

Cisco Nexus Dashboard and Services Deployment and Upgrade Guide, Release 3.1.x

First Published: 2024-03-07

Last Modified: 2024-08-12

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)

© 2024–2025 Cisco Systems, Inc. All rights reserved.



CONTENTS

Trademarks ?

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

PART I **Preparing to Deploy Nexus Dashboard** 3

CHAPTER 2 **Deployment Overview and Requirements** 5

Deployment Overview 5

CHAPTER 3 **Prerequisites: Nexus Dashboard** 9

Prerequisites and Guidelines 9

Communication Ports 15

Fabric Connectivity 17

Node Distribution Across Sites 23

Services Co-location Use Cases 24

Pre-Installation Checklist 27

CHAPTER 4 **Prerequisites: Fabric Controller** 31

Requirements for Fabric Controller 31

Communication Ports for Fabric Controller 33

CHAPTER 5 **Prerequisites: Orchestrator** 45

Requirements for Orchestrator 45

Communication Ports for Orchestrator 46

Fabric Requirements for Orchestrator 46

Pod Profile and Policy Group 47
Configuring Fabric Access Global Policies 48
Configuring Fabric Access Interface Policies 49

CHAPTER 6 **Prerequisites: Insights** 53
Requirements for Insights 53
Communication Ports for Insights 54
Fabric Requirements for Insights 55

PART II **Deploying the Cluster** 59

CHAPTER 7 **Deploying as Physical Appliance** 61
Prerequisites and Guidelines 61
Physical Node Cabling 64
Deploying Nexus Dashboard as Physical Appliance 66

CHAPTER 8 **Deploying in VMware ESX** 81
Prerequisites and Guidelines 81
Deploying Nexus Dashboard Using VMware vCenter 84
Deploying Nexus Dashboard Directly in VMware ESXi 103

CHAPTER 9 **Deploying in Linux KVM** 119
Prerequisites and Guidelines 119
Deploying Nexus Dashboard in Linux KVM 120

CHAPTER 10 **Deploying in Amazon Web Services** 137
Prerequisites and Guidelines 137
Deploying Nexus Dashboard in AWS 139

CHAPTER 11 **Deploying in Microsoft Azure** 151
Prerequisites and Guidelines 151
Generating an SSH Key Pair in Linux or MacOS 152
Generating an SSH Key Pair in Windows 153

Deploying Nexus Dashboard in Azure 156

CHAPTER 12 **Onboarding Fabrics** 165

Onboarding ACI Fabrics 165

Onboarding NDFC Fabrics 166

Onboarding NX-OS Switches 168

PART III **Upgrading or Migrating to This Release** 171

CHAPTER 13 **Upgrading Existing ND Cluster to This Release** 173

Prerequisites and Guidelines 173

Supported Upgrade Paths 175

Upgrading Nexus Dashboard 178

Troubleshooting Upgrades 183

Resolving Nexus Dashboard Orchestrator Configuration Issues 183

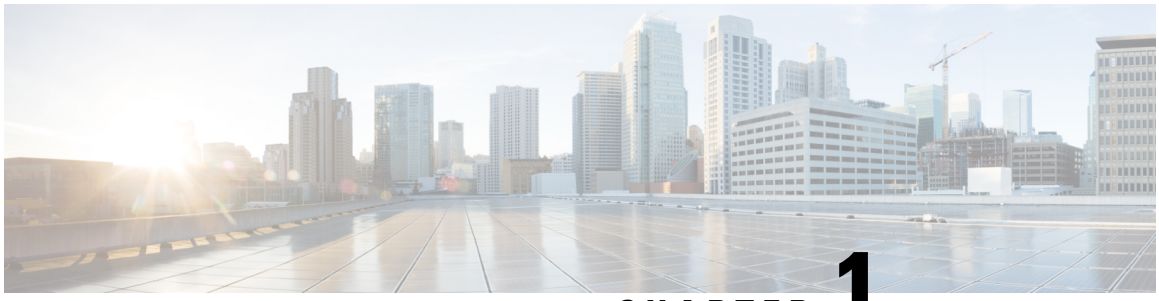
Resolve Configuration Inconsistencies In Templates 184

Resolve Configuration Drifts 184

CHAPTER 14 **Migrating From DCNM to NDFC** 187

Prerequisites and Guidelines 187

Migrate Existing DCNM Configuration to NDFC 189



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
3.1(1)	First release of this document. Individual platform and services installation guides have been combined into this document to reflect the unified deployment workflow in this release.	--

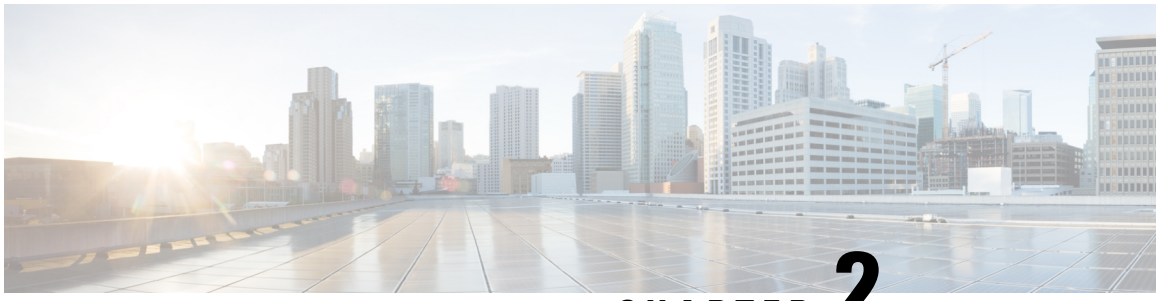
New and Changed Information



PART I

Preparing to Deploy Nexus Dashboard

- Deployment Overview and Requirements, on page 5
- Prerequisites: Nexus Dashboard, on page 9
- Prerequisites: Fabric Controller, on page 31
- Prerequisites: Orchestrator, on page 45
- Prerequisites: Insights, on page 53



CHAPTER 2

Deployment Overview and Requirements

• Deployment Overview, on page 5

Deployment Overview

Nexus Dashboard Platform

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 Insights, Orchestrator, and Fabric Controller. 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.



Note This document describes how to initially deploy a Nexus Dashboard cluster, enable one or more services, and onboard the fabrics to be managed by those services. After your cluster is up and running, see the Nexus Dashboard *configuration and operation articles* as well as the service-specific documentation for day-to-day operation.

Nexus Dashboard Services and Unified Installation in Release 3.1(1)

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. Nexus Dashboard services (such as Insights, Orchestrator, and Fabric Controller) can be enabled with the Nexus Dashboard platform to 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, Cisco NDFC, or Standalone NX-OS switches without a controller.

Prior to release 3.1(1), Nexus Dashboard shipped with only the platform software and no services included, which you would then download, install, and enable separately after the initial platform deployment. Beginning with release 3.1(1), the platform and the individual services have been unified into a single image and can be deployed and enabled during the initial cluster configuration for a simpler, more streamlined experience.

Note While prior releases of Nexus Dashboard platform often supported multiple versions of the individual services, because of the unified installation, each Nexus Dashboard release supports a single specific version of each service, which you can choose to enable during cluster deployment.

In addition, the unified installation ensures that only the supported combinations of services can be enabled during cluster deployment, which prevents the possibility of unsupported service combinations. Detailed information about service cohosting is available in the Nexus Dashboard Cluster Sizing tool.

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.

This guide describes the initial deployment of the Nexus Dashboard software, which is common for physical as well as virtual or cloud form factors; if you are deploying a physical cluster, see *Nexus Dashboard Hardware Setup Guide* for the UCS servers' hardware overview, specification, and racking instructions..

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" article in the Nexus Dashboard *documentation library*.

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 nodes of 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.

- Physical appliance (`.iso`)

  This form factor refers to the Cisco UCS physical appliance hardware with the 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 Nexus Dashboard hardware is described in *Nexus Dashboard Hardware Setup Guide* for the specific UCS model.

- Virtual appliance
  - VMware ESX (`.ova`)

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

  - Linux KVM (`.qcow2`)

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

- Public cloud
  - Amazon Web Services (`.ami`)

    Cloud form factor that allows you to deploy a Nexus Dashboard cluster using AWS instances.
  - Microsoft Azure (`.arm`)

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



Note Not all services are supported on all form factors, and some form factors support only certain services. When planning your deployment, ensure to check the list of "Prerequisites and Guidelines" in one of the following sections of this document specific to the form factor you are deploying. A quick reference of the supported form factors, services, scale, and cluster sizing requirements are available in the Nexus Dashboard Cluster Sizing tool.

Scale and Cluster Sizing Guidelines

A basic Nexus Dashboard deployment typically consists of 1 or 3 `primary` nodes, which are required for the cluster to come up. Depending on services and scale requirements, 3-node clusters can be extended with additional `secondary` nodes to support cohosting of services and higher scale. For physical clusters, you can also add up to 2 `standby` nodes for easy cluster recovery in case of a primary node failure.



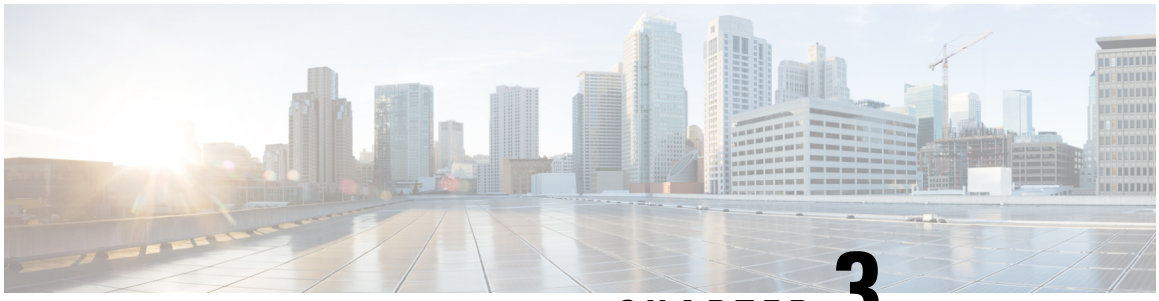
Note

- Single-node deployments are supported for a limited number of services and cannot be extended to a 3-node cluster after the initial deployment.
- Single-node deployments do not support additional `secondary` or `standby` nodes.
- If you deploy a single-node cluster and want to extended it to a 3-node cluster or add `secondary` nodes, you will first need to redeploy it as a 3-node base 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 "Troubleshooting" article in the Nexus Dashboard *documentation library*.
- 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.

Exact number of additional secondary nodes required for your specific use case is available from the Nexus Dashboard Cluster Sizing tool.



CHAPTER 3

Prerequisites: Nexus Dashboard

- Prerequisites and Guidelines, on page 9
- Communication Ports, on page 15
- Fabric Connectivity, on page 17
- Node Distribution Across Sites, on page 23
- Services Co-location Use Cases, on page 24
- Pre-Installation Checklist, on page 27

Prerequisites and Guidelines

Note This section describes requirements and guidelines that are common for all services enabled in your Nexus Dashboard cluster. Additional service-specific requirements are listed in the following sections of this document.

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, as well as impact regular services functionality.

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.

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 will need to provide the following information during cluster configuration:

- **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.

Nexus Dashboard External Networks

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, as described in the following sections.

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 • 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 a Nexus Dashboard cluster running only Nexus Dashboard Fabric Controller service, which can be deployed using the same subnets for the data and management networks. Note that only Nexus Dashboard Fabric Controllers running in SAN controller mode can be deployed using the same subnets for the data and management networks; Nexus Dashboard Fabric Controllers running in LAN controller mode must be deployed in different subnets for the data and management networks.

- 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.
- 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.
- 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 on the switches to which the nodes are connected .



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 on the switch ports where the nodes are connected.

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

Note RTT requirements for connectivity from the Nexus Dashboard cluster to the site controllers or switches depends on the specific service you plan to enable, see the "Network Requirements" sections in the service-specific chapters below.

Table 3: Cluster RTT Requirements

Connectivity	Maximum RTT
Between nodes within the same Nexus Dashboard cluster	50 ms
Between nodes in one cluster and nodes in a different cluster if the clusters are connected via multi-cluster connectivity For more information about multi-cluster connectivity, see *Cisco Nexus Dashboard Infrastructure Management*.	500 ms
Between external DNS servers and the Nexus Dashboard cluster	5 seconds

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.

    Mixed 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.
  - Kubernetes internal core services will start in IPv4 mode.
  - DNS will serve and forward both IPv4 and IPv6 requests.
  - 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 and Azure do not support pure IPv6 mode.

- You must provide IPv6 management network addresses when initially configuring the nodes.

  After the nodes 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 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.

BGP Configuration and Persistent IPs

Some prior releases of Nexus Dashboard allowed you to configure one or more persistent IP addresses for services (such as Insights and Fabric Controller) 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

While that is still supported, this release also allows you to configure the Persistent IPs feature 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 *Infrastructure Management* document.
- Ensure that the persistent IP addresses you allocate do not overlap with any of the nodes' management or data subnets.
- Ensure that you fulfill the service-specific persistent IP requirements listed in the service-specific sections that follow.

The total number of persistent IPs required for each service is listed in the service-specific requirements sections that follow.

Communication Ports

The following ports are required by the Nexus Dashboard cluster.

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

Table 4: 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
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 5: 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, 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
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

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
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

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

Connecting via External Layer 3 Network

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

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

  Note that if the fabric connectivity is from the Nexus Dashboard's management interface, you must configure specific static routes or ensure that the management interface is part of the same IP subnet of the APIC interfaces.
- 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 in-band 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 1: Connecting via Layer 3 Network, Day-2 Operations Applications

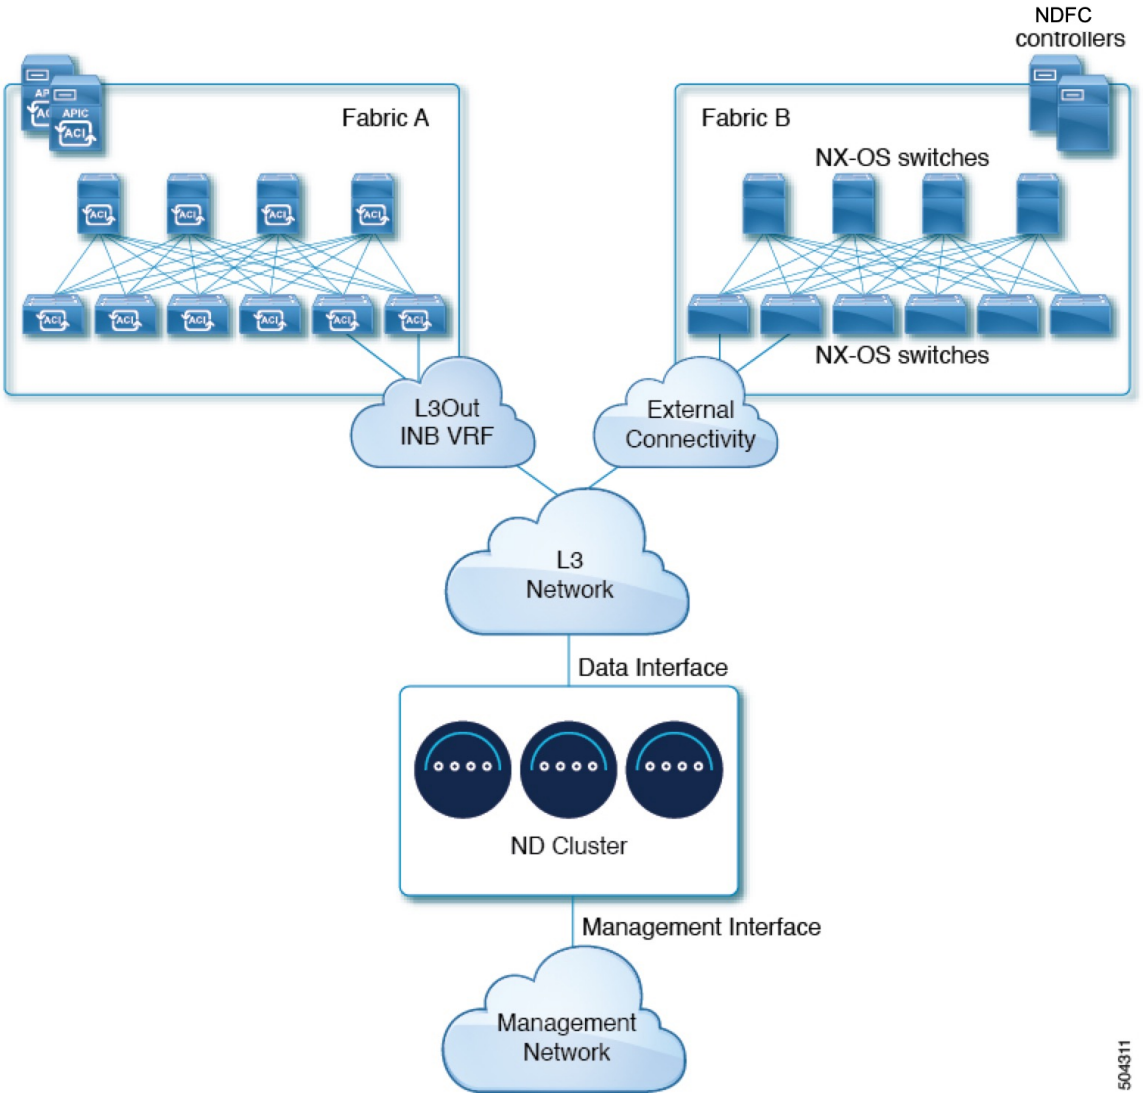
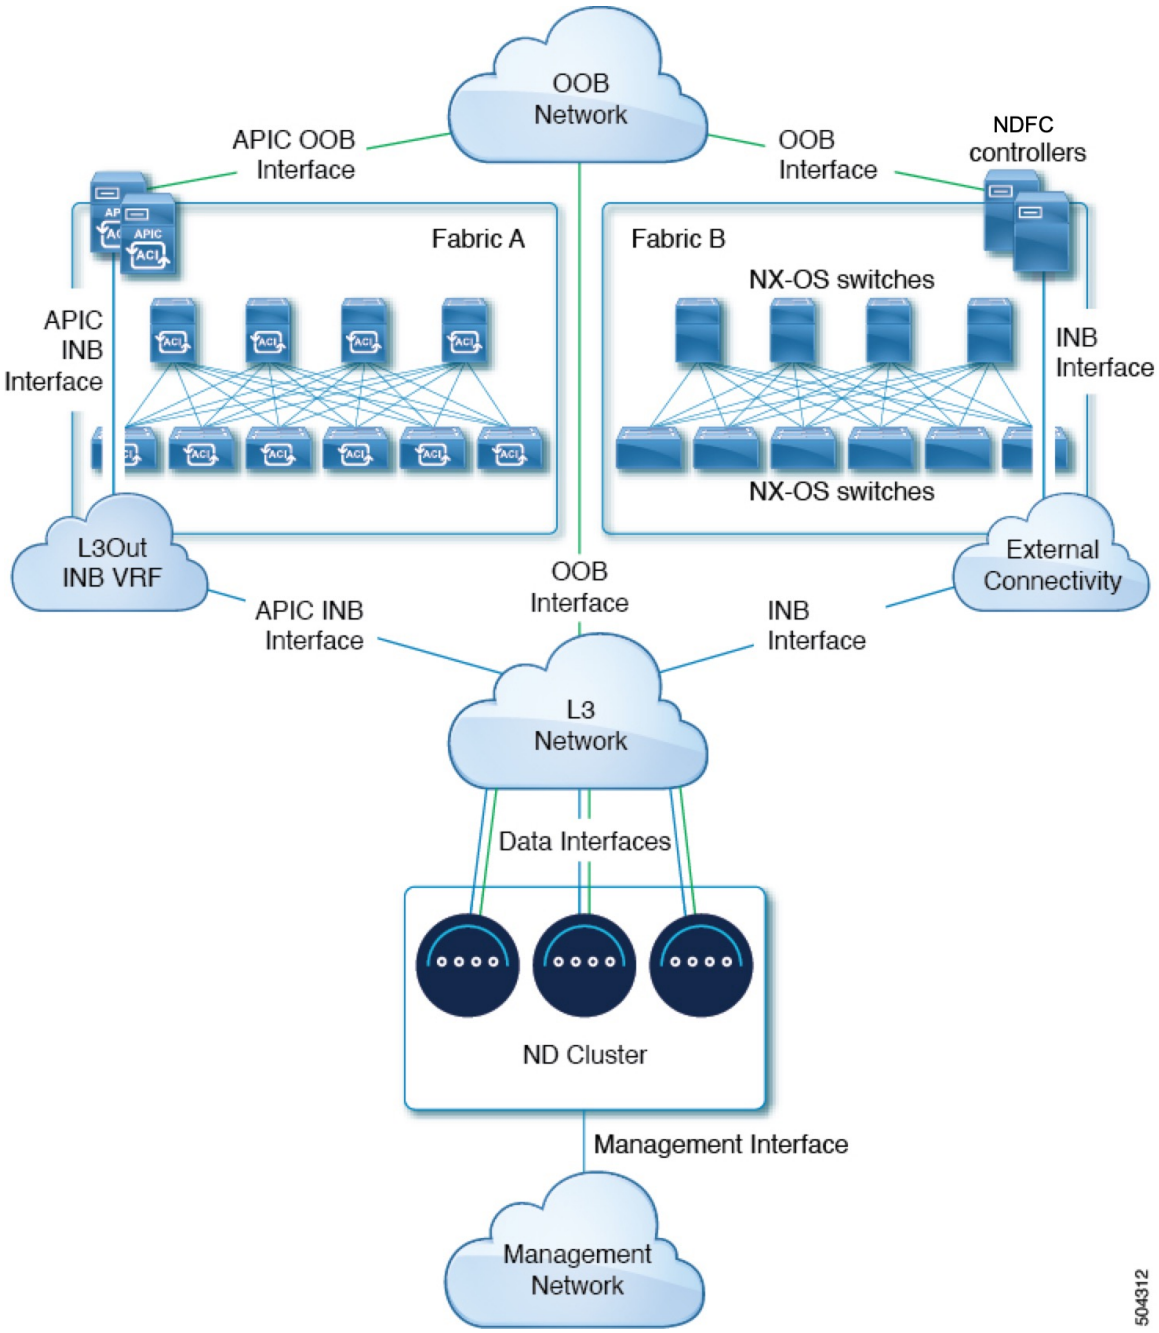


Figure 2: Connecting via Layer 3 Network, Nexus Dashboard Orchestrator



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 either the data interface or the management interface to either the in-band or out-of-band (OOB) interface of each site's APIC or both.

  Note that if the fabric connectivity is from the Nexus Dashboard's management interface, you must configure specific static routes or ensure that the management interface is part of the same IP subnet of the APIC interfaces.

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

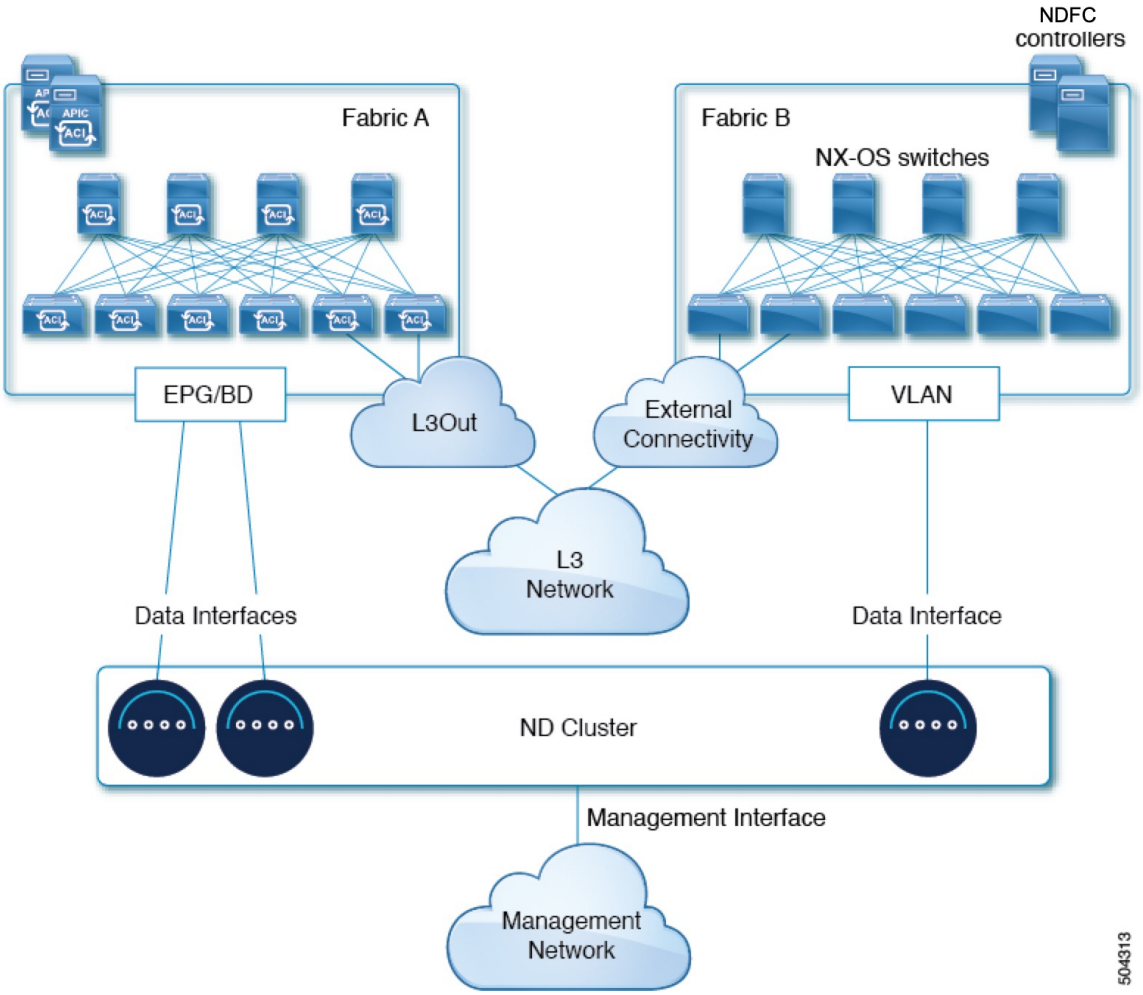
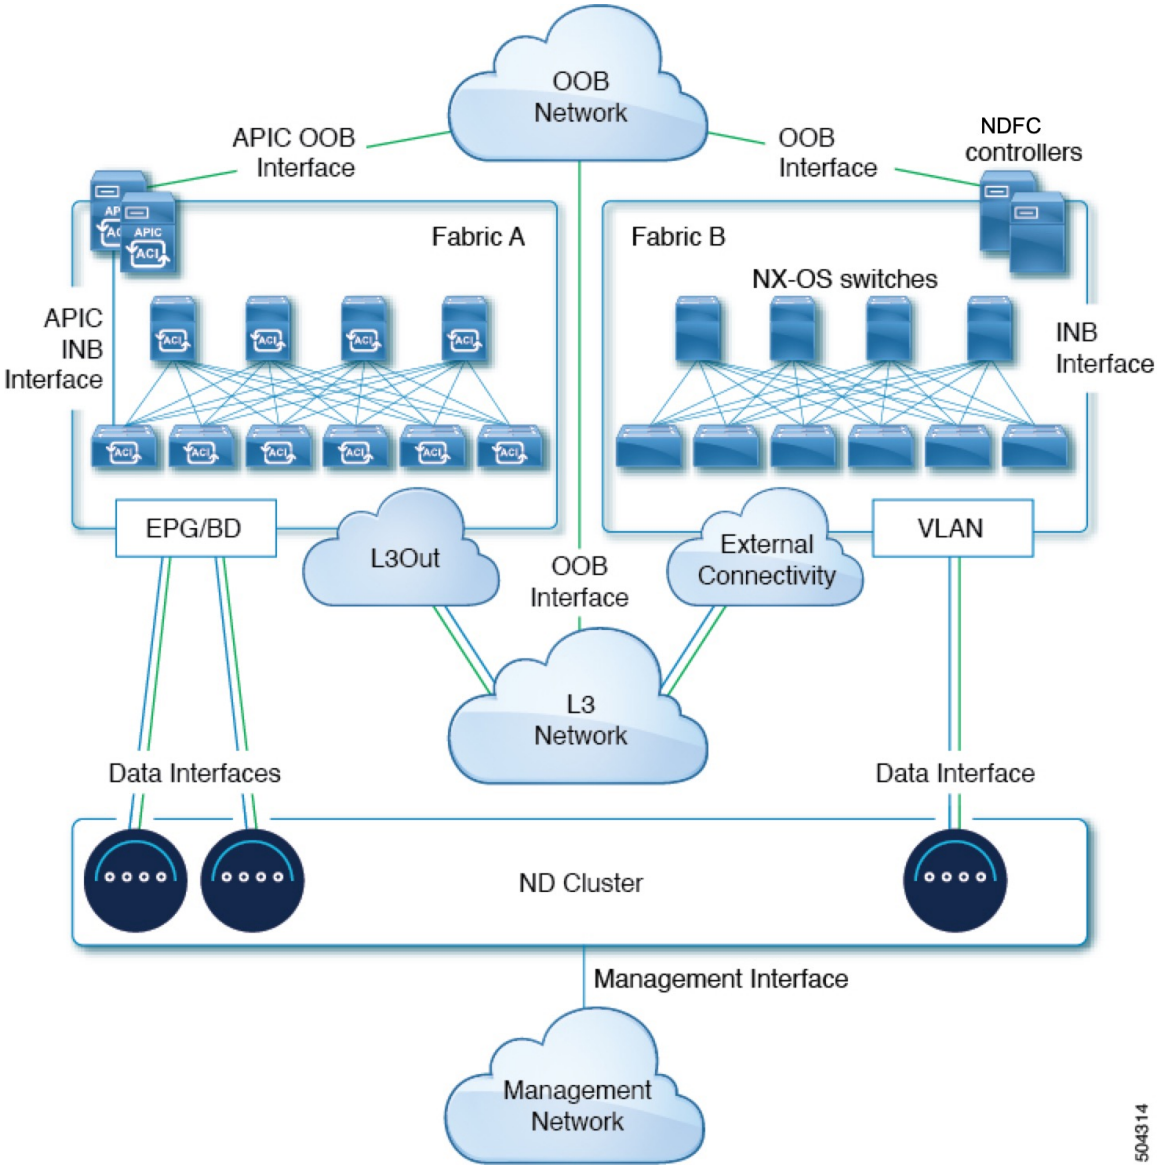


Figure 4: Connecting Directly to Leaf Switches, Nexus Dashboard Orchestrator



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.

Node Distribution for Nexus Dashboard Insights

For Nexus Dashboard Insights, we recommend a centralized, single-site deployment. This service does not support recovery 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

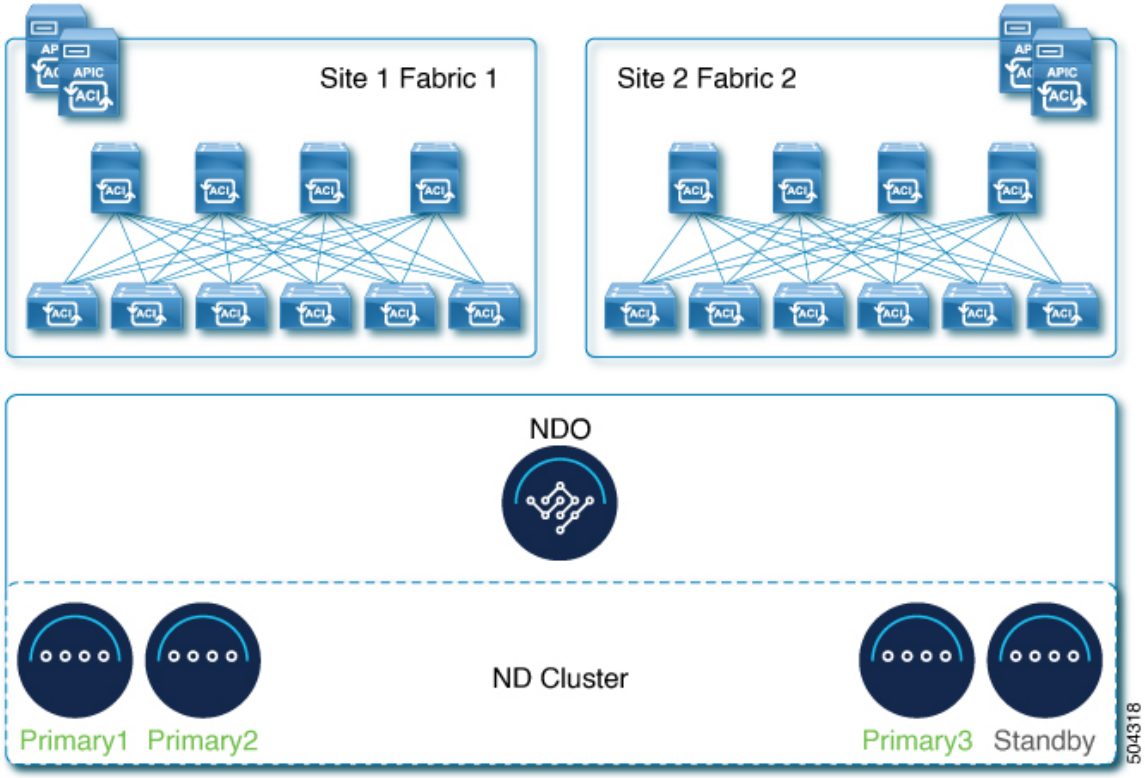
For Nexus Dashboard Fabric Controller, we recommend a centralized, single-site deployment. This service does not support recovery if two `primary` nodes are not available and so it gains no redundancy benefits from a distributed cluster, which could instead expose the cluster to interconnection failures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

Note Standby nodes are supported only for physical clusters. For virtual clusters, you can simply bring up a new VM with identical settings as the failed node.

Figure 5: Node Distribution Across Two Sites for Nexus Dashboard Orchestrator



Services Co-location Use Cases

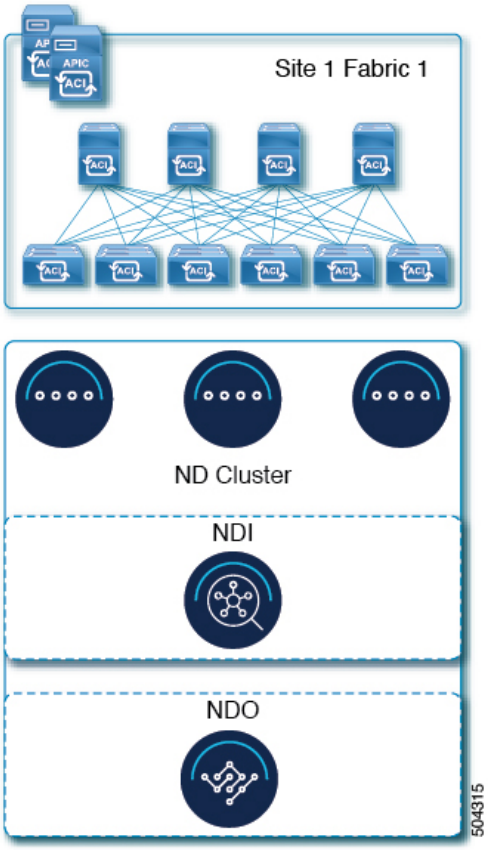
This section describes a number of recommended deployment scenarios for specific single-service or multiple services co-hosting use cases.

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.

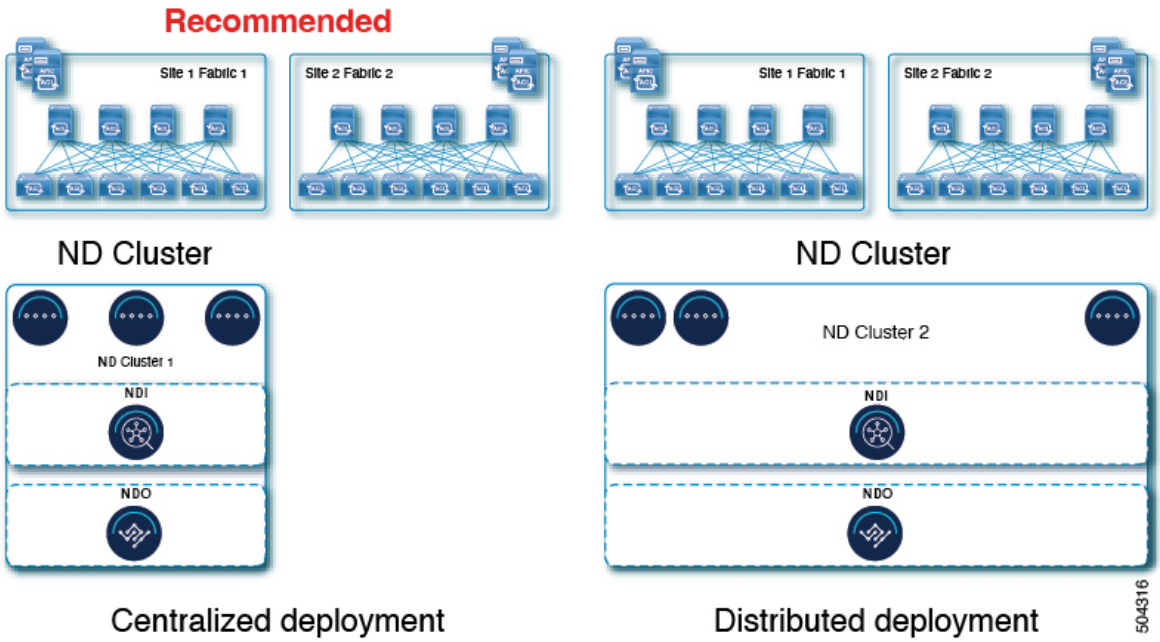
Figure 6: Single Site, Nexus Dashboard Insights and Orchestrator



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:

