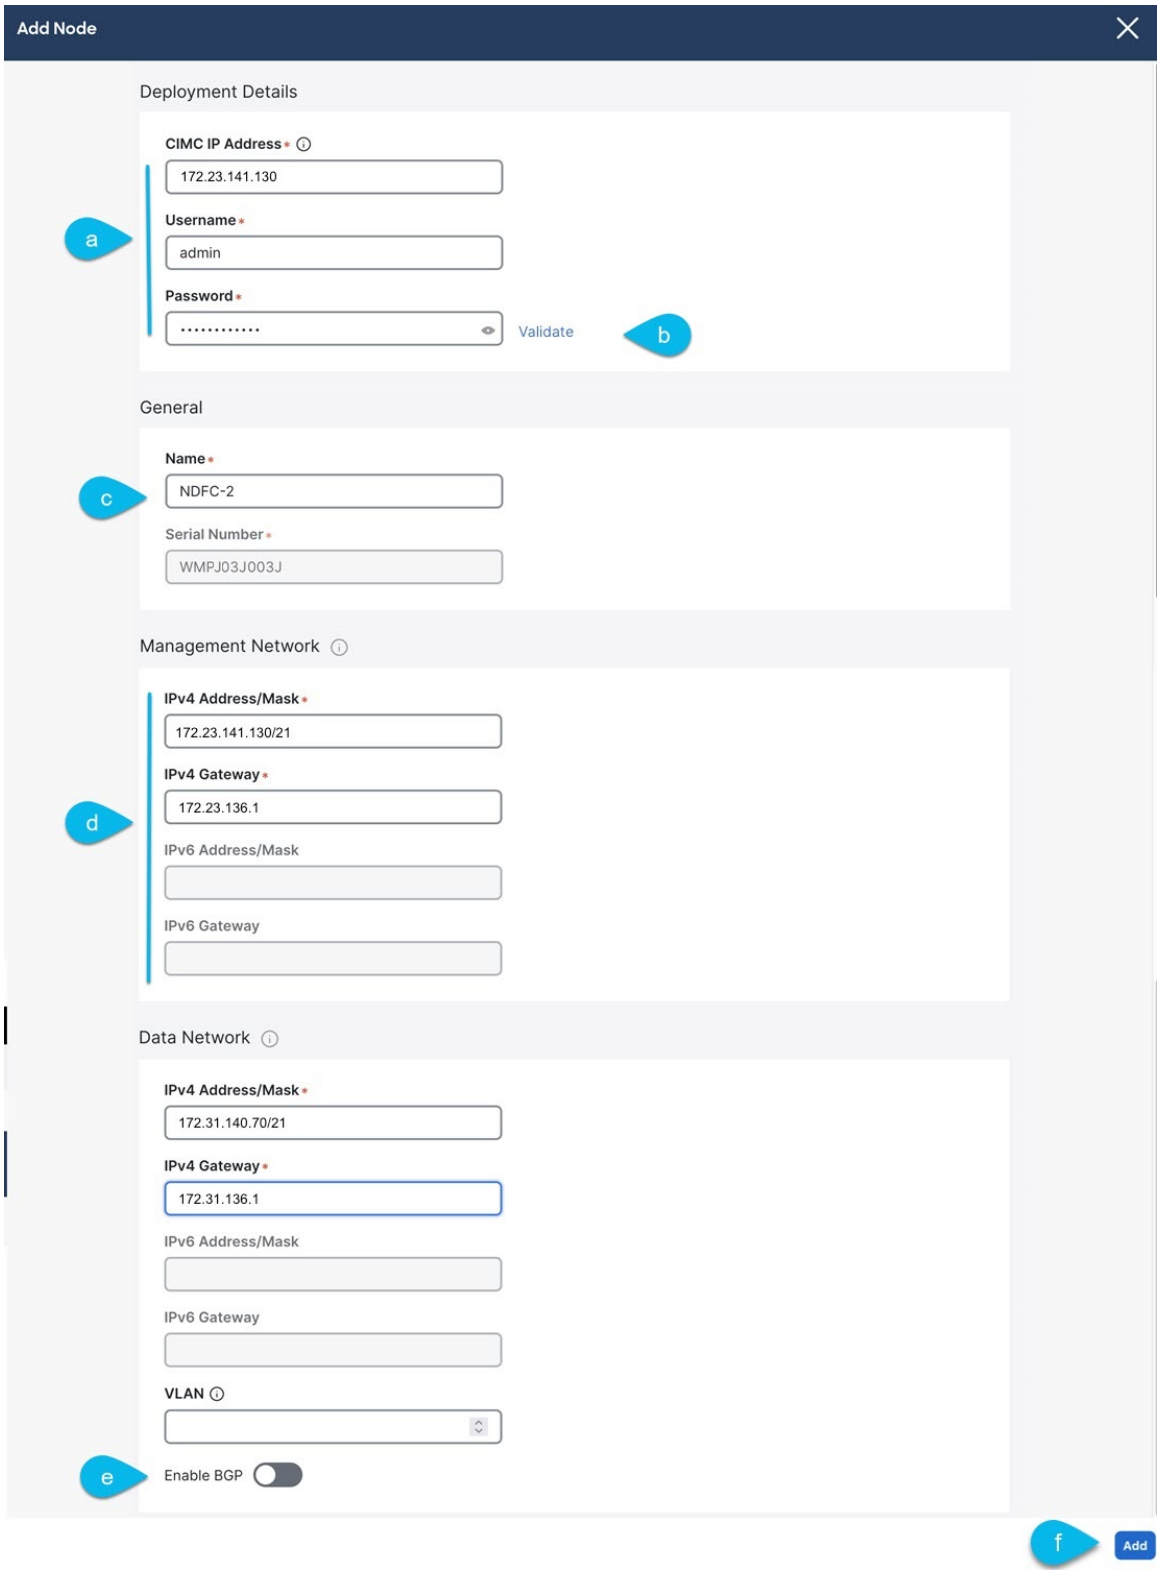
  The router ID must be an IPv4 address, for example `1.1.1.1`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

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

Step 6 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 **CIMC IP Address**, **Username**, and **Password** for the second node.
b) Click **Validate** to verify connectivity to the node.

After network connectivity is validated, you can provide the other required information for the node.

c) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated after CIMC connectivity is validated.

d) In the **Data Network** area, provide the node's **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

  The router ID must be an IPv4 address, for example `2.2.2.2`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

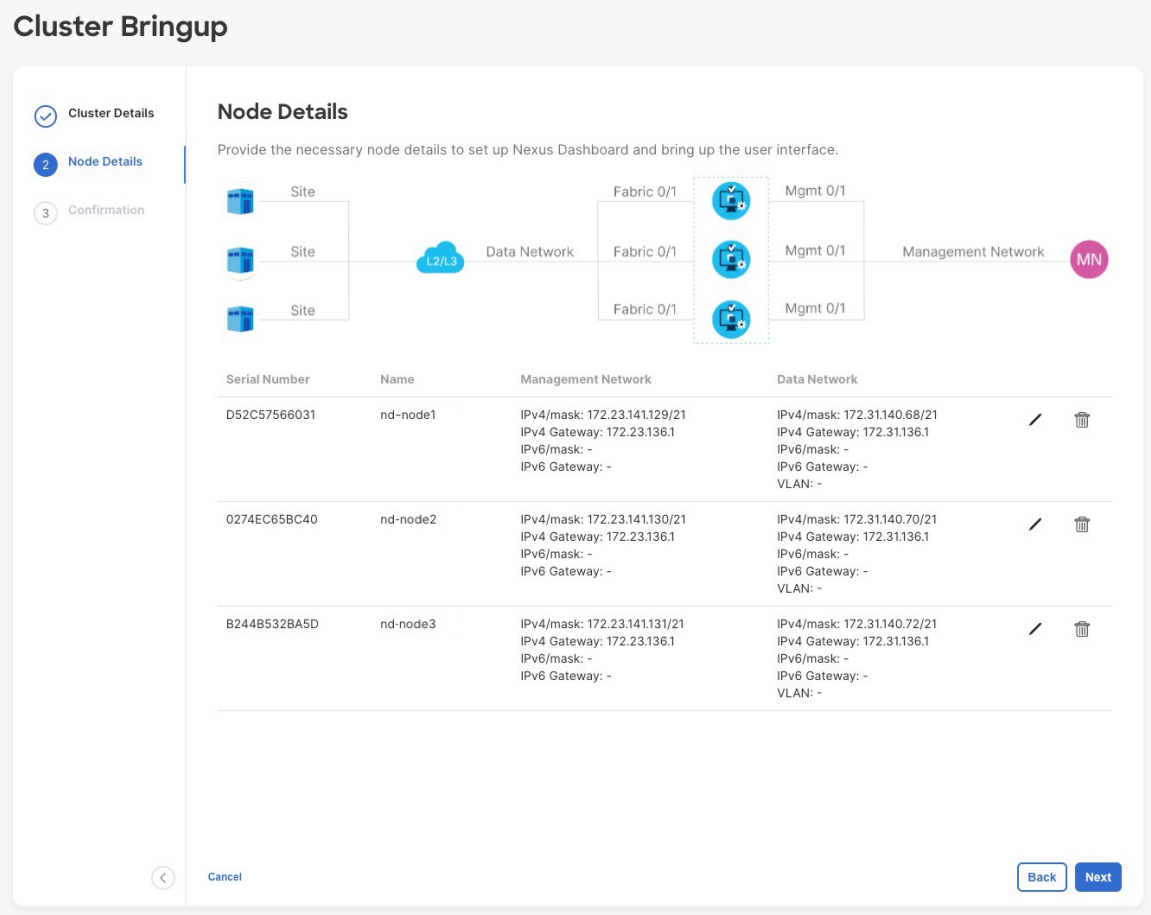
f) Click **Add** to save the changes.

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

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

Step 8 In the **Node Details** page, click **Next** to continue.

After you have provided the management and data network information for all nodes, you can proceed to the final **Confirmation** screen.



Step 9 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 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 as the `rescue-user` using the password you provided during node deployment 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:

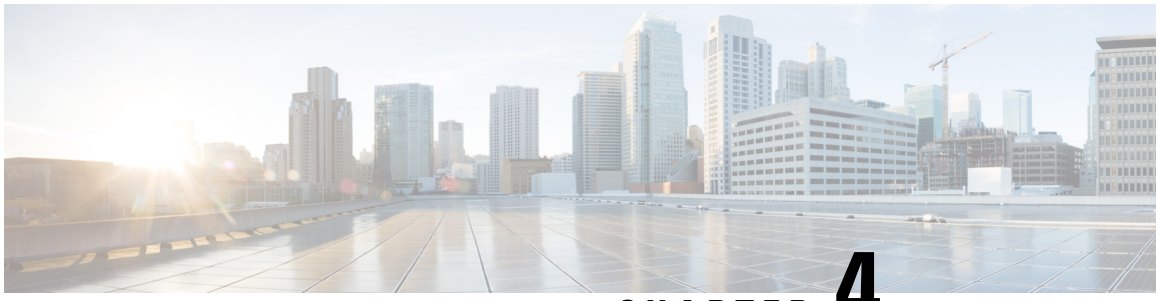
```
$ acs health
All components are healthy
```

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 Configure the **Network Scale** parameters for your cluster.

This is described in the **Infrastructure Management** > **Cluster Configuration** section of the *Cisco Nexus Dashboard User Guide*, which is also available directly from your Nexus Dashboard's Help Center.



CHAPTER 4

Deploying in VMware ESX

- Prerequisites and Guidelines, on page 61
- Deploying Nexus Dashboard Using VMware vCenter, on page 64
- Deploying Nexus Dashboard Directly in VMware ESXi, on page 80

Prerequisites and Guidelines

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

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

  Scale and services support and co-hosting vary based on the cluster form factor and the specific services you plan to deploy. You can use the Nexus Dashboard Capacity Planning tool to verify that the virtual form factor satisfies your deployment requirements.

Note This document describes how to initially deploy the base Cisco Nexus Dashboard cluster. If you want to expand an existing cluster with additional `worker` nodes, see the *Infrastructure Management* article instead, which is also available from the Cisco Nexus Dashboard UI.

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

Standby nodes are not supported with this cluster form factor.

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

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

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.

- The disk must have I/O latency of 20ms or less.
- When deploying in VMware ESX, you can deploy two types of nodes:
  - Data Node—node profile with higher system requirements designed for specific services that require the additional resources.
  - App Node—node profile with a smaller resource footprint that can be used for most services.



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

Ensure you have enough system resources:

Table 16: Deployment Requirements

Data Node Requirements	App Node Requirements

• VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3
• VMware vCenter 7.0.1, 7.0.2, 7.0.3 if deploying using vCenter
• Each VM requires the following:
  • 32 vCPUs with physical reservation of at least 2.2GHz
  • 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
	• VMware vCenter 7.0.1, 7.0.2, 7.0.3 if deploying using vCenter
	• Each VM requires the following:
	  • 16 vCPUs with physical reservation of at least 2.2GHz
	  • 64GB of RAM with physical reservation
	  • 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume
	  Some services require `App` nodes to be deployed on faster SSD storage while other services support HDD. Check the Nexus Dashboard Capacity Planning tool to ensure that you use the correct type of storage.
	  Note
	  Beginning with Nexus Dashboard release 3.0(1i) and Nexus Dashboard Insights release 6.3(1), you can use the OVA-App node profile for the Insights service. However, you must change from the default 500GB disk requirement to 1536GB when deploying node VMs which will be used for hosting Insights.
	• 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 vCenter for Virtual Guest VLAN Tagging (VGT) mode.

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

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

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

Snapshots of powered on VMs are not supported.

- 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.

- This release does not support VMware snapshots of the Nexus Dashboard VMs.
- Cisco does not support the use of nested virtualization environments. Deploying Nexus Dashboard on a virtual machine that is itself running on a virtualized hypervisor (for example, KVM on ESXi) is not a supported configuration and may result in performance degradation or system instability.
- You can choose to deploy the nodes directly in ESXi or using vCenter.

  If you want to deploy using vCenter, following the steps described in Deploying Nexus Dashboard Using VMware vCenter, on page 64.

  If you want to deploy directly in ESXi, following the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 80.



Note If you plan to deploy Nexus Dashboard Insights using the OVA-App node profile, you must deploy using vCenter.

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

Deploying Nexus Dashboard Using VMware vCenter

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter. If you prefer to deploy directly in ESXi, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 80 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 61.

Procedure

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

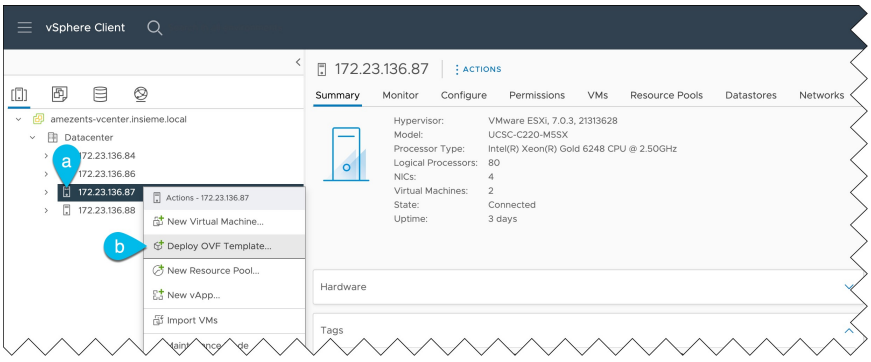
https://software.cisco.com/download/home/286327743/type/286328258/

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.

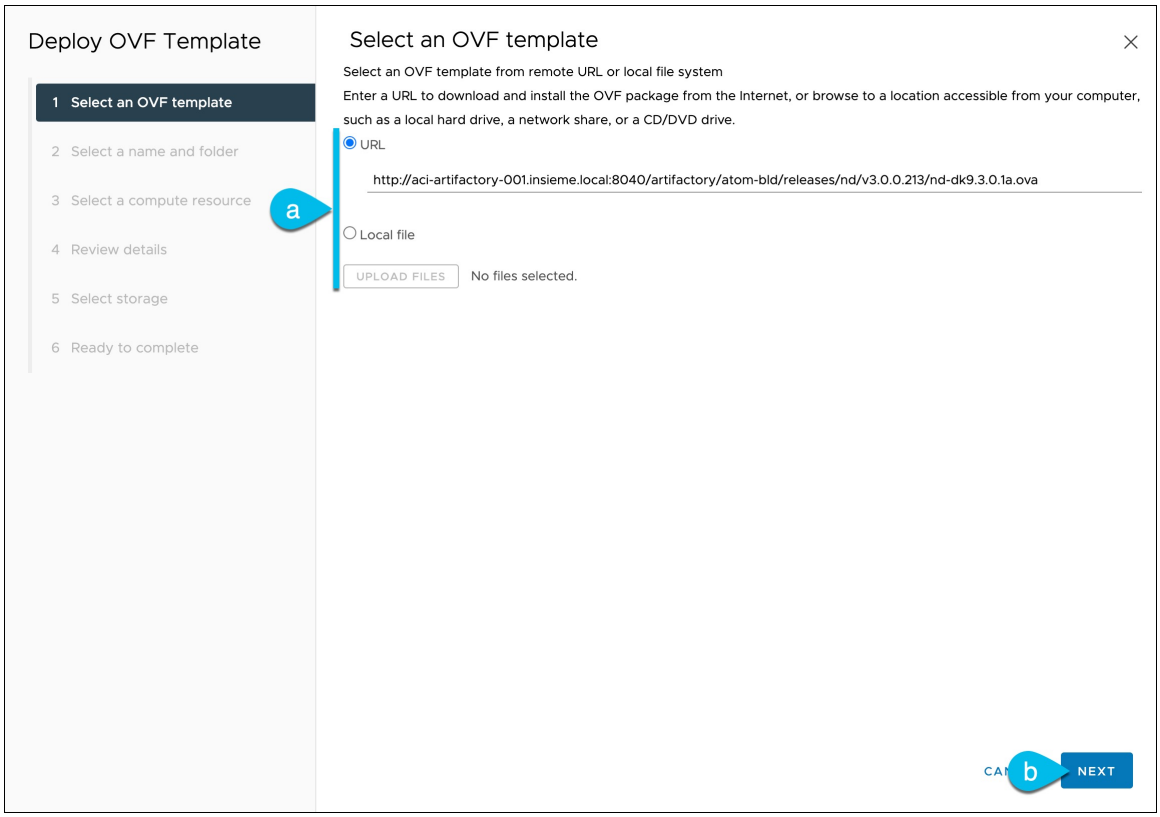
Step 3 Start the new VM deployment.



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.



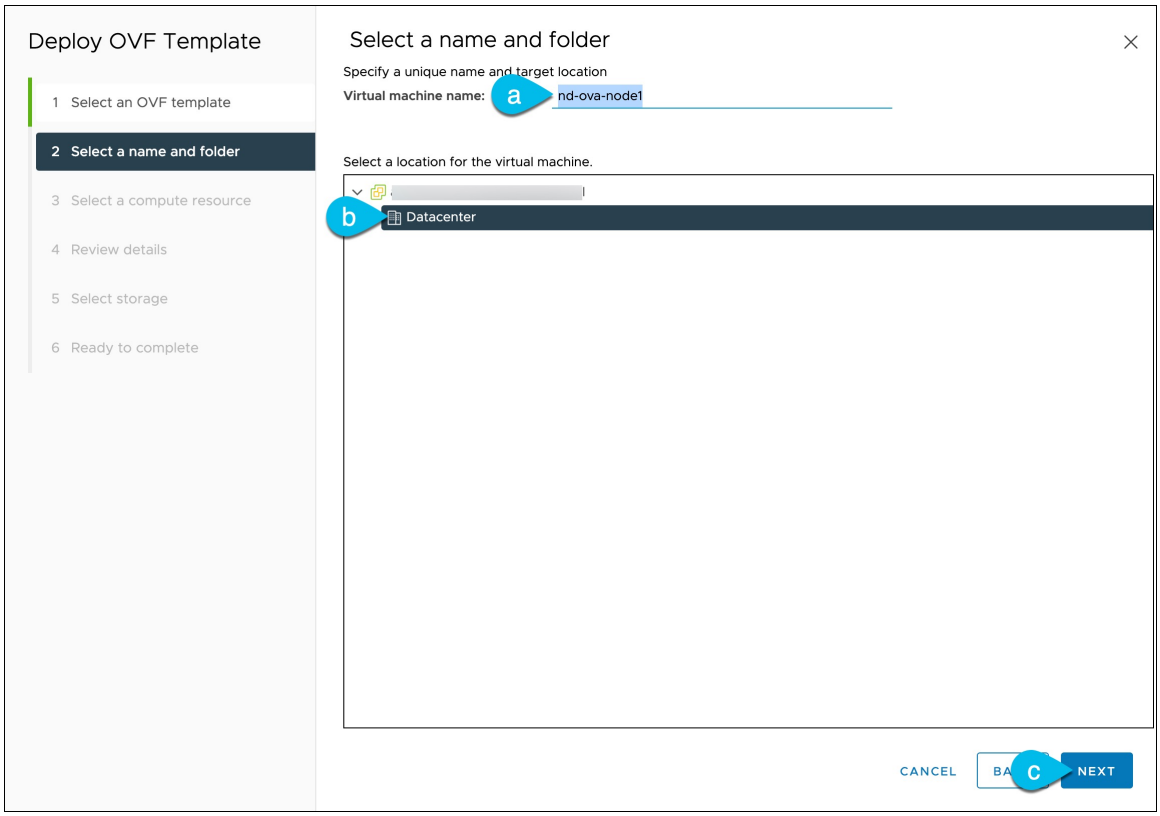
a) Provide the location of the image.

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.



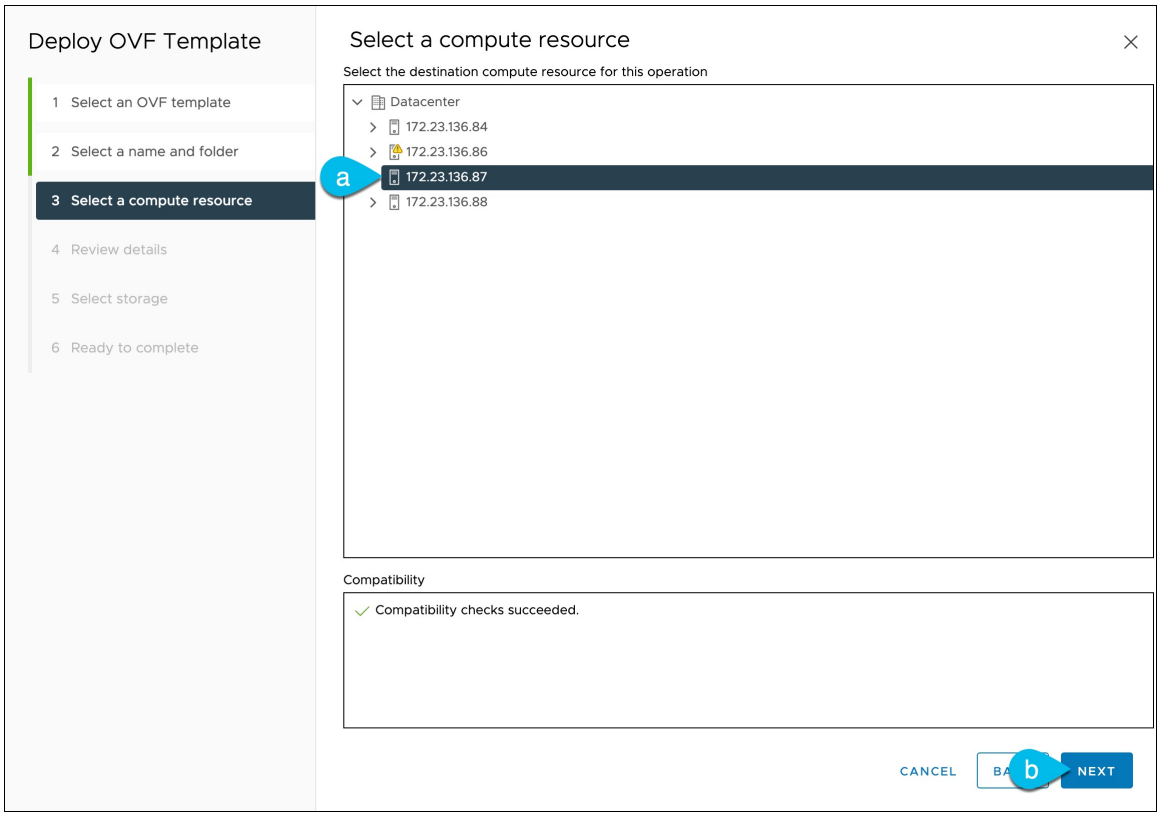
a) Provide the name for the virtual machine.

   For example, `nd-ova-node1`.

b) Select the location for the virtual machine.

c) Click **Next** to continue

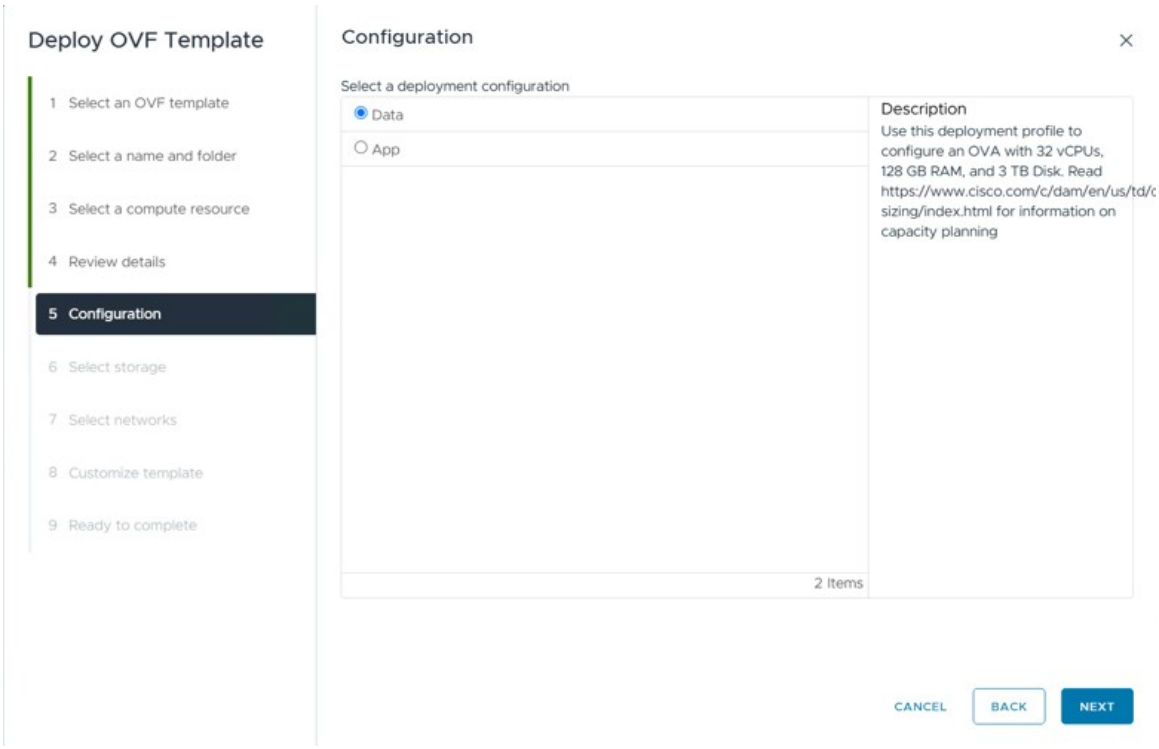
Step 6 In the **Select a compute resource** screen, select the ESX host.



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.

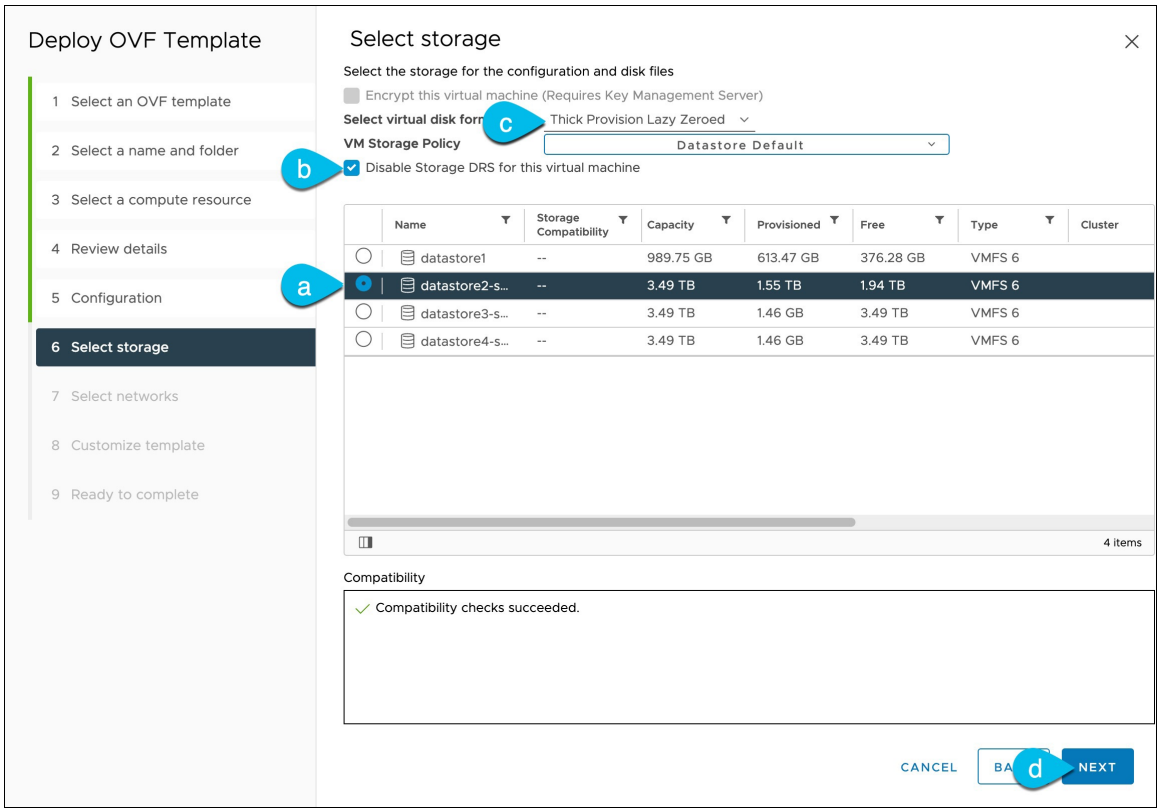


a) Select either `App` or `Data` node profile based on your use case requirements.

For more information about the node profiles, see Prerequisites and Guidelines, on page 61.

b) Click **Next** to continue

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



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

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

Nexus Dashboard does not support VMware DRS. We recommend that you check the **Disable Storage DRS for this virtual machine** option in case DRS is enabled at the ESXi cluster level.

c) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

d) Click **Next** to continue

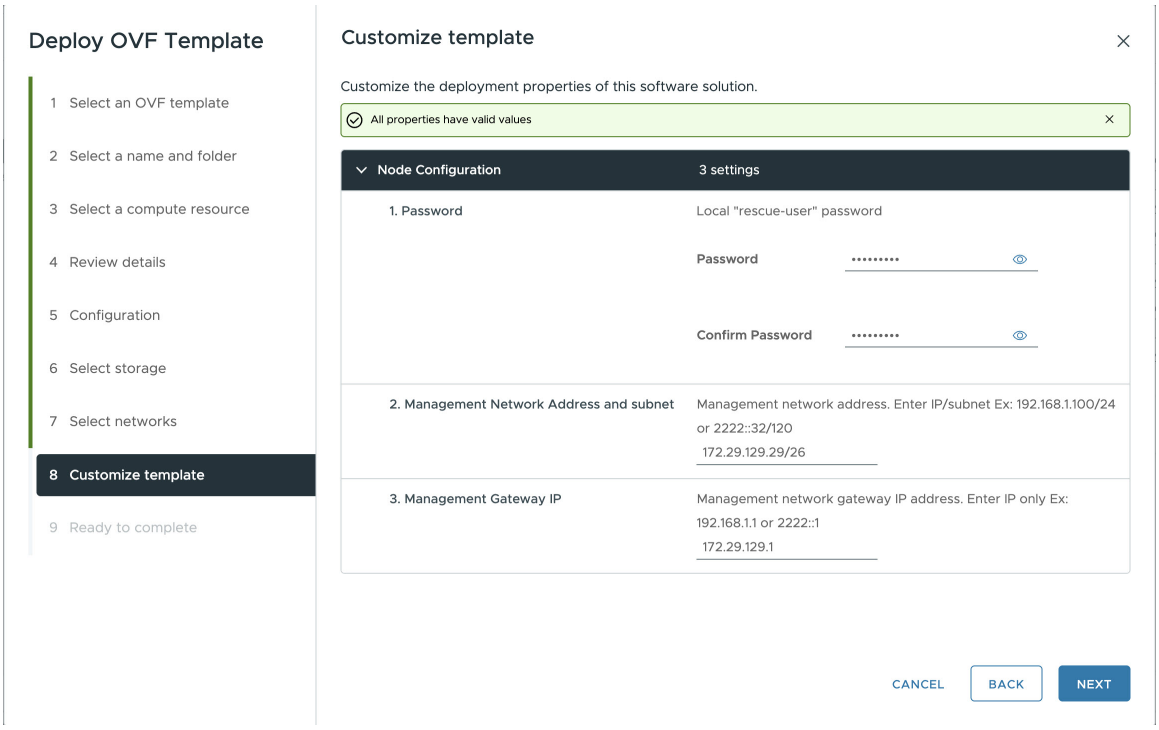
Step 10 In the **Select networks** screen, choose the VM network for the Nexus Dashboard's Management and Data networks and click **Next** to continue.

There are two networks required by the Nexus Dashboard cluster:

- **fabric0** is used for the Nexus Dashboard cluster's Data Network
- **mgmt0** is used for the Nexus Dashboard cluster's Management Network.

For more information about these networks, see Prerequisites and Guidelines, on page 6 in the "Deployment Overview and Requirements" chapter.

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



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

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

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

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

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 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'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.

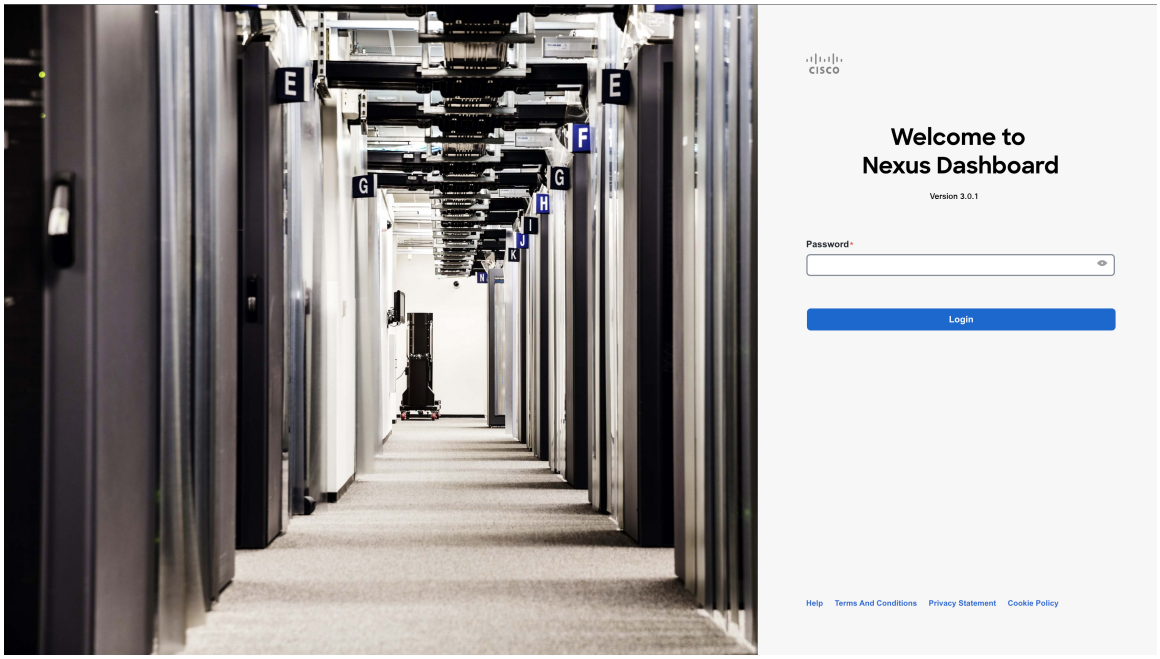
To disable time synchronization:

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 guest time with host`** option.

Step 16 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

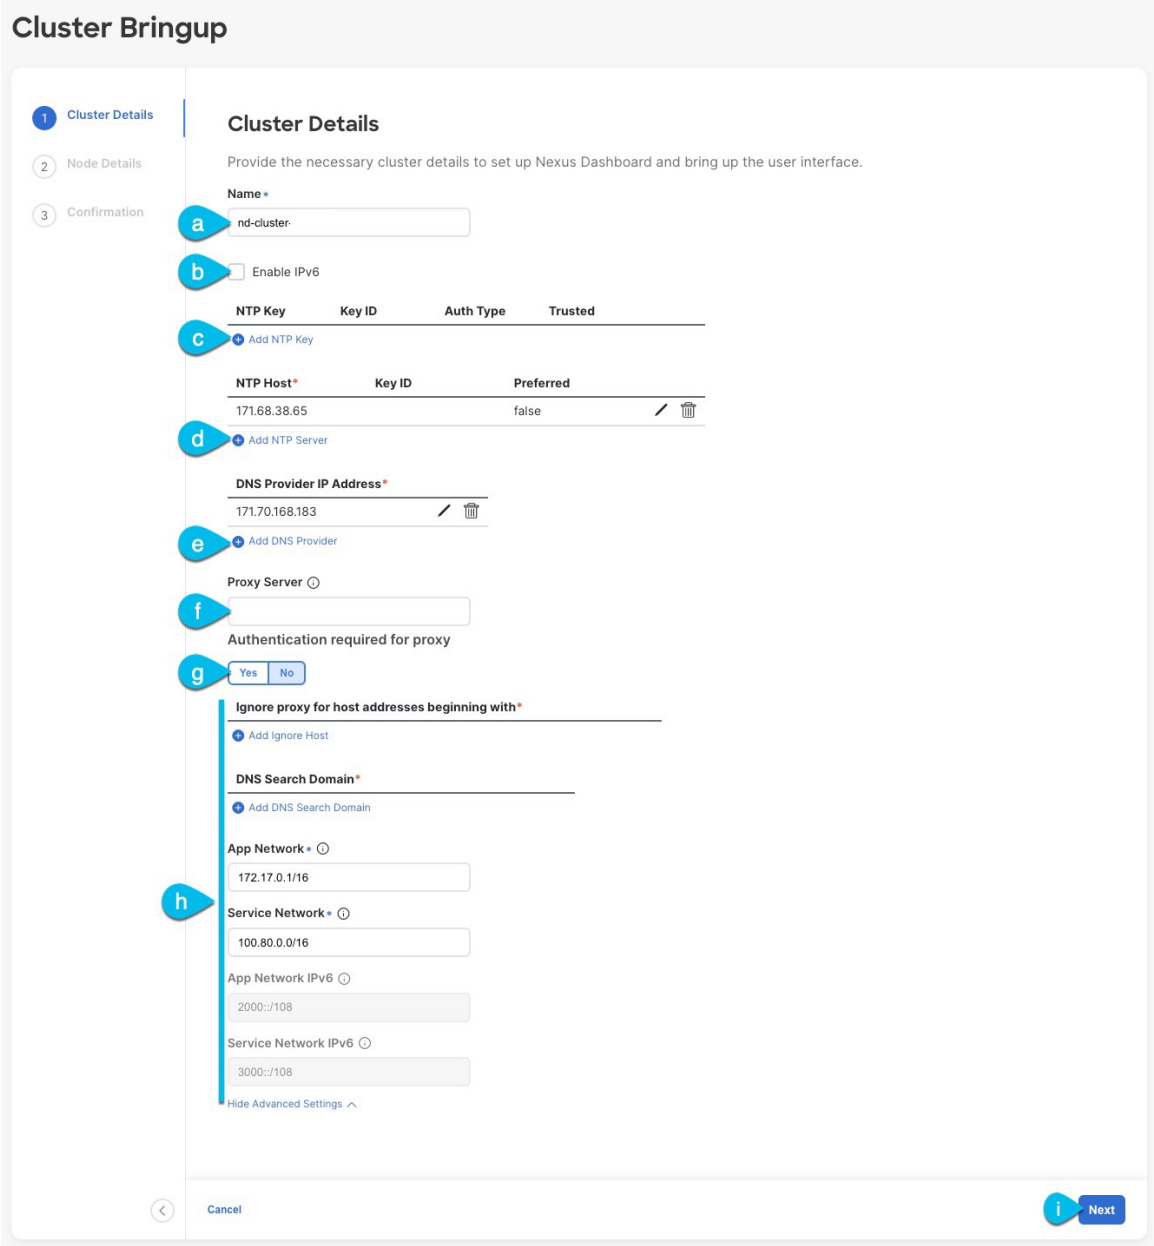
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Login**



Step 17 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

b) (Optional) If you want to enable IPv6 functionality for the cluster, check the **Enable IPv6** checkbox.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

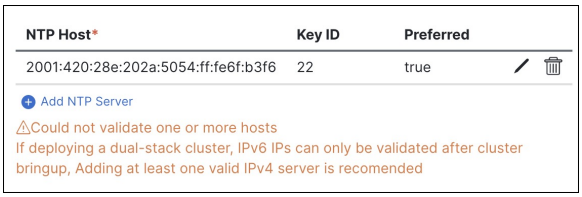
d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

Note
If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and provided an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

The proxy server must have the following URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
```

```
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

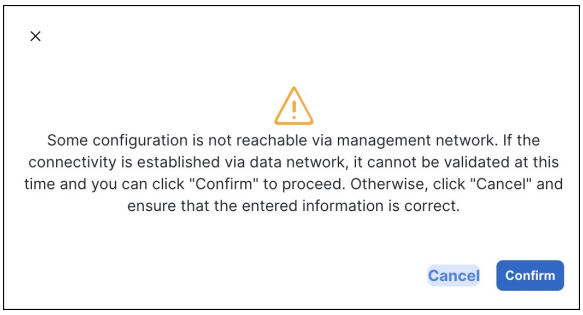
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

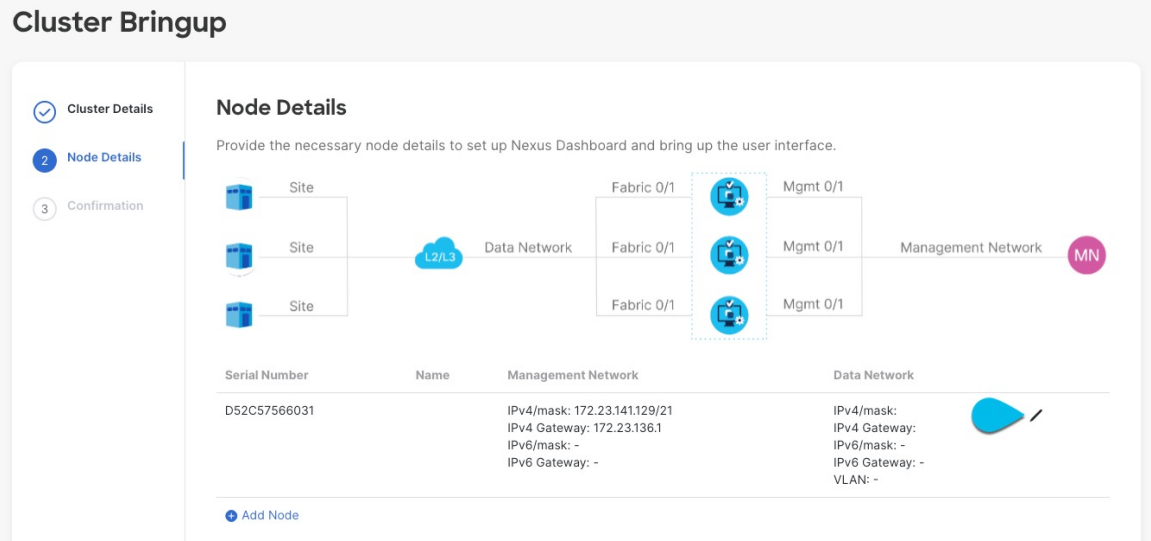
i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



Step 18 In the **Node Details** screen, update the current node's information.

You have defined the Management network and IP address for the node into which you are currently logged in during the initial node configuration in earlier steps, but you must also provide the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.



a) Click the **Edit** button next to the first node.

b) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated.

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

c) 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, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

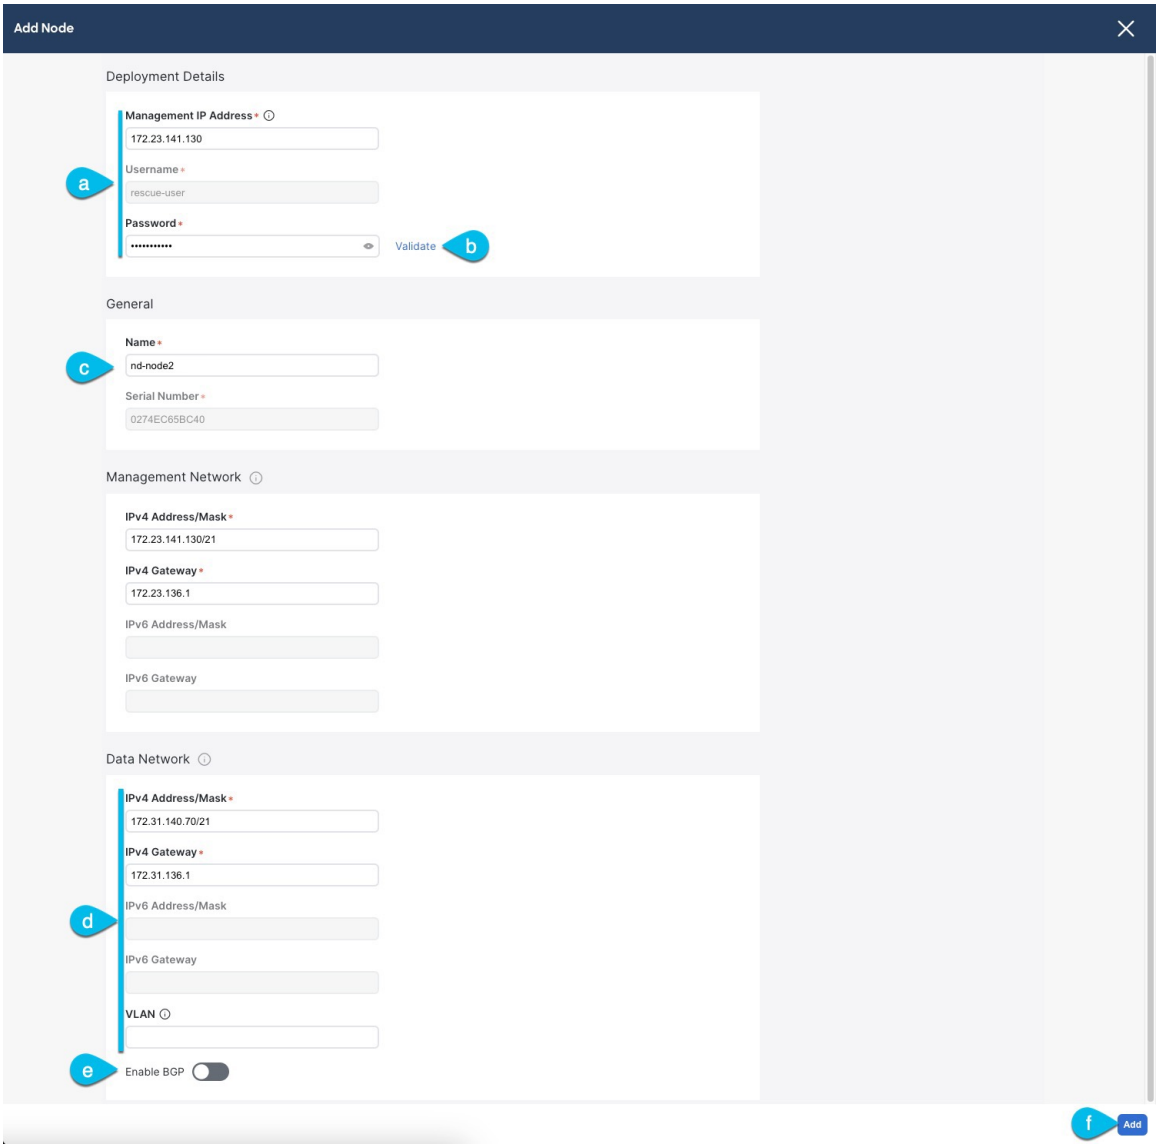
  The router ID must be an IPv4 address, for example `1.1.1.1`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

e) Click **Update** 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.

After network connectivity is validated, you can provide the other required information for the node.

c) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated during the management network information validation in the previous step.

d) In the **Data Network** area, provide the node's **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

  The router ID must be an IPv4 address, for example `2.2.2.2`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

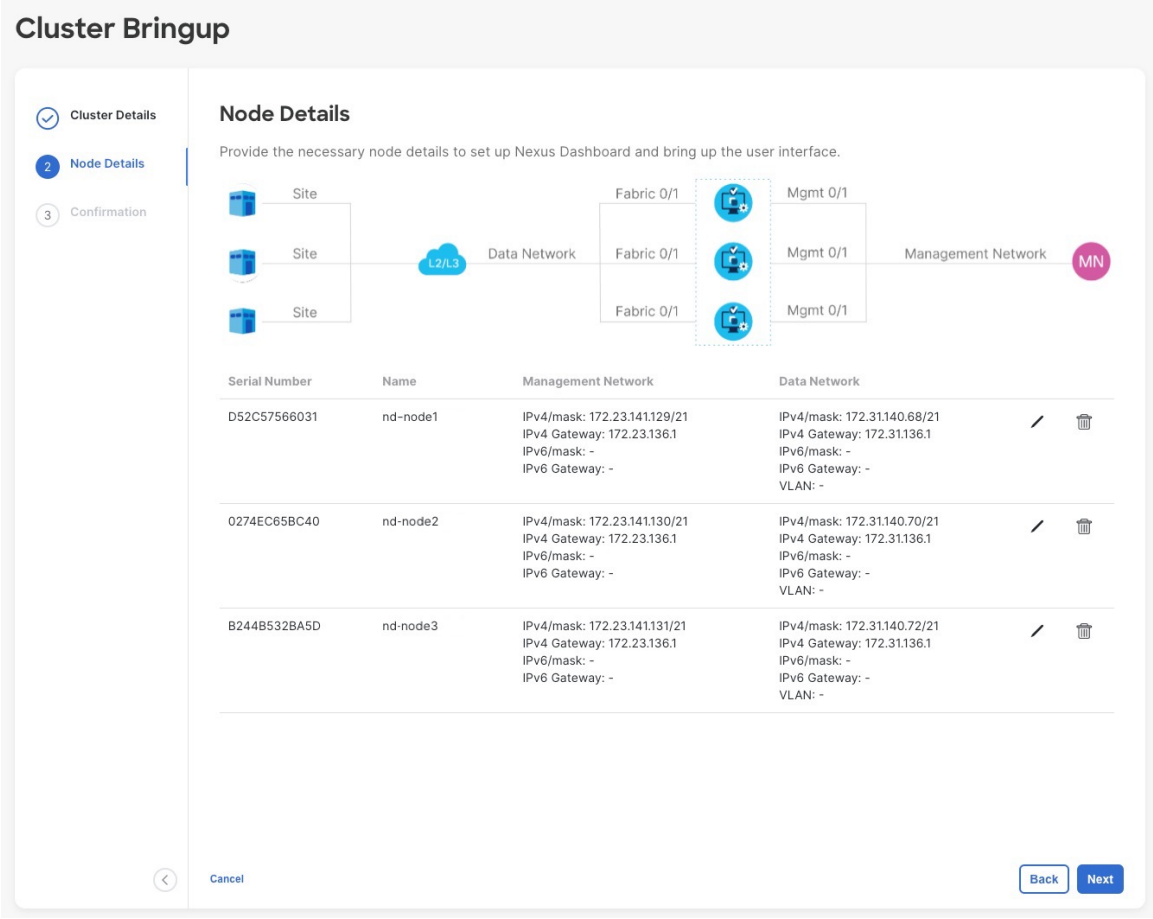
f) Click **Add** to save the changes.

Step 20 Repeat the previous step to add the 3rd node.

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

Step 21 In the **Node Details** page, click **Next** to continue.

After you have provided the management and data network information for all nodes, you can proceed to the final **Confirmation** screen.



Step 22 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 23 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 as the `rescue-user` using the password you provided during node deployment 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 24 Configure the **Network Scale** parameters for your cluster.

This is described in the **Infrastructure Management** > **Cluster Configuration** section of the *Cisco Nexus Dashboard User Guide*, which is also available directly from your Nexus Dashboard's Help Center.

Deploying Nexus Dashboard Directly in VMware ESXi

This section describes how to deploy Cisco Nexus Dashboard cluster directly in VMware ESXi. If you prefer to deploy using vCenter, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 80 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 61.

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/

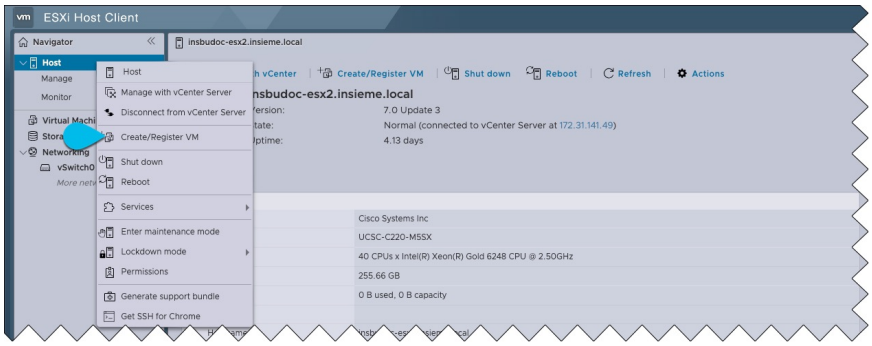
b) Choose the Nexus Dashboard release version you want to download.

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

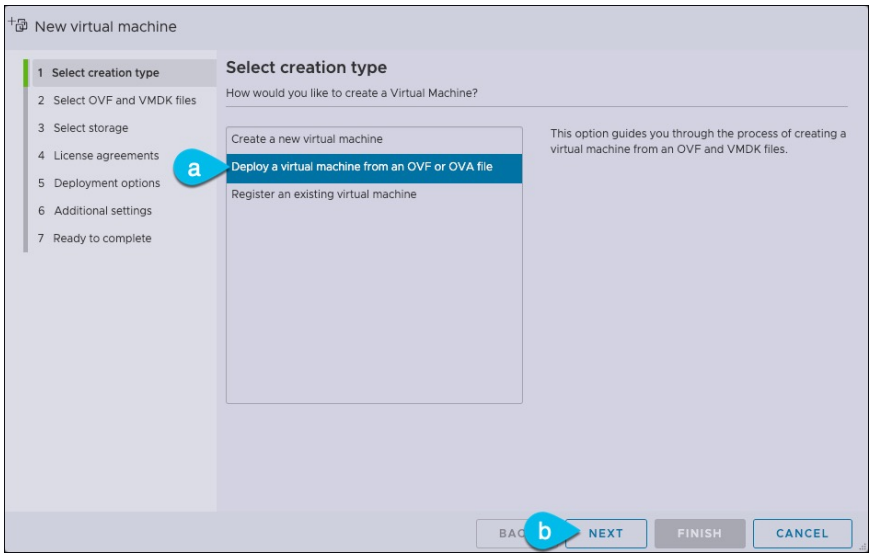
Step 2 Log in to your VMware ESXi.

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 Deployment Overview and Requirements, on page 3.

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

  Node profiles are described in Prerequisites and Guidelines, on page 61.

- 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 step to configure the initial information for the second and third nodes.

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

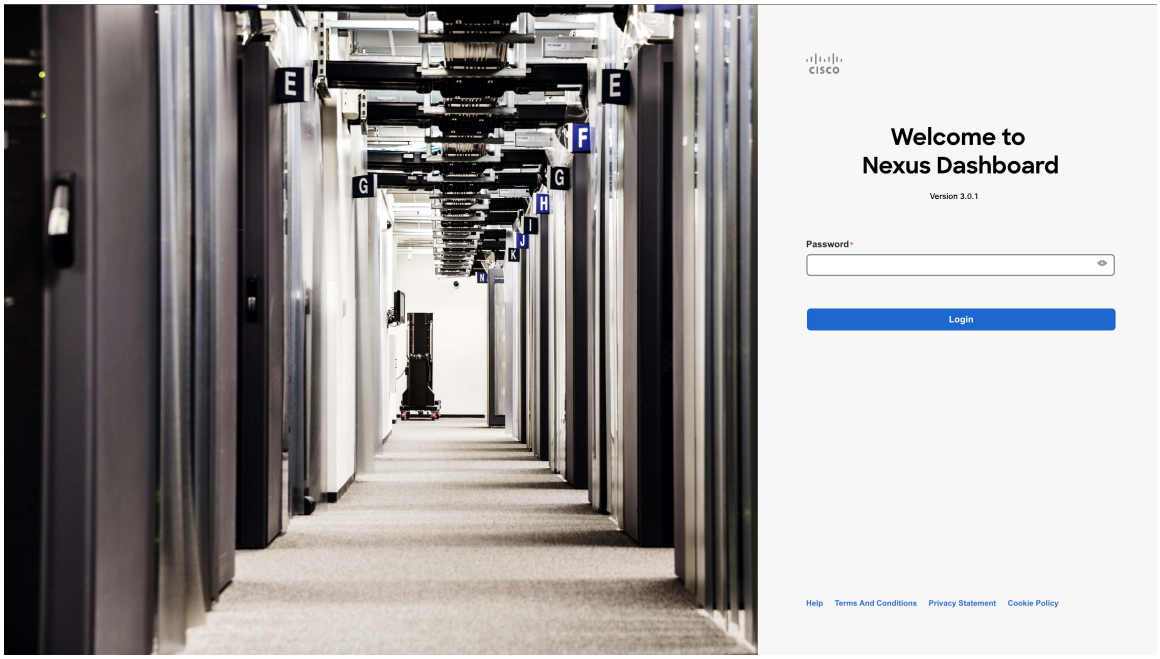
Note
You must provide the same password for all nodes or the cluster creation will fail.

The steps to deploy the second and third nodes are identical with the only exception being that you must indicate that they are not the **Cluster Leader**.

Step 15 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

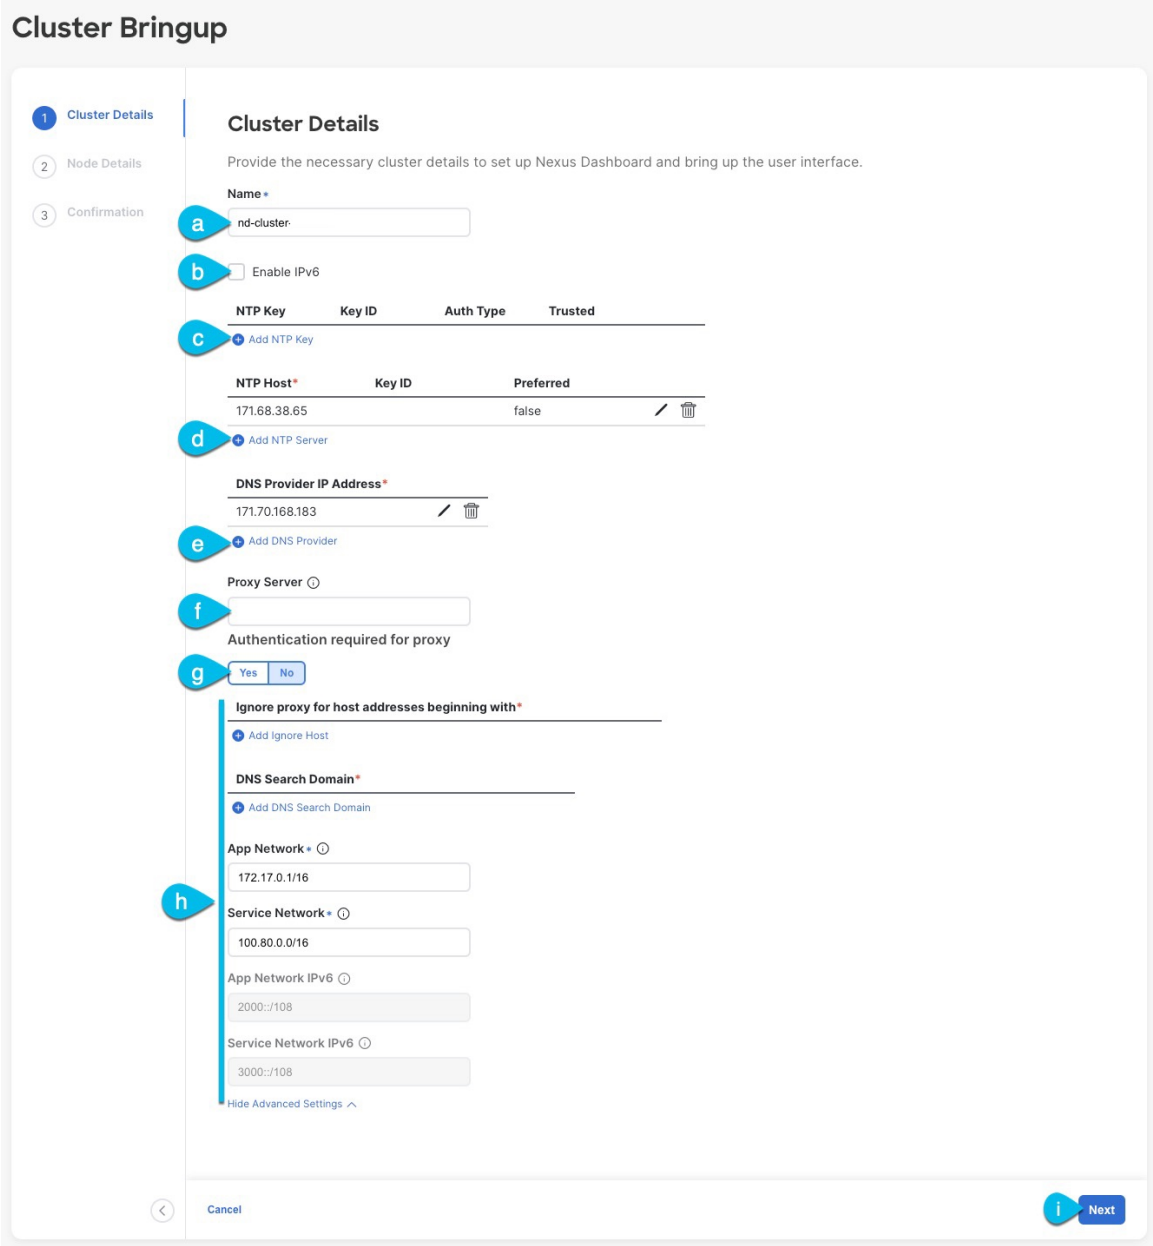
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

Enter the password you provided in a previous step and click **Login**



Step 16 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

b) (Optional) If you want to enable IPv6 functionality for the cluster, check the **Enable IPv6** checkbox.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

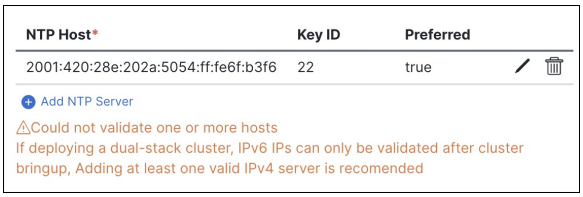
d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

Note
If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and provided an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

The proxy server must have the following URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
```

```
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

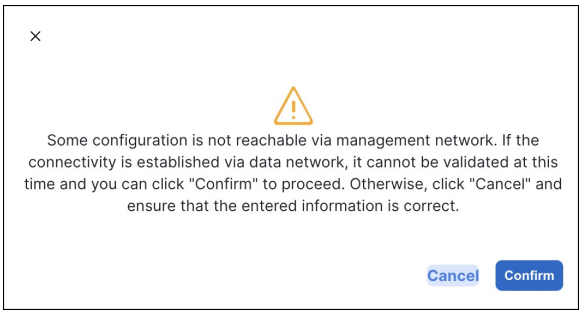
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

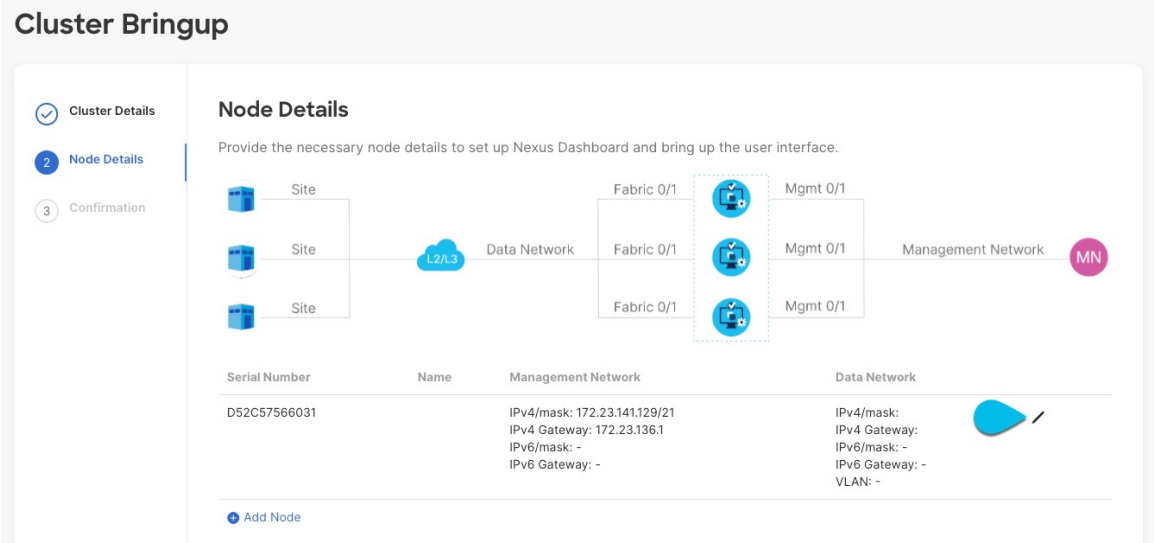
i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



Step 17 In the **Node Details** screen, update the current node's information.

You have defined the Management network and IP address for the node into which you are currently logged in during the initial node configuration in earlier steps, but you must also provide the Data network information for the node before you can proceed with adding the other `primary` nodes and creating the cluster.



a) Click the **Edit** button next to the first node.

b) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated.

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

c) 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, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

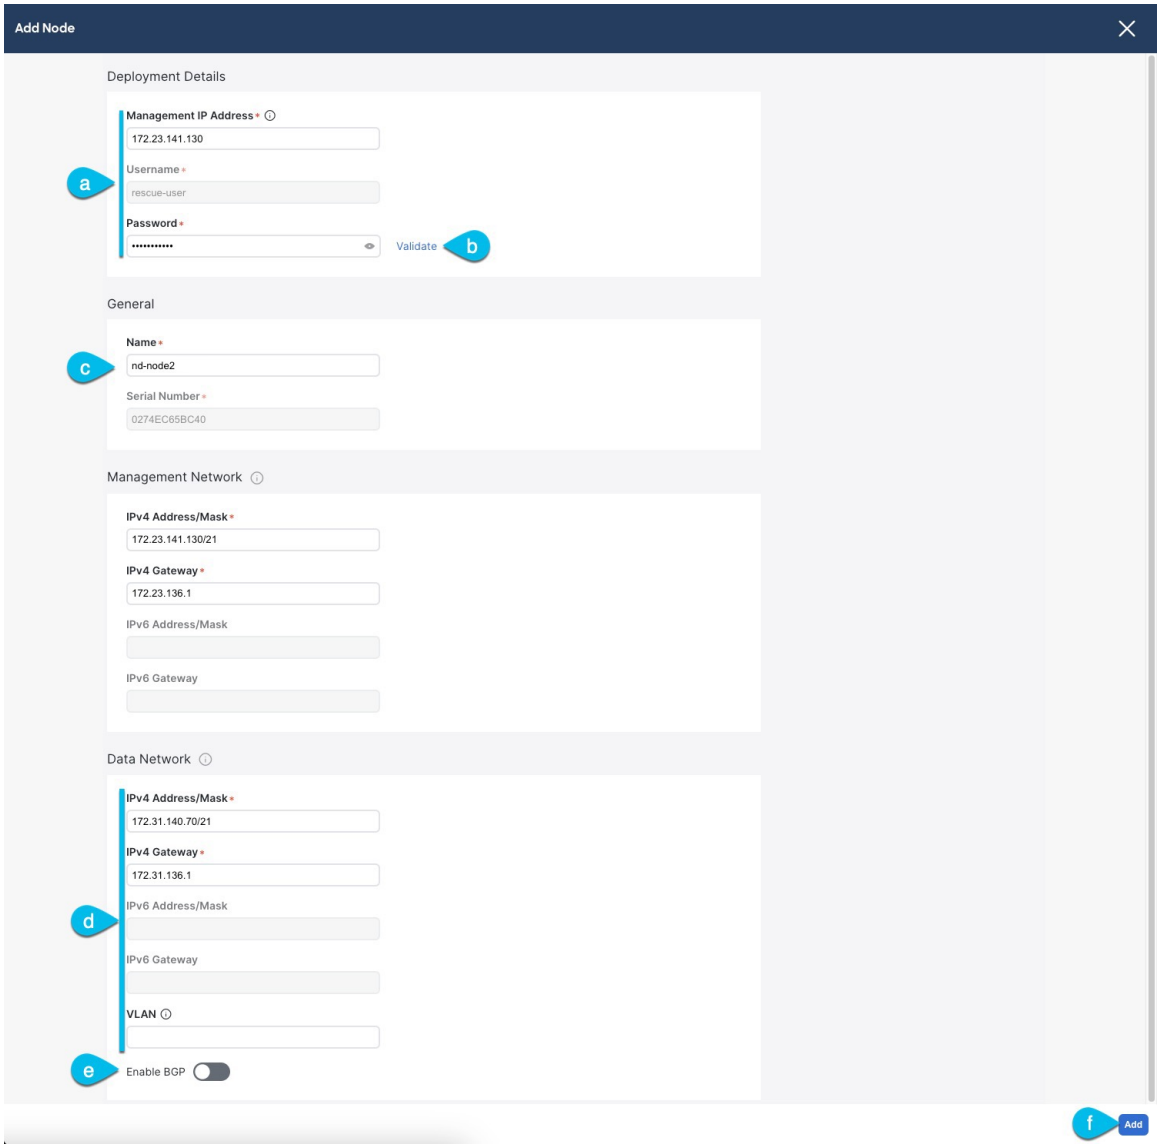
  The router ID must be an IPv4 address, for example `1.1.1.1`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

e) Click **Update** 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.

After network connectivity is validated, you can provide the other required information for the node.

c) Provide the **Name** for the node.

The node's **Serial Number** and the **Management Network** information are automatically populated during the management network information validation in the previous step.

d) In the **Data Network** area, provide the node's **Data Network** information.

You must provide the data network IP address, netmask, and gateway. Optionally, you can also provide the VLAN ID for the network. For most deployments, you can leave the VLAN ID field blank.

If you had enabled IPv6 functionality in a previous screen, 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 IPv4/IPv6 stack.

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- For pure IPv6, the **Router ID** of this node.

  The router ID must be an IPv4 address, for example `2.2.2.2`

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

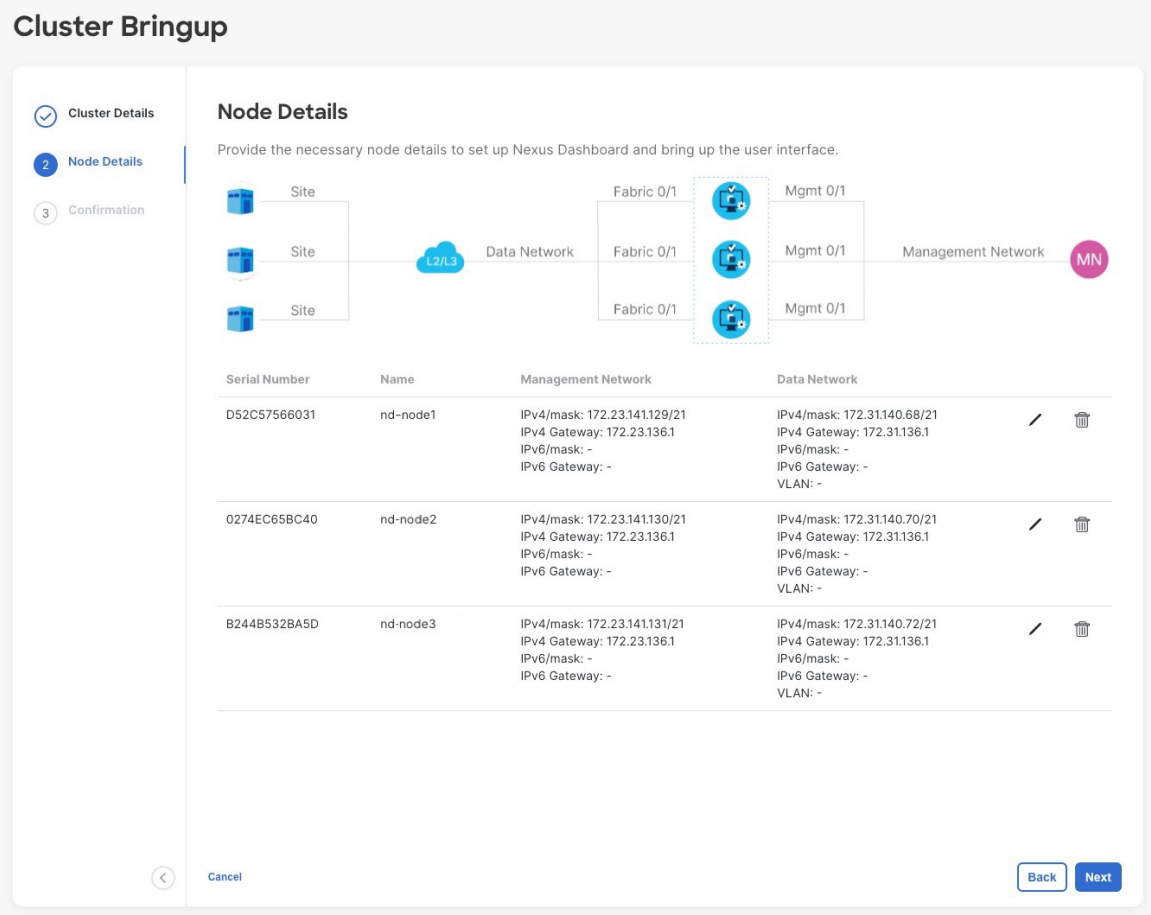
f) Click **Add** to save the changes.

Step 19 Repeat the previous step to add the 3rd node.

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

Step 20 In the **Node Details** page, click **Next** to continue.

After you have provided the management and data network information for all nodes, you can proceed to the final **Confirmation** screen.



Step 21 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 22 **Note**

Step 23 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 as the `rescue-user` using the password you provided during node deployment 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.

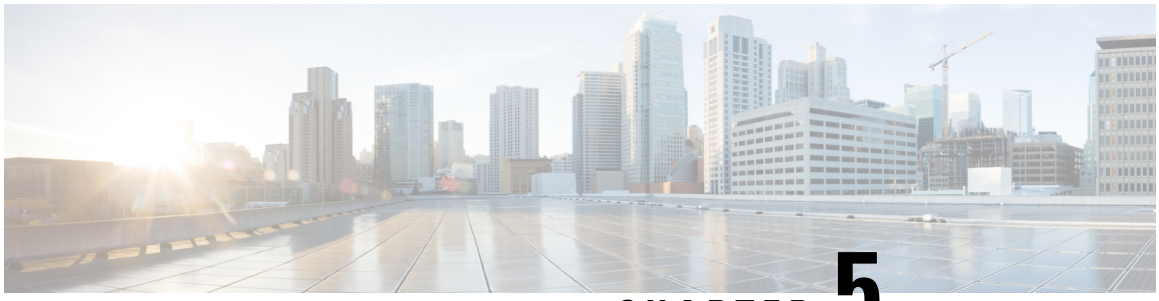
After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node of the Nexus Dashboard cluster.

Note
There may be an issue during the bootstrap process on 3-node vND (ESX) clusters which can cause the 'acs health' command to show the following error: `'k8s: services not in desired state - aaamgr,cisco-intersightdc,eventmonitoring,infra-kafka,kafka,mongodb,sm,statscollect'`

Contact Cisco TAC and open a case referring to open Bug ID CSCwf65557 requesting `root` access to run the workaround command on each node.

Step 24 Configure the **Network Scale** parameters for your cluster.

This is described in the **Infrastructure Management** > **Cluster Configuration** section of the *Cisco Nexus Dashboard User Guide*, which is also available directly from your Nexus Dashboard's Help Center.



CHAPTER 5

Deploying in Amazon Web Services

- Prerequisites and Guidelines, on page 93
- Deploying Nexus Dashboard in AWS, on page 95

Prerequisites and Guidelines

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

- Ensure that the AWS form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.

Note Standby and worker nodes are not supported with this cluster form factor.

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.
- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- 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.

Note By default only PEM-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins by logging in to each node using the generated key and running the required command as described in the setup section below.

Deploying 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 93.

Procedure

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. The CIDR can be a `/24` or `/25` subnet and will be segmented to be used across the availability zones.

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.

f) 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.

g) In the **Password** and **Confirm 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.

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

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

i) 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.

j) Click **Next** to continue.

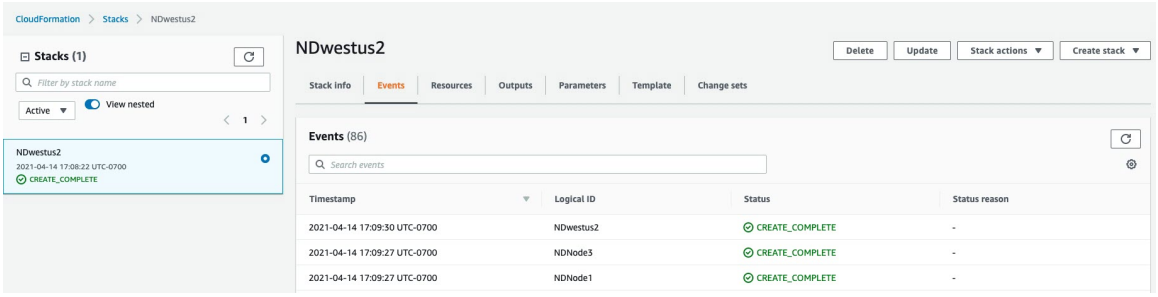
Step 6 In the **Advanced options** screen, simply click **Next**.

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

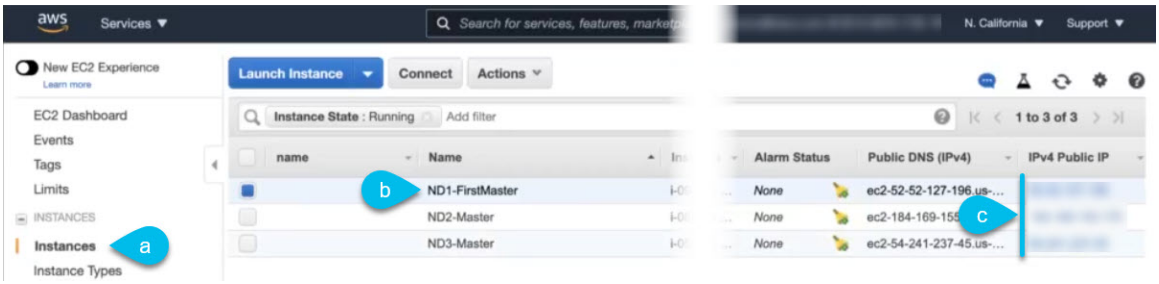
Step 8 Wait for the deployment to complete, then start the VMs.

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 Note down all nodes' public IP addresses.



a) After all instances are deployed, navigate to the AWS console's **EC2** > **Instances** page.

b) Note down which node is labeled as `FirstMaster`.

You will use this node's public IP address to complete cluster configuration.

c) Note down all nodes' public IP addresses.

You will provide this information to the GUI bootstrap wizard in the following steps.

Step 10 Enable password-based login on all nodes.

By default only PEM-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins.

Note
You must enable password-based login on all nodes before proceeding to cluster bootstrap described in the following steps or you will not be able to complete the cluster configuration.

a) SSH into one of the instances using its public IP address and the PEM file.

Use the PEM file you created for this as part of Prerequisites and Guidelines, on page 93.

```
# ssh -i <pem-file-name>.pem rescue-user@<node-public-ip>
```

b) Enable password-based login.

On each node, run the following command:

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

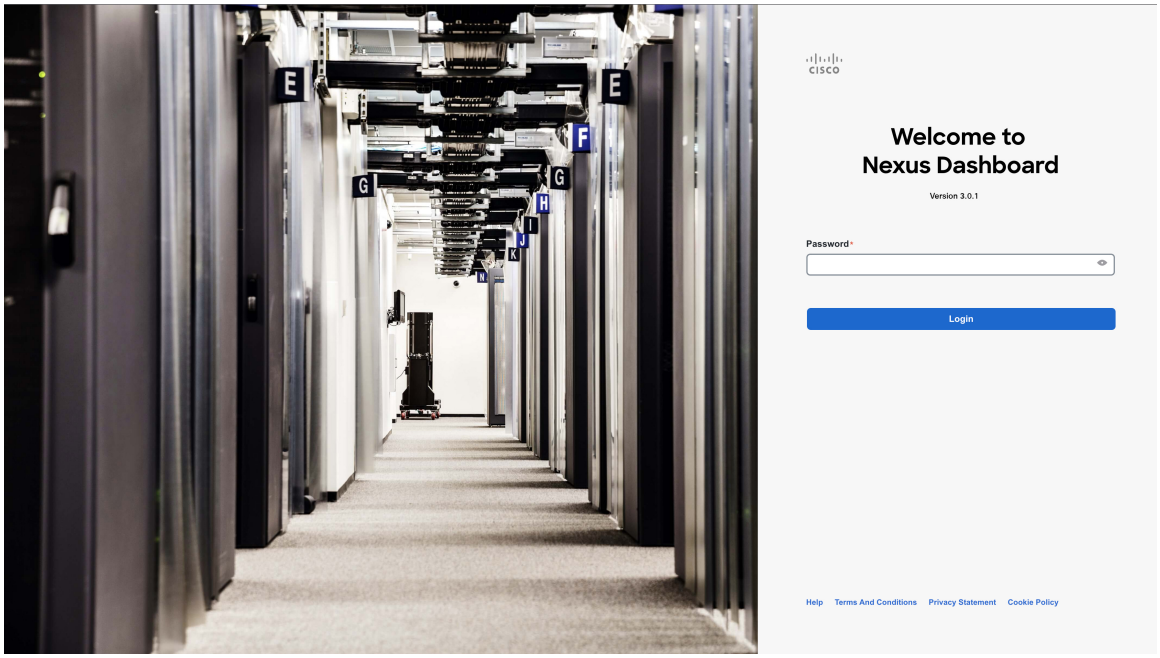
c) Repeat this step for the other two instances.

Step 11 Open your browser and navigate to `https://<first-node-public-ip>` to open the GUI.

Note
You must use the public IP address of the first node (`FirstMaster`) or cluster configuration cannot be completed.

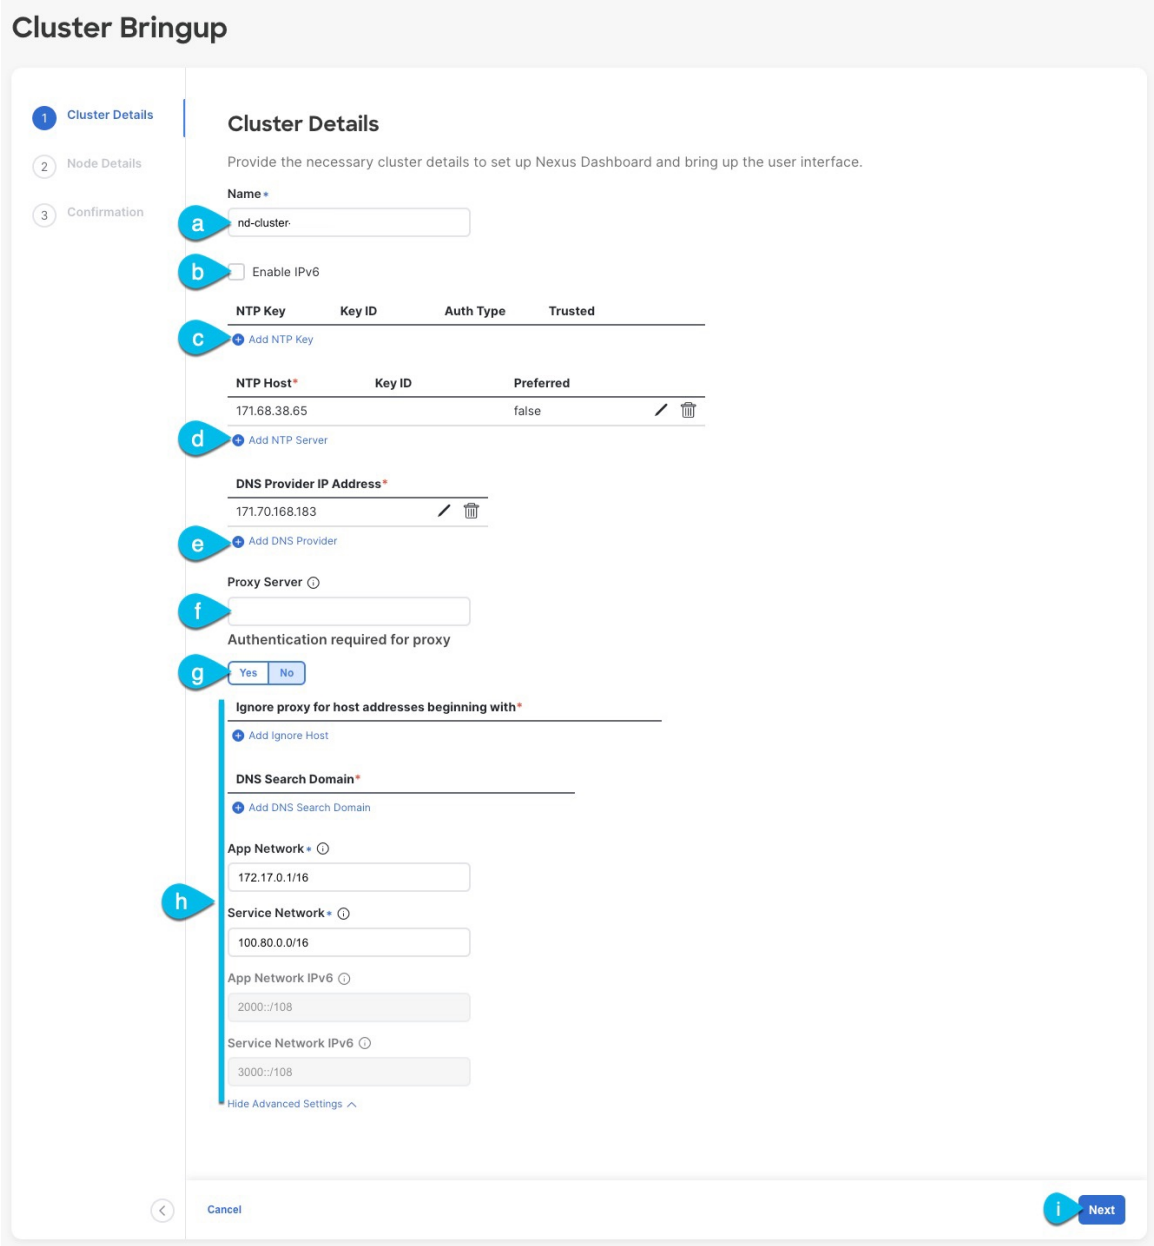
The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided for the first node and click **Login**



Step 12 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

b) (Optional) If you want to enable IPv6 functionality for the cluster, check the **Enable IPv6** checkbox.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

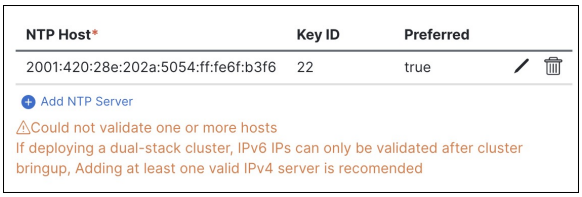
d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

Note
If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and provided an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

The proxy server must have the following URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
```

```
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

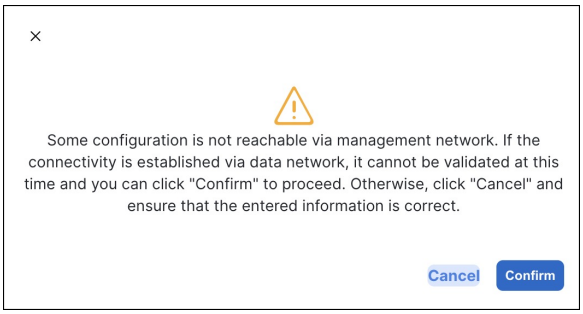
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



Step 13 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 **Name**.

The **Management Network** and **Data Network** information will be already populated from the VPC subnet you have configured before deploying the cluster.

The cluster creates six subnets from the given VPC CIDR, from which the data and management networks will be allocated for the cluster's three nodes.

c) Leave IPv6 addresses and VLAN fields blank.

Cloud Nexus Dashboard clusters do not support these options.

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

Step 14 Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.
b) In the **Credentials** section, provide the node's **Public IP Address** and the password you provided during template deployment, then click **Verify**.

The IP address and password are used to pull that node's **Management Network** and **Data Network** information, which will be populated in the fields below.

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

Step 15 Repeat the previous step to add the 3rd node.

Step 16 In the **Node Details** page, click **Next** to continue.

Step 17 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 18 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 as the `rescue-user` using the password you provided during node deployment 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
```

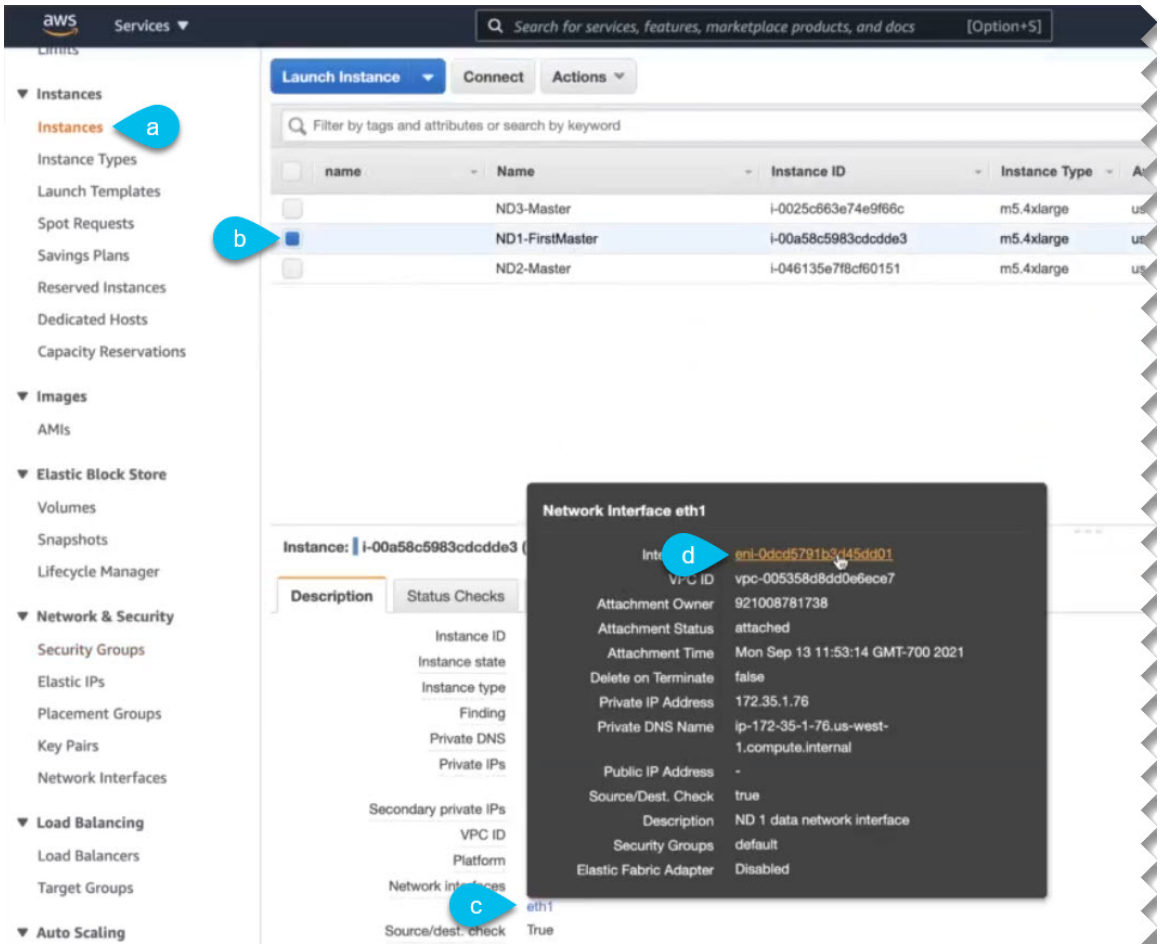
b) Log in to the Nexus Dashboard GUI.

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 19 Update the nodes' security group with required ports.

This step describes how to update the Nexus Dashboard nodes' instances with the required port configuration for on-boarding Cisco NDFC sites. If you do not plan on on-boarding any NDFC sites to your Nexus Dashboard cluster, you can skip this step.

Navigate to one of the nodes' data interface:



a) In the AWS console, navigate to **Instances**.
b) Select one of the Nexus Dashboard instances.

You will make changes to the default security group, so you only need to select one of the nodes.

c) Click the data interface (`eth1`).
d) Click the **Interface ID**.

The **Network Interface** properties page opens.

e) In the **Network Interface** page, click `default` in the **Security groups** column of the interface.

Add the new rules:

a) In the default security group's page, select the **Inbound rules** tab.
b) Click **Edit inbound rules**.
c) In the **Edit inbound rules** page, click **Add rule** to add a new inbound security rule, then provide the details to allow inbound communication on port 443.

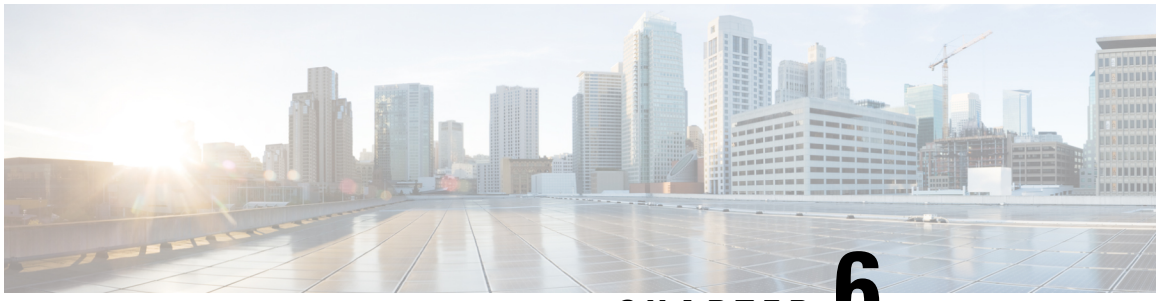
   Provide the following information for the new rule:

   - For **Type**, select `Custom TCP`.
   - For **Port range**, enter `443`.
   - For **Source**, provide the IP addresses of the NDFC controllers which you plan to onboard to your Nexus Dashboard.

d) Still in the **Edit inbound rules** page, click **Add rule** to add another inbound security rule, then provide the details to allow inbound communication on port 9092.

   Provide the following information for the new rule:

   - For **Type**, select `Custom TCP`.
   - For **Port range**, enter `9092`.
   - For **Source**, provide the IP addresses of the NDFC controllers which you plan to onboard to your Nexus Dashboard.



CHAPTER 6

Deploying in Microsoft Azure

- Prerequisites and Guidelines, on page 105
- Deploying Nexus Dashboard in Azure, on page 110

Prerequisites and Guidelines

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

- Ensure that the Azure form factor supports your scale and services requirements.

  Scale and services support and co-hosting vary based on the cluster form factor. You can use the Nexus Dashboard Capacity Planning tool to verify that the cloud form factor satisfies your deployment requirements.

Note Standby and worker nodes are not supported with this cluster form factor.

- Review and complete the general prerequisites described in the Deployment Overview, on page 3.
- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- Have appropriate access privileges for your Azure account and subscription.
- Have created a resource group for your Nexus Dashboard cluster resources.

Note The resource group must be empty and not contain any existing objects. Resource groups with existing objects cannot be used for Nexus Dashboard deployment.

To create a resource group:

- In the Azure portal, navigate to **All Resources** > **Resource Groups**.
- Click **+Add** to create a new resource group.

- In the **Create a resource group** screen, provide the name of the subscription you will use for your Nexus Dashboard cluster, the name for the resource group (for example, `nd-cluster`), and the region.

- Create an SSH key pair.

  A key pair consists of a private key and a public key, you will be asked to provide the public key when creating the Nexus Dashboard nodes.



Note You will need to use the same machine where you create the public key for a one-time login into each node to enable general SSH login during cluster deployment procedure.

Creating SSH keys is described in Generating an SSH Key Pair in Linux or MacOS, on page 106 and Generating an SSH Key Pair in Windows, on page 107 sections below.

Generating an SSH Key Pair in Linux or MacOS

These procedures describe how to generate an SSH public and private key pair in Linux or MacOS. For instructions on generate an SSH public and private key pair in Windows, see Generating an SSH Key Pair in Windows, on page 107.

Procedure

Step 1 On your Linux virtual machine or Mac, create a public and private key pair using ssh-keygen, directing the output to a file.

```
# ssh-keygen -f filename
```

For example:

```
# ssh-keygen -f azure_key
```

Output similar to the following appears. Press the Enter key without entering any text when you are asked to enter a passphrase (leave the field empty so that there is no passphrase).

```
Generating public/private rsa key pair.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in azure_key.
Your public key has been saved in azure_key.pub.
The key fingerprint is:
SHA256:gTsQIIAadjgNsgcguifIloh4XGpVWMdcXVV6U0dyBNs
...
```

Step 2 Locate the public and private key files that you saved.

```
# ls
```

Two files should be displayed, where:

- The file with the .pub suffix contains the public key information

- The file with the same name, but with no suffix, contains the private key information

For example, if you directed the output to a file named `azure_key`, you should see the following output:

```
# ls
azure_key
azure_key.pub
```

In this case:

- The `azure_key.pub` file contains the public key information
- The `azure_key` file contains the private key information

Step 3 Open the public key file and copy the public key information from that file, without the `username@hostname` information at the end.

Note
The private key file is not used in the installation process. However, you might need it for other reasons, such as logging into your Nexus Dashboard nodes through SSH.

Generating an SSH Key Pair in Windows

These procedures describe how to generate an SSH public and private key pair in Windows. For instructions on generate an SSH public and private key pair in Linux, see Generating an SSH Key Pair in Linux or MacOS, on page 106.

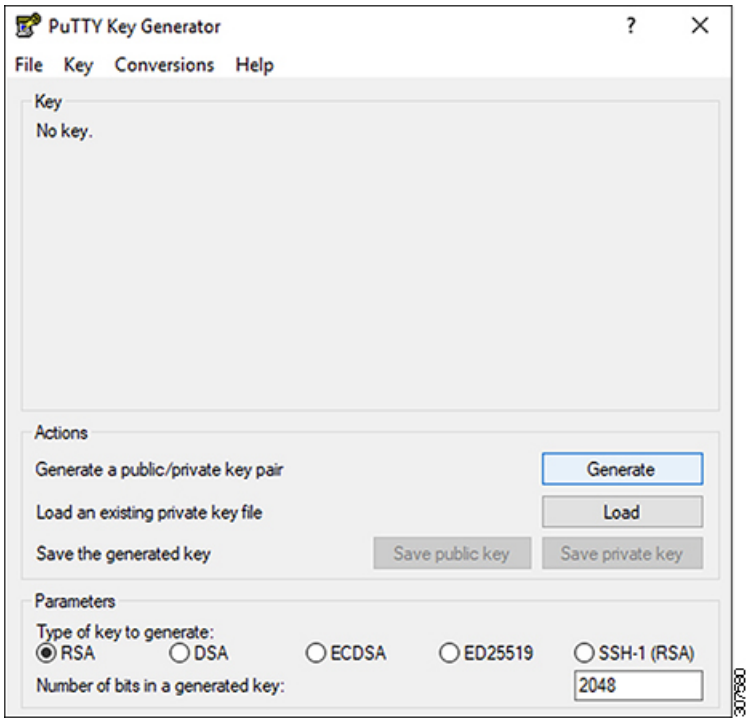
Procedure

Step 1 Download and install the PuTTY Key Generator (puttygen):

https://www.puttygen.com/download-putty

Step 2 Run the PuTTY Key Generator by navigating to **Windows** > **Start Menu** > **All Programs** > **PuTTY** > **PuTTYgen**.

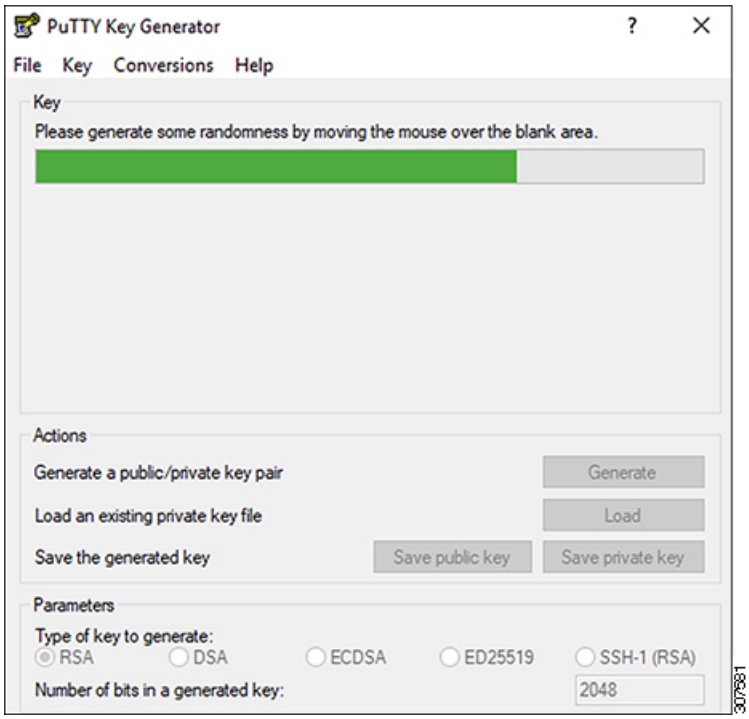
You will see a window for the PuTTY Key Generator on your screen.



Step 3 Click **Generate**.

A screen appears, asking you to move the mouse over the blank area to generate a public key.

Step 4 Move your cursor around the blank area to generate random characters for a public key.



Step 5 Save the public key.

a) Navigate to a folder on your laptop where you want to save the public key file and create a text file for this public key.

b) Copy the information in the PuTTY Key Generator.

Copy the public key information in the window, with these inclusions and exclusions:

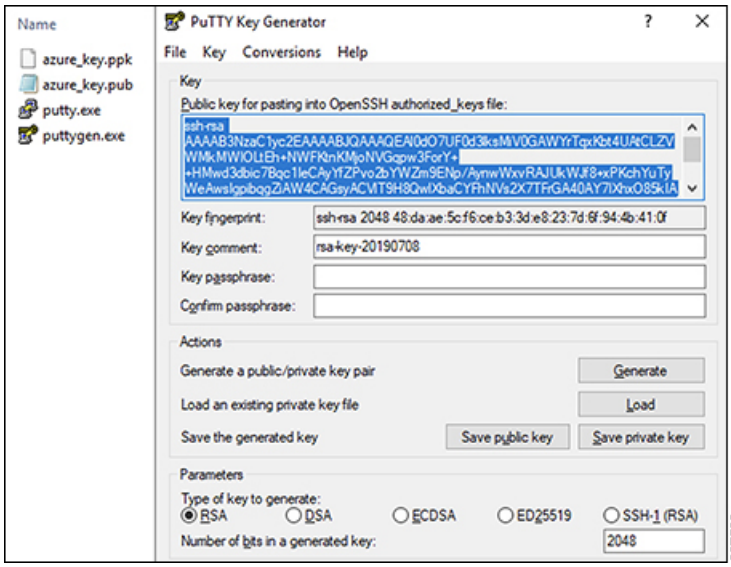
- Including the **ssh-rsa** text at the beginning of the public key.
- Excluding the following text string at the end:

```
== rsa-key-<date-stamp>
```

Truncate the key so that it does not include the == **rsa-key-<date-stamp>** text string at the end.

Note
In the next set of procedures, you will paste the public key information into the Azure ARM template. If the form does not accept the key in this format, add == back to the end of the key, as this format is required in some regions.

If the key is not in the correct format, the Nexus Dashboard will not complete its installation.



c) Paste the information in the public key text file that you created in 5.a, on page 109 and save the file, giving it a unique file name.

This public key text file will now contain a key that is on a single line of text. You will need the information in this public key text file in the next set of procedures.

Note
Do not save the public key using the **Save public key** option in the PuTTY Key Generator. Doing so saves the key in a format that has multiple lines of text, which is not compatible with the Nexus Dashboard deployment process.

Step 6 Save the private key.

a) Click **Save private key**.

A screen appears, asking if you want to save the file without a passphrase. Click **Yes** on this screen.

b) Navigate to a folder on your laptop and save the private key file, giving it a unique file name.

Note
The private key file is not used in the installation process. However, you might need it for other reasons, such as logging into your Nexus Dashboard nodes through SSH.

Deploying 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 105.

Procedure

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

a) Log into your Azure account and browse to https://marketplace.microsoft.com/en-us/.
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 Provide **Basic** information.

a) From the **Subscription** dropdown, select the subscription you want to use for this.
b) From the **Resource group** dropdown, select the resource group you created for this as part of Prerequisites and Guidelines, on page 105.
c) From the **Region** dropdown, select the region where the template will be deployed.
d) In the **Password** and **Confirm Password** fields, provide the admin password for the nodes.

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.

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

e) In the **SSH public key** field, paste the public key from the key pair you generated as part of the Prerequisites and Guidelines, on page 105 section.
f) Click **Next** to proceed to the next screen.

Step 3 Provide **ND Settings** information.

a) Provide the **Cluster Name**.
b) In the **Image Version** dropdown, verify that the correct version is selected.
c) In the **Virtual Network Name** field, provide the name for a VNET that will be created for your cluster.

The VNET must not already exist and will be created for you during deployment. If you provide an already existing VNET, the deployment cannot proceed.

d) In the **Subnet Address Prefix** field, provide a subnet within the VNET.

The subnet must be a `/24` subnet and it must be different from the default VNET subnet you defined when creating the VNET.

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) Click **Next** to proceed to the next screen.

Step 4 In the **Review + create** page, review information and click **Create** to deploy the cluster.

Step 5 Wait for the deployment to complete, then start the VMs.

Step 6 Note down all nodes' public IP addresses.

After all instances are deployed, navigate to the Azure console, select each VM, and note down all nodes' public IP addresses. You will provide this information to the GUI bootstrap wizard in the following steps.

Also note which is the "first" node, which will be indicated by the node's VM name `vm-node1-<cluster-name>`. You will use this node's public IP address to complete cluster configuration.

Step 7 Enable password-based login on all nodes.

By default only key-based login is enabled for each node. To be able to SSH into the nodes using a password, as required by the GUI setup wizard, you will need to explicitly enable password-based logins.

Note
You must enable password-based login on all nodes before proceeding to cluster bootstrap described in the following steps or you will not be able to complete the cluster configuration.

a) SSH in to one of the nodes as `rescue-user`.

Note
You must use the same machine as you used to create the public key for the deployment during the Prerequisites and Guidelines, on page 105 section.

You can log in as `rescue-user` using the password you provided in template's **Basic** settings:

```
# ssh rescue-user@<node-public-ip>
```

b) Enable password-based login.

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

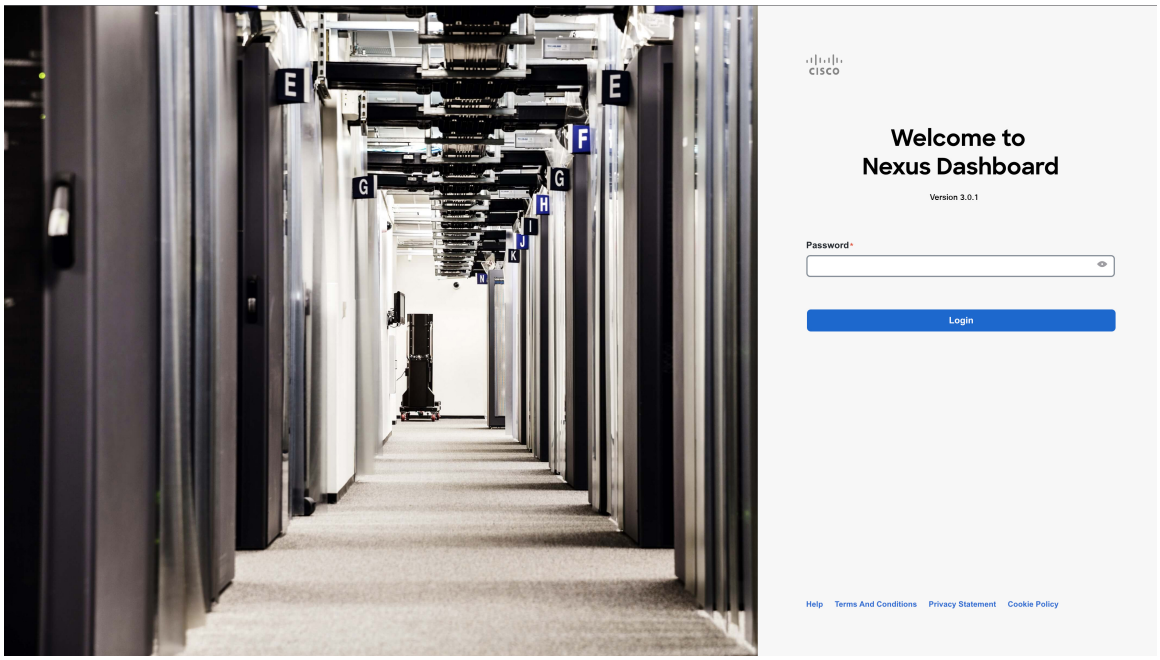
c) Repeat this step for the other two nodes.

Step 8 Open your browser and navigate to `https://<first-node-public-ip>` to open the GUI.

Note
You must use the public IP address of the first node (`vm-node1-<cluster-name>`) or cluster configuration cannot be completed.

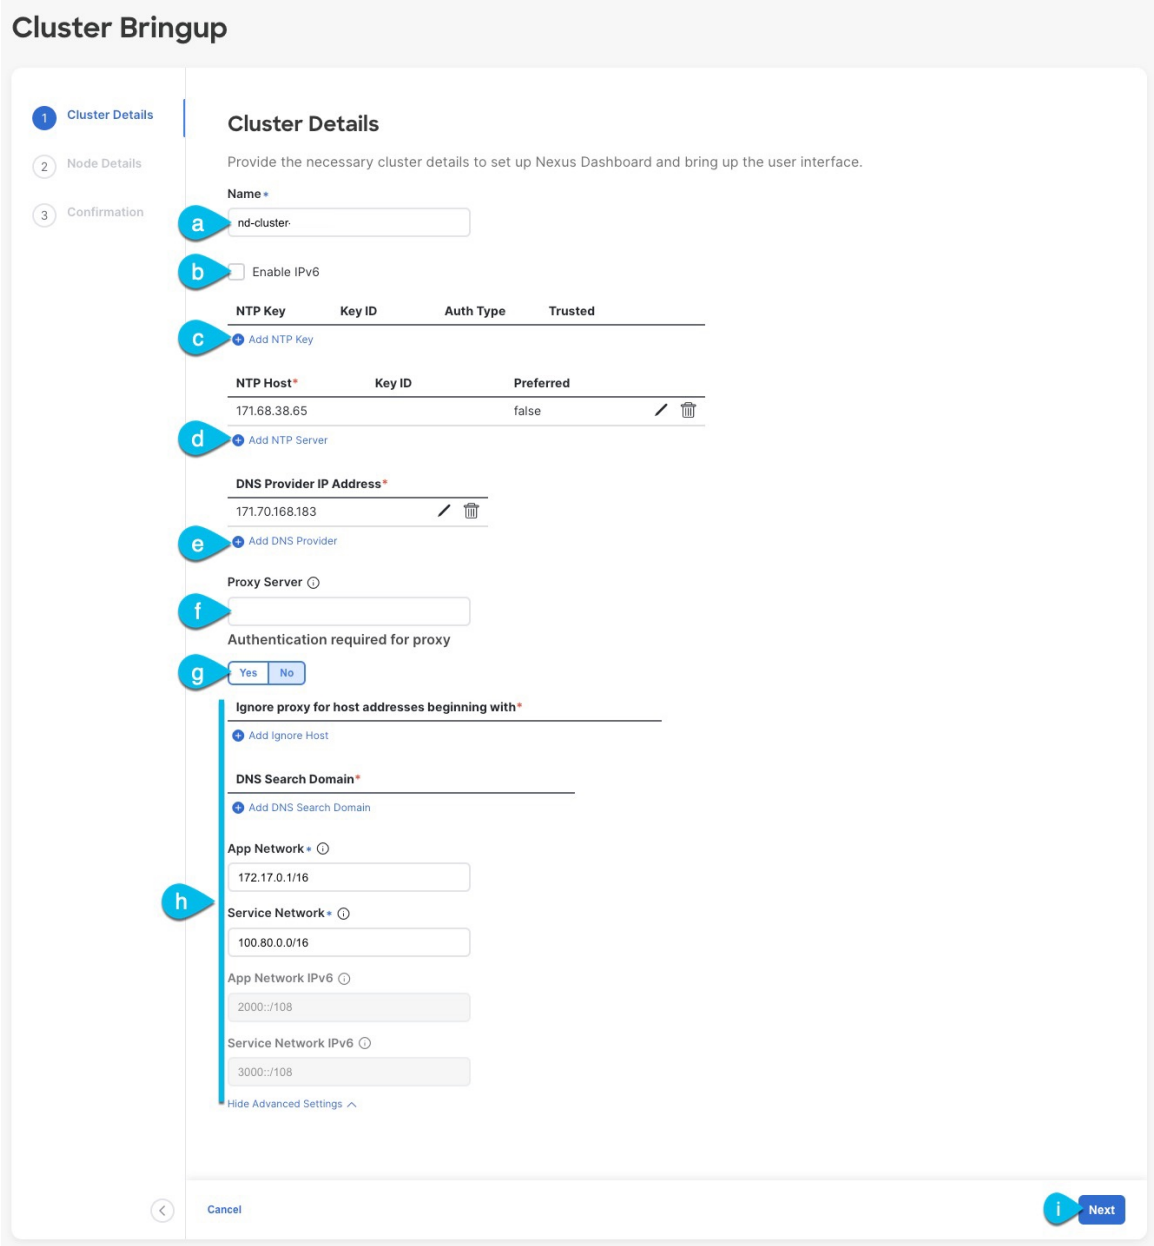
The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

Enter the password you provided for the first node and click **Login**



Step 9 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

b) (Optional) If you want to enable IPv6 functionality for the cluster, check the **Enable IPv6** checkbox.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

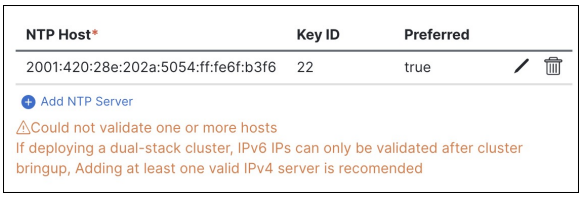
d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

Note
If the node into which you are logged in is configured with only an IPv4 address, but you have checked **Enable IPv6** in a previous step and provided an IPv6 address for an NTP server, you will get the following validation error:



This is because the node does not have an IPv6 address yet (you will provide it in the next step) and is unable to connect to an IPv6 address of the NTP server.

In this case, simply finish providing the other required information as described in the following steps and click **Next** to proceed to the next screen where you will provide IPv6 addresses for the nodes.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

The proxy server must have the following URLs enabled:

```
dcappcenter.cisco.com
svc.intersight.com
```

```
svc.ucs-connect.com
svc-static1.intersight.com
svc-static1.ucs-connect.com
```

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

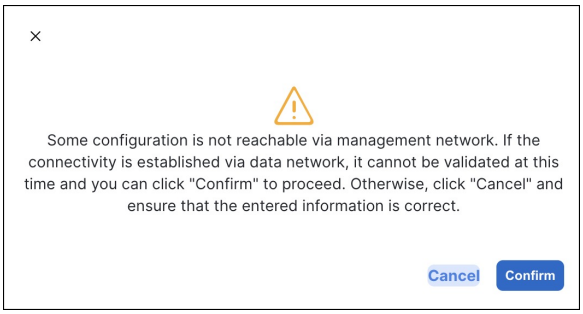
  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

i) Click **Next** to continue.

Note
If your node has only an IPv4 management address but you have checked **Enabled IPv6** and provided an IPv6 NTP server address, ensure that the NTP address is correct and click **Confirm** to proceed to the next screen where you will provide the nodes' IPv6 addresses.



Step 10 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 **Name**.

The **Management Network** and **Data Network** information will be already populated from the VNET subnet you have configured before deploying the cluster.

The cluster creates six subnets from the given VNET, from which the data and management networks will be allocated for the cluster's three nodes.

c) Leave IPv6 addresses and VLAN fields blank.

Cloud Nexus Dashboard clusters do not support these options.

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

Step 11 Click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) Provide the node's **Name**.
b) In the **Credentials** section, provide the node's **Public IP Address** and the password you provided during template deployment, then click **Verify**.

The IP address and password are used to pull that node's **Management Network** and **Data Network** information, which will be populated in the fields below.

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

Step 12 Repeat the previous step to add the 3rd node.

Step 13 In the **Node Details** page, click **Next** to continue.

Step 14 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 15 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 as the `rescue-user` using the password you provided during node deployment 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:

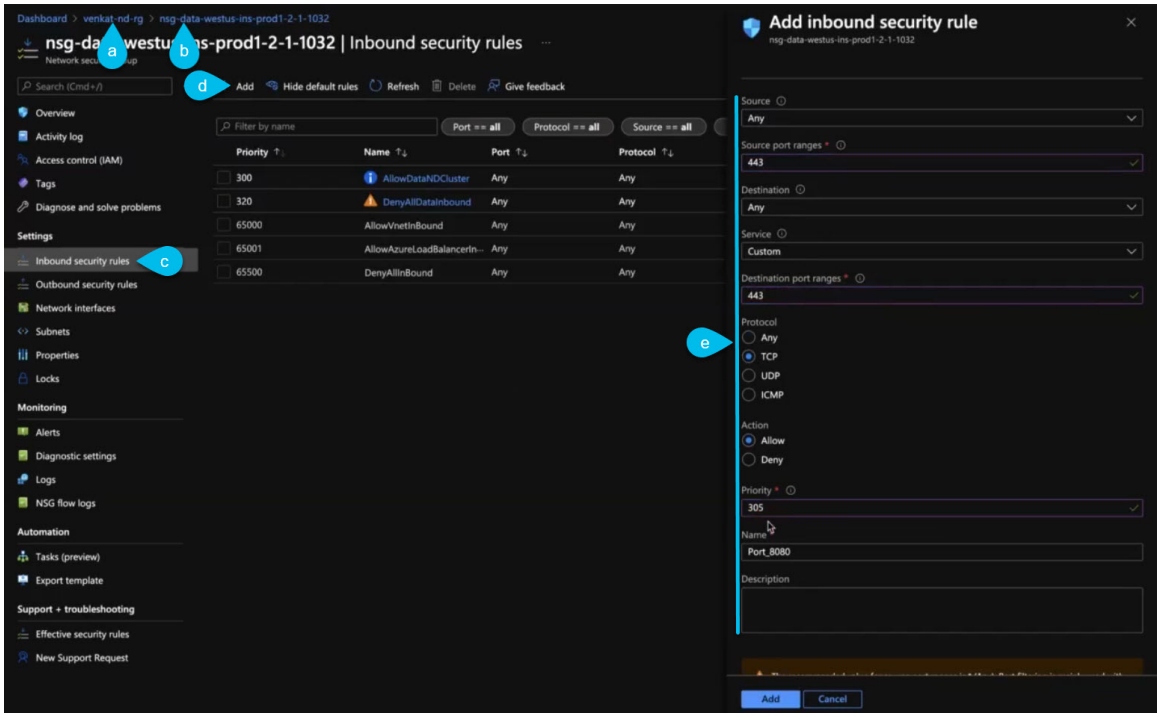
```
$ acs health
All components are healthy
```

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 16 Update the nodes' security group with required ports.

This step describes how to update the Nexus Dashboard nodes' instances with the required port configuration for on-boarding Cisco NDFC sites. If you do not plan on on-boarding any NDFC sites to your Nexus Dashboard cluster, you can skip this step.



a) In the Azure portal, navigate to the resource group where you deployed your Nexus Dashboard.

This is the same resource group you selected in Step 2.

b) Select the security group attached to the nodes' data interfaces.

The name of the security group will begin with `nsg-data-<region>-...`.

c) In the security group's setting navigation bar, select **Inbound security rules**.

d) Click **+Add** to add a new inbound security rule, then provide the details to allow inbound communication on port 443.

Provide the following information for the new rule:

- For **Source**, select `Any`.
- For **Source port ranges**, enter `443`.
- For **Destination**, select `Any`.
- For **Destination port ranges**, enter `443`.
- For **Protocol**, choose `TCP`.

- For **Action**, choose `Allow`.
- For **Priority**, choose a priority between `300` and `320`.

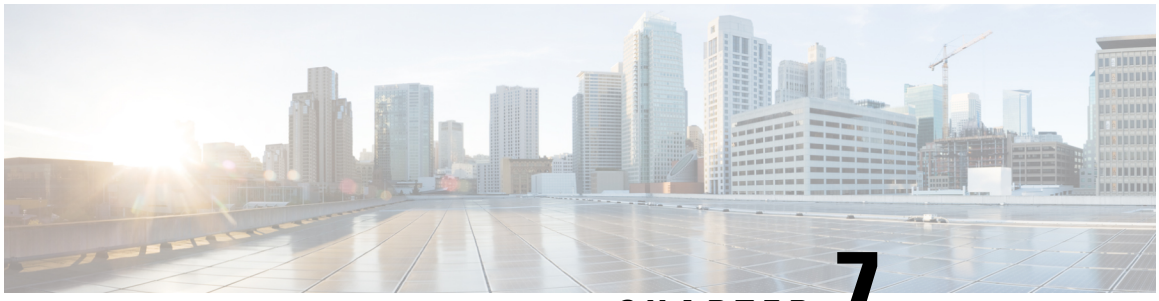
  For example, `305`.
- Provide a **Name** for the rule.

e) Click **+Add** to add a new inbound security rule, then provide the details to allow inbound communication on port 9092.

Repeat the previous substep to add another rule with the following details:

- For **Source**, select `Any`.
- For **Source port ranges**, enter `9092`.
- For **Destination**, select `Any`.
- For **Destination port ranges**, enter `9092`.
- For **Protocol**, choose `TCP`.
- For **Action**, choose `Allow`.
- For **Priority**, choose a priority between `300` and `320`.

  For example, `310`.
- Provide a **Name** for the rule.



CHAPTER 7

Deploying in Existing Red Hat Enterprise Linux Installation

- Prerequisites and Guidelines, on page 119
- Deploying Nexus Dashboard in Existing Red Hat Enterprise Linux Installation, on page 121
- Uninstalling Nexus Dashboard Software, on page 130
- Troubleshooting Nexus Dashboard Deployments in RHEL, on page 130

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.

  The guide is available from the Nexus Dashboard UI or online at *Cisco Nexus Dashboard User Guide*
- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- Ensure that your server is running Red Hat Enterprise Linux (RHEL) release 8.4 or 8.6.

  This release of Nexus Dashboard support both physical and virtual deployments of RHEL.

Note For non-production deployments, such as labs or testing, you can deploy in other Linux distributions by passing an additional parameter to the installer in **Step 4** of the Deploying Nexus Dashboard in Existing Red Hat Enterprise Linux Installation, on page 121 section:

```
./nd-installer setup ./examples/nd-linux-input.yaml
skip-os-version-check
```

- Ensure that your server has at least 20GB of available disk space in the `/var` partition.

  The installation process requires the additional disk space to extract temporary installation files.
- You can deploy only a single-node or three-node (all `master` nodes) cluster in RHEL.

  Adding `worker` or `standby` node is not supported for this cluster form factor.

- Clusters deployed in RHEL support only Nexus Dashboard Fabric Controller (NDFC), Release 12.1(1) or later service with **SAN Controller** deployment type.

  You must deploy one of the other form factors if you want to run any other Nexus Dashboard services or another deployment type of NDFC. For more information about services supported one each Nexus Dashboard cluster form factor, see Cisco Nexus Dashboard Cluster Sizing and Nexus Dashboard and Services Compatibility Matrix.



Note Standby and worker nodes are not supported with this cluster form factor.

- Ensure that the following system-level requirements are satisfied:
  - An existing Linux user on each cluster node, which you will provide to the installer and which can be used to manage and troubleshoot the node.

    Only one system user can connect to the Nexus Dashboard node's system. For more information, see Troubleshooting Nexus Dashboard Deployments in RHEL, on page 130 after the deployment.

  - System clocks across all nodes must be synchronized.

    You can use a system utility such as `chrony` to ensure proper time synchronization between the nodes.



Note By default, the Nexus Dashboard installer for RHEL verifies that the system clock is synchronized using `chrony`. If you use a different system to synchronize the clock, you can use `./nd-installer setup input.yaml` **`skip-ntp-check`** during installation to bypass the default verification.

  - Skopeo package is installed.

    Skopeo is outside the scope of this document, but in short you can use `yum install skopeo` command to install the package.

  - Swap file is disabled.

    You can disable swap by removing its entry from the `/etc/fstab` file and restarting the server.

  - The `firewalld` and `libvirt` services must be installed but stopped and disabled prior to deploying Nexus Dashboard software.

Note The following additional system-level changes will be applied when you deploy Nexus Dashboard software to allow executables from additional directories and the cluster's own SSH server:

```
/usr/bin/chcon -R -t bin_t /mnt/atom
/usr/bin/chcon -R -t bin_t /mnt/linux
/usr/bin/chcon -R -t bin_t /opt/apic-sn

/usr/bin/chcon -t ssh_home_t -R /data/services/issh/ssh_host_rsa_key
/usr/bin/chcon -t ssh_home_t -R /data/services/issh/intssh
/usr/sbin/semanage port -a -t ssh_port_t -p tcp 1022
```

- Ensure you have enough system resources.

  When deploying in RHEL, you can deploy two types of nodes:

Table 17: Deployment Requirements

Default Node Profile	Large Node Profile
• 16 vCPUs • 64 GB of RAM • 500GB SSD storage for the data volume and an additional 100GB for the system volume. All nodes must be deployed on SSD or faster storage. • Two network interfaces in addition to the RHEL's management interface.	• 32 vCPUs • 128 GB of RAM • 3TB SSD storage for the data volume and an additional 100GB for the system volume. The data volume may be a combination of multiple drives (such as in RAID configuration) as long as the drive is presented as a single device to the operating system. All nodes must be deployed on SSD or faster storage. • Two network interfaces in addition to the RHEL's management interface.

Deploying Nexus Dashboard in Existing Red Hat Enterprise Linux Installation

This section describes how to configure and bring up a Nexus Dashboard cluster in RHEL.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 119.

Procedure

Step 1 Obtain the Cisco Nexus Dashboard software archive package (tarball).

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 release version you want to download.

d) Click the **Download** icon next to the Nexus Dashboard tarball (`nd-rhel-<version>.tar`).

Step 2 Extract the downloaded archive.

```
tar -xvf nd-rhel-<version>.tar
```

Step 3 Modify the installation `yaml` file.

The distribution tarball includes a sample YAML file (`./nd-linux/examples/nd-linux-input.yaml`), which you can modify to provide the values appropriate to your deployment.

For example, the following sample node configuration YAML file highlights the specific fields which you must provide:

- For `blkdev`, provide the SSD device(s) for the node's system and data volumes.

  You must provide at least one SSD device for ND installation. The total size of this device must meet the system and data volume requirements listed in Prerequisites and Guidelines, on page 119. You can choose to provide two or more devices to satisfy system and data volumes' size requirements. The order of the devices in the YAML file does not matter – the smaller disk will be used for the system volume and the rest of device(s) combined will be used for the data volume.

  Note
  All provided devices will be erased and consumed for the Nexus Dashboard node

  For more information on the node device requirements, see Prerequisites and Guidelines, on page 119.

- For `oobNetwork`, provide the management network information:
  - For `uplinks`, provide the names of the network interfaces connected to the Nexus Dashboard management network.

    These interfaces must be dedicated exclusively to the Nexus Dashboard.
  - For `ipNet`, provide the node's management network IPv4 address and netmask in the `172.23.152.214/24` format.
  - For `gatewayIP`, provide the node's management network IPv4 gateway.
- For `inbandNetwork`, provide the names of the network interfaces connected to the Nexus Dashboard data network.

  You must provide only the interface(s) in the `uplinks` section (and not the network or gateway information) as the rest of the configuration is defined during the GUI bootstrap process.
- For `firstMaster`, ensure that only one of the node is set to `true` and the other 2 nodes are set to `false`.

  You will use the `firstMaster` node to complete the cluster bootstrap process using the GUI.
- For `clusterName`, provide the name of the cluster.

- For `installProfile`, choose either `Default` or `Large`.

  For more information on the node profile requirements, see Prerequisites and Guidelines, on page 119.

- For `serviceUser`, provide an existing Linux account name which will be used for managing and troubleshooting the Nexus Dashboard node.

  Note
  The `serviceUser` must be different from the system's `root` user, which you must have created before starting the installation as mentioned in Prerequisites and Guidelines, on page 119.

```
# Node role. This release supports 'Master' nodes only.
nodeRole: Master

# Block devices. One or more un-partitioned devices that meet profile requirements for data and
system volumes.
blkdev:
  - type: SSD
    name: "/dev/sdb"
  - type: SSD
    name: "/dev/sdc"

# Networking. You must provide 2 interfaces exclusively to ND that are separate from the Linux
management interface.
oobNetwork:
  uplinks:
  - ens924
  ipNet: 172.23.152.214/24
  gatewayIP: 172.23.152.1

# Data network interface only, the rest of the data network configuration is provided during UI
bootstrap
inbandNetwork:
  uplinks:
  - ens956

# Set to 'true' for one of the nodes in the cluster.
firstMaster: true

# Cluster name.
clusterName: nd-cluster

# Installation profile: 'Default' or 'Large'.
installProfile: Default

# Linux user account name. This must not be the 'root' user.
# Only this user will have privileges to execute certain ND diagnostics commands.
serviceUser: nduser
```

Step 4 Install Nexus Dashboard node software.

```
cd nd-linux
./nd-installer setup ./examples/nd-linux-input.yaml
```

You will be asked to provide the password, which will be used for the Nexus Dashboard cluster `admin` account.

Note
By default, the installer verifies that the system clock is synchronized using `chrony`. If you use a different system to synchronize the clock, you can use `./nd-installer setup ./examples/nd-linux-input.yaml` **`skip-ntp-check`** to bypass the default verification.

Step 5 Repeat the previous steps to deploy the 2nd and 3rd nodes.

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

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

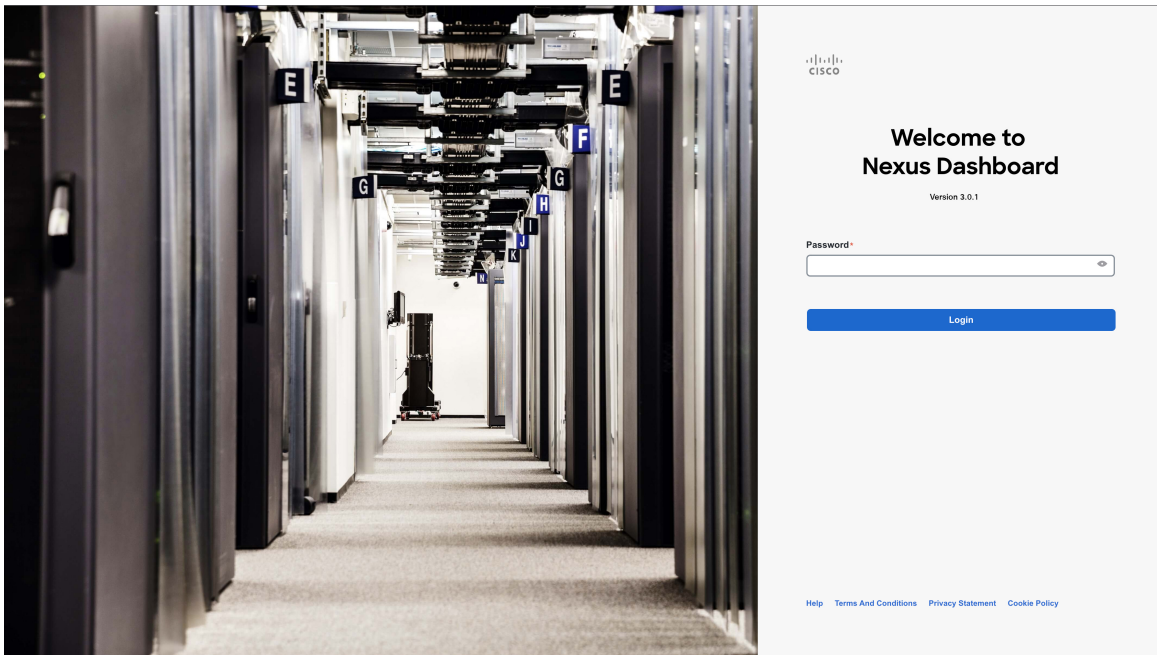
Note
When providing the node details in the configuration YAML file for the 2nd and 3rd nodes, ensure that `firstMaster` parameter is set to `false`.

Step 6 Wait for all three nodes to finish deploying.

Step 7 Open your browser and navigate to `https://<first-node-mgmt-ip>` to open the GUI.

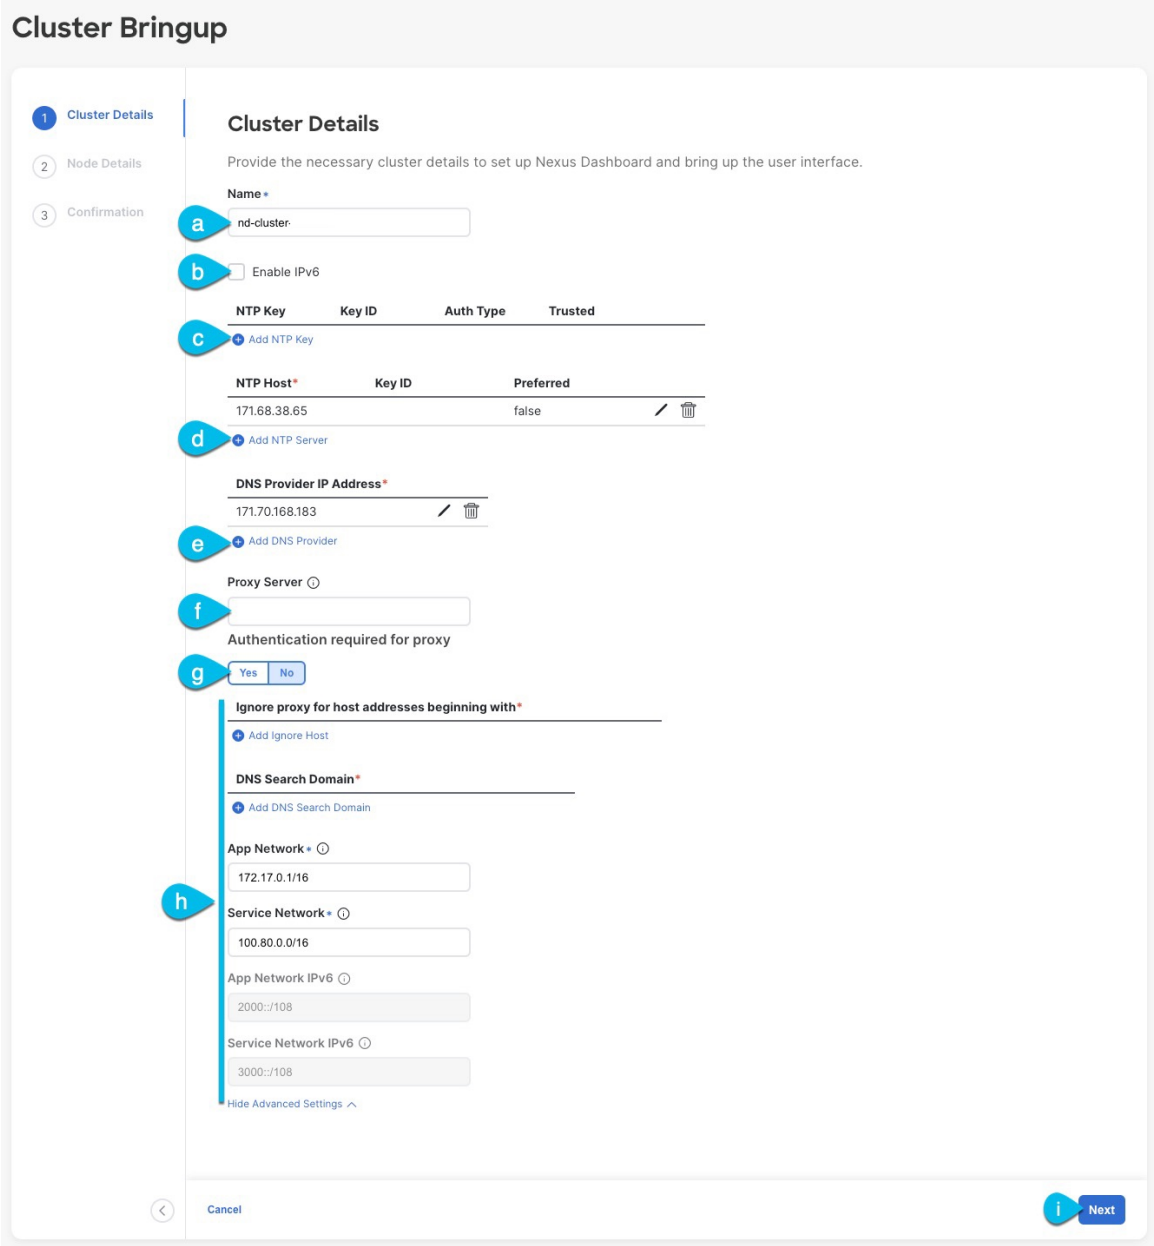
The rest of the configuration workflow takes place from one of the node's GUI. You must use the IP address you provided for the node which you designated as `firstMaster`.

Enter the password you provided in a previous step and click **Begin Setup**



Step 8 Provide the **Cluster Details**.

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



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

b) Do not **Enable IPv6**.

Note
This form factor does not support IPv6 functionality.

c) (Optional) If you want to enable NTP server authentication, click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.

- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

Note
For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 6.

After you've entered the information, click the checkmark icon to save it.

d) Click **+Add NTP Host** to add one or more NTP servers.

In the additional fields, provide the following information:

- **NTP Host** – you must provide an IP address; fully qualified domain name (FQDN) are not supported.
- **Key ID** – if you want to enable NTP authentication for this server, provide the key ID of the NTP key you defined in the previous step.
- Choose whether this NTP server is **Preferred**.

After you've entered the information, click the checkmark icon to save it.

If you want to provide additional NTP servers, click **+Add NTP Host** again and repeat this substep.

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

After you've entered the information, click the checkmark icon to save it.

f) Provide a **Proxy Server**.

For clusters that do not have direct connectivity to Cisco cloud, we recommend configuring a proxy server to establish the connectivity. This allows you to mitigate risk from exposure to non-conformant hardware and software in your fabrics.

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

g) (Optional) If your proxy server required authentication, change **Authentication required for Proxy** to `Yes` and provide the login credentials.

h) (Optional) Expand the **Advanced Settings** category and change the settings if required.

Under advanced settings, you can configure the following:

- Provide one or more search domains by clicking **+Add DNS Search Domain**.

  After you've entered the information, click the checkmark icon to save it.

- Provide custom **App Network** and **Service Network**.

  The application overlay network defines the address space used by the application's services running in the Nexus Dashboard. The field is pre-populated with the default `172.17.0.1/16` value.

  The services network is an internal network used by the Nexus Dashboard and its processes. The field is pre-populated with the default `100.80.0.0/16` value.

  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

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

i) Click **Next** to continue.

Step 9 In the **Node Details** screen, provide the node's information.

a) Click the **Edit** button next to the first node.

b) In the **Password** field, enter the password for this node and click **Validate**.

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

The hostname of the RHEL server where the node software is installed is used for the node's **Name**.

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

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

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

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

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

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

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

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

BGP configuration is required for the Persistent IPs feature required by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in detail in the "Persistent IP Addresses" sections of the *Nexus Dashboard User's Guide*.

Note
You can enable BGP at this time or in the Nexus Dashboard GUI after the cluster is deployed.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

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

Step 10 In the **Node Details** screen, click **Add Node** to add the second node to the cluster.

The **Node Details** window opens.

a) In the **Deployment Details** section, provide the node's **Management IP Address** and **Password**, then click **Verify**.

This is the password you provided to the `./nd-installer setup` command during installation in Step 4.

Verifying the IP and password will auto-populate the **Name**, **Serial Number**, and **Management Network** information for the node.

The hostname of the RHEL server where the node software is installed is used for the node's **Name**.

b) Provide the node's **Name**.

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

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

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

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

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

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

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

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Nexus Dashboard Insights with NDFC fabrics. This feature is described in more detail in Prerequisites and Guidelines, on page 6 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.

When you enable BGP, you must also provide the following information:

- **ASN** (BGP Autonomous System Number) of this node.

  You can configure the same ASN for all nodes or a different ASN per node.

- **BGP Peer Details**, which includes the peer's IPv4 or IPv6 address and peer's ASN.

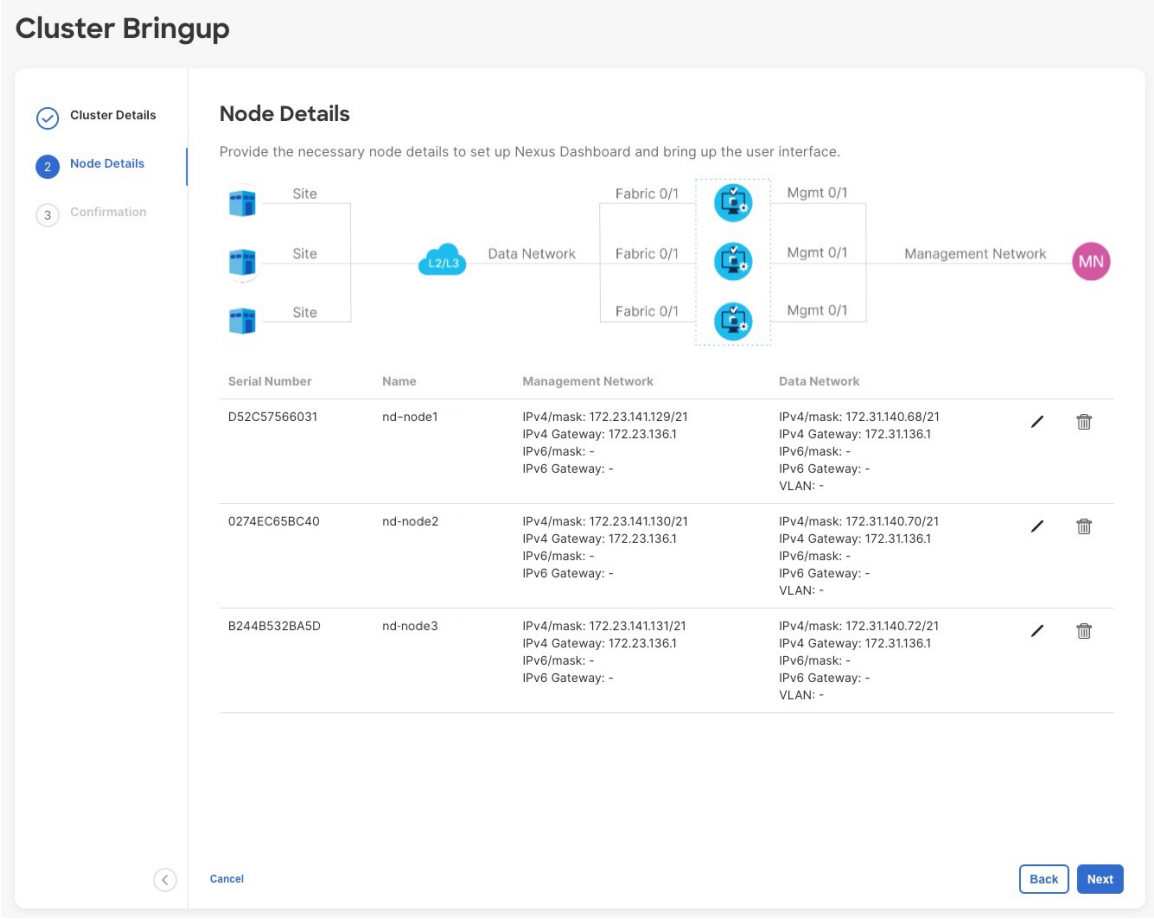
f) Click **Add** to save the changes.

Step 11 Repeat the previous step to add the 3rd node.

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

Step 12 In the **Node Details** screen, click **Next** to continue.

After you have provided the management and data network information for all nodes, you can proceed to the final **Confirmation** screen.



Step 13 In the **Confirmation** screen, review and verify the configuration information and click **Configure** to create the cluster.

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

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

Step 14 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 using the `serviceUser` you provided in the configuration YAML and run the following command to verify cluster health:

a) After logging in to the Linux system, connect to the node using the `/usr/bin/attach-nd` command.

This command can be used only by the `serviceUser` user.

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 password you provided to the `./nd-installer setup` command during installation in Step 4.

Uninstalling Nexus Dashboard Software

When the Nexus Dashboard node software is deployed, the uninstaller is copied into the `/usr/bin` directory.

If at any point you want to uninstall the software, simply run the following command as the `root` user:

```
/usr/bin/nd-installer uninstall
```

Note If you log in to the RHEL system using SSH, you must connect to the system's management IP address to uninstall; you must not use the Nexus Dashboard's management IP addresses.

This will remove the software and undo the file system changes done during the installation process.

Troubleshooting Nexus Dashboard Deployments in RHEL

This section describes common troubleshooting steps for Nexus Dashboard software deployed in RHEL.

Procedure

Step 1 Check installation logs.

Nexus Dashboard installation logs are available in the following directory:

```
/logs/ndlinux/
```

Step 2 Access the Nexus Dashboard environment after installation.

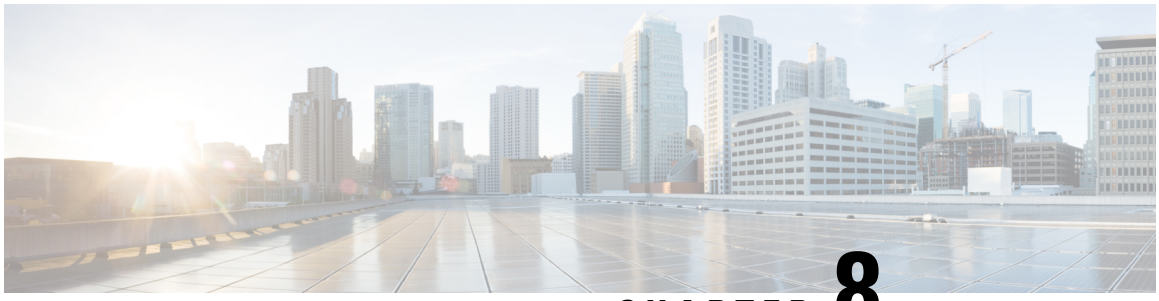
a) Log in to your RHEL system using the Nexus Dashboard user you provided in the YAML configuration file during installation.

b) Access the Nexus Dashboard environment.

```
/usr/bin/attach-nd
```

c) Use the common Nexus Dashboard troubleshooting commands.

After you access the Nexus Dashboard environment, you can use all the common Nexus Dashboard commands described in the "Troubleshooting" section of the *Cisco Nexus Dashboard User Guide*.



CHAPTER 8

Upgrading Nexus Dashboard

- Prerequisites and Guidelines, on page 133
- Upgrading Nexus Dashboard, on page 135

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.

  The upgrade process is the same for all Nexus Dashboard form factors. However, if your existing cluster is deployed using physical servers, VMware ESX, Linux KVM, Azure, or AWS, you will use the target release's ISO image (`nd-dk9.<version>.iso`) to upgrade; if your existing cluster is deployed in Red Hat Enterprise Linux, you will use the RHEL-specific image (`nd-rhel-<version>.tar`).

- Ensure that you have read the *Release Notes* and upgrade guides for any services you run in the existing cluster and plan to run on the target release for service-specific changes in behavior, guidelines, and issues that may affect your upgrade.

  You can find the service-specific documents at the following links:

  - Nexus Dashboard Fabric Controller Release Notes
  - Nexus Dashboard Fabric Controller Upgrade Guide
  - Nexus Dashboard Insights Release Notes
  - Nexus Dashboard Insights Upgrade Guide
  - Nexus Dashboard Orchestrator Release Notes
  - Nexus Dashboard Orchestrator Upgrade Guide

- If you are upgrading a physical Nexus Dashboard cluster, ensure that the nodes have the supported CIMC version for the target Nexus Dashboard release.

  Supported CIMC versions are listed in the *Nexus Dashboard Release Notes* for the target release.

  CIMC upgrade is described in detail in the "Troubleshooting" sections of the *Nexus Dashboard User Guide*.

- If you are upgrading a virtual Nexus Dashboard cluster deployed in VMware ESX, ensure that the ESX version is still supported by the target release.

  Supported ESX versions are listed in the Prerequisites and Guidelines, on page 61 section for VMware deployments.

Note If you need to upgrade the ESX server, you must do that before upgrading your Nexus Dashboard to the target release:

1. Upgrade one of the ESX hosts as you typically would with your existing Nexus Dashboard node VM running.
2. After the host is upgraded, ensure that the Nexus Dashboard cluster is still operational and healthy.
3. Repeat the upgrade on the other ESX hosts one at a time.
4. After all ESX hosts are upgraded and the existing Nexus Dashboard cluster is healthy, proceed with upgrading your Nexus Dashboard to the target release as described in the following section.

- You must perform configuration backups of your Nexus Dashboard and services before the upgrade to safeguard data and minimize any potential risk before proceeding with the upgrade.
- You must be running Nexus Dashboard release 2.2(1) or later to upgrade directly to release 3.0(1).

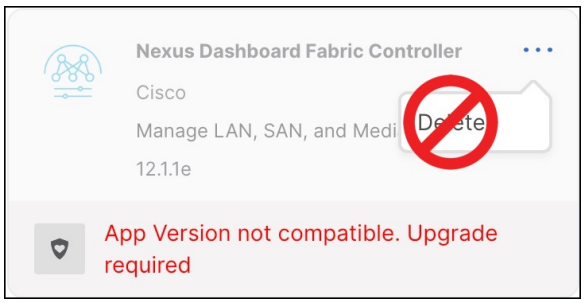
Note Additional upgrade path and minimum version requirements may depend on the specific services you have installed in your cluster. Check the service-specific *Release Notes* and *Upgrade Guides* linked above for more information.

- You must disable all services running in the cluster before upgrading to release 3.0(1) or later.

Note After you have upgraded your Nexus Dashboard to release 3.0(1) or later, you must also upgrade all existing services to the versions compatible with the new Nexus Dashboard version.

When upgrading the services, you must not re-enable the existing versions that you had disabled prior to the Nexus Dashboard upgrade.

You must also not delete the existing versions of the service, which may now display an `App Version not compatible. Upgrade required.` error:



- You must have valid DNS and NTP servers configured and reachable by all cluster nodes.
- Ensure that your current Nexus Dashboard cluster is healthy.

  You can check the system status on the **Overview** page of the Nexus Dashboard's **Admin Console** or by logging in to one of the nodes as `rescue-user` and ensuring that the `acs health` command returns `All components are healthy`.
- Ensure that no configuration changes are made to the cluster, such as adding `worker` or `standby` nodes, while the upgrade is in progress.
- Nexus Dashboard does not support platform downgrades.

  If you want to downgrade to an earlier release, you will need to deploy a new cluster and reinstall the services.

Upgrading Nexus Dashboard

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

Note The following steps illustrate the upgrade workflow from Nexus Dashboard release 2.2.2. If you are upgrading from release 2.3(x), the UI may differ slightly but the upgrade workflow and functionality remains the same.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 133

Procedure

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 for your target release.

Note

- If your Nexus Dashboard is deployed in Red Hat Enterprise Linux, you must download the `.tar` image (`nd-rhel-<version>.tar`) to perform the upgrade.

  More information about RHEL deployments is available in Deploying in Existing Red Hat Enterprise Linux Installation, on page 119.

- For all other form factors, you must download the `.iso` image (`nd-dk9.<version>.iso`) to perform the upgrade.

  For example, even if you used the virtual form factors, such as the `.ova` image for deploying in VMware ESX, or a cloud provider's marketplace for initial cluster deployment, you must still use the `.iso` image for upgrades.

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 and navigate to the **Admin Console**.

Step 3 Disable any existing services installed in the cluster.

Note
You must disable all services before you upgrade the cluster.

a) From the main navigation menu, select **Services**.

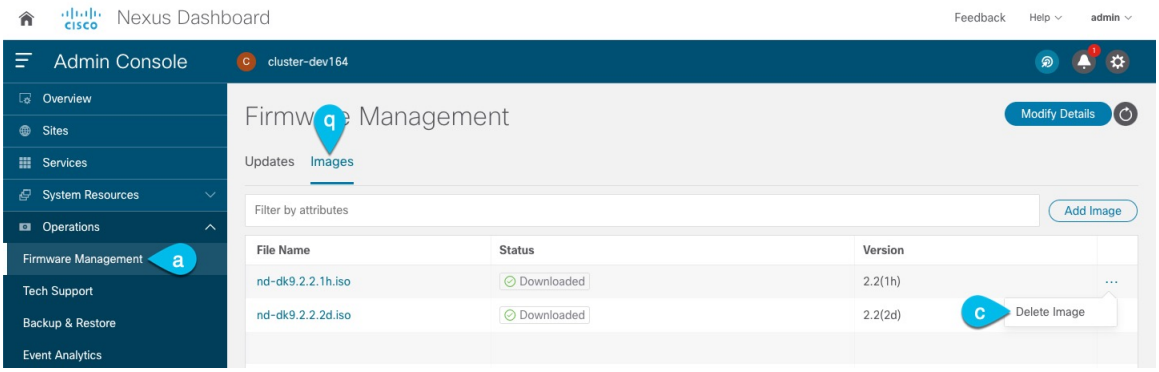
b) In the service's tile, click the **Actions** (**…**) menu and choose **Disable**.

c) Repeat this step for all other services deployed in the cluster.

Step 4 Delete any existing upgrade images from your cluster.

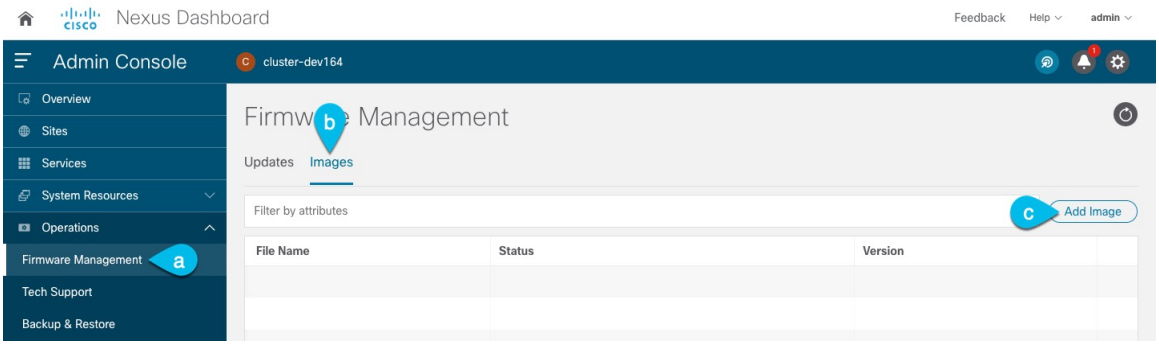
If this is the first time you're upgrading your cluster, you can skip this step.

If you had previously upgraded your cluster to your current version, you must delete all previous upgrade images.



a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Images** tab.
c) From the actions menu (**…**) next to an existing upgrade image, choose **Delete Image**.
d) Repeat this step for all existing upgrade images.

Step 5 Upload the new image to the cluster.



a) Navigate to **Operations** > **Firmware Management**.
b) Select the **Images** tab.
c) Click **Add Image**.

Step 6 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` or `.tar` image you downloaded in the first step.

   For RHEL deployments, use the `.tar` file to upgrade. For all other deployment factors, use the `.iso` file.

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

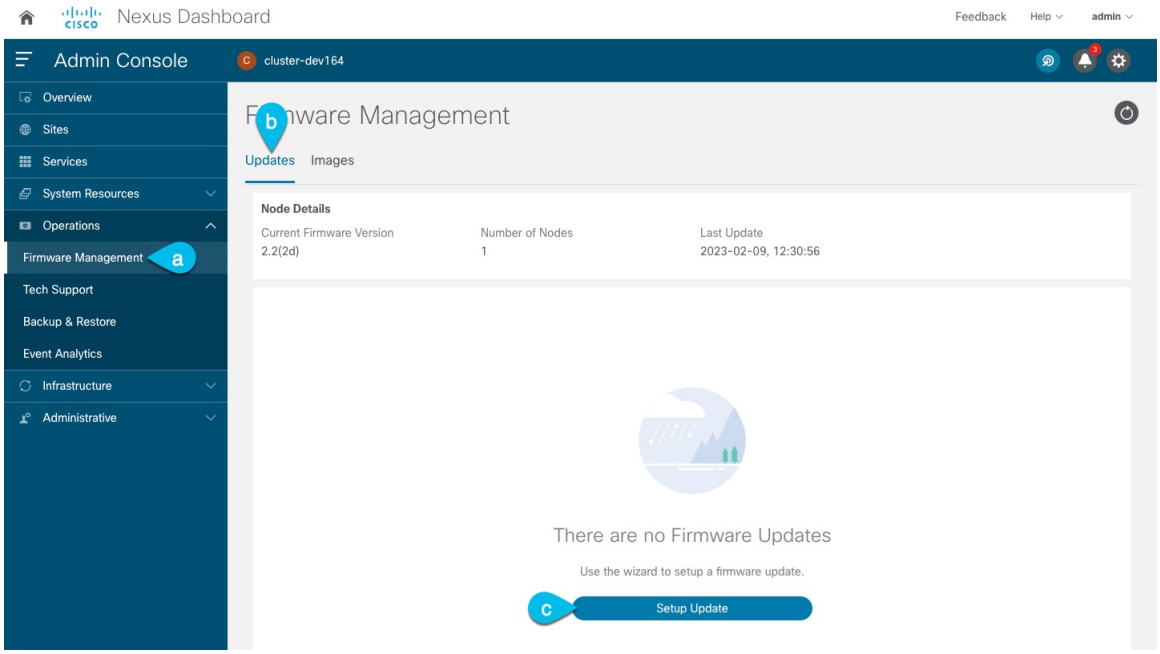
c) Click **Upload** to add the image.

   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 7 Wait for the image status to change to `Downloaded`.

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

Step 8 Set up the update.



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

Note
If you had upgraded your cluster before, the page shows the previous upgrade's details instead. In that case, click the **Modify Details** button in the top right of the page to provide new upgrade information.

The **Firmware Update** screen opens.

Step 9 Start the upgrade.

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 **Validate**.

The setup goes through a number of preparation and validation stages to ensure successful upgrade. This may take several minutes to complete.

c) After the validation is complete, click **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. This step may take up to 20 minutes.

Step 10 Activate the new image.

If you never navigated away from the upgrade screen, simply click **Activate** to activate the new image.

Otherwise:

a) Navigate back to the **Operations** > **Firmware Management** screen
b) In the **Last Update Status** tile, click **Continue**.

In some of the earlier Nexus Dashboard versions, this link may be called **View Details** instead.

c) In the **Firmware Update** > **Install** screen, click **Activate**.

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 11 Upgrade the individual services deployed in the cluster.

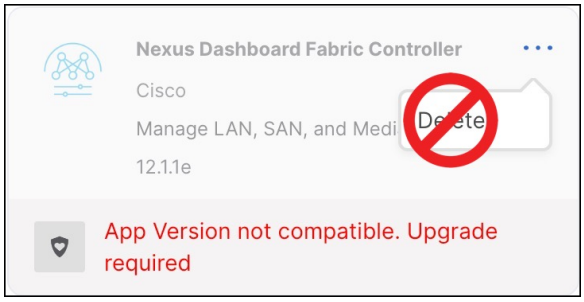
Ensure that you have read the upgrade guides specific to the services that you are upgrading:

- Nexus Dashboard Fabric Controller Upgrade Guide
- Nexus Dashboard Insights Upgrade Guide
- Nexus Dashboard Orchestrator Upgrade Guide

Note
After you have upgraded your Nexus Dashboard to release 3.0(1) or later, you must upgrade all existing services to the versions compatible with the new Nexus Dashboard version.

When upgrading the services, you must not re-enable the existing versions that you had disabled prior to the Nexus Dashboard upgrade.

You must also not delete the existing versions of the service, which may now display an `App Version not compatible. Upgrade required.` error:



Step 12 (Optional) Migrate to the new UCS-C225-M6 hardware.

Note
If you do not plan to replace your Nexus Dashboard nodes with the new UCS-C225-M6 servers, you can skip this step.

To migrate your existing Nexus Dashboard cluster deployed using UCS-C220-M5 hardware, you can simply add a new UCS-C225-M6 node as a `stadby` node to the existing cluster and fail over one of the older nodes. Then repeat the process one node at a time for the remaining nodes of the older cluster. Adding and using `stadby` nodes is described in detail in the "Infrastructure Management" chapter of the *Nexus Dashboard User Guide*.

Upgrading Nexus Dashboard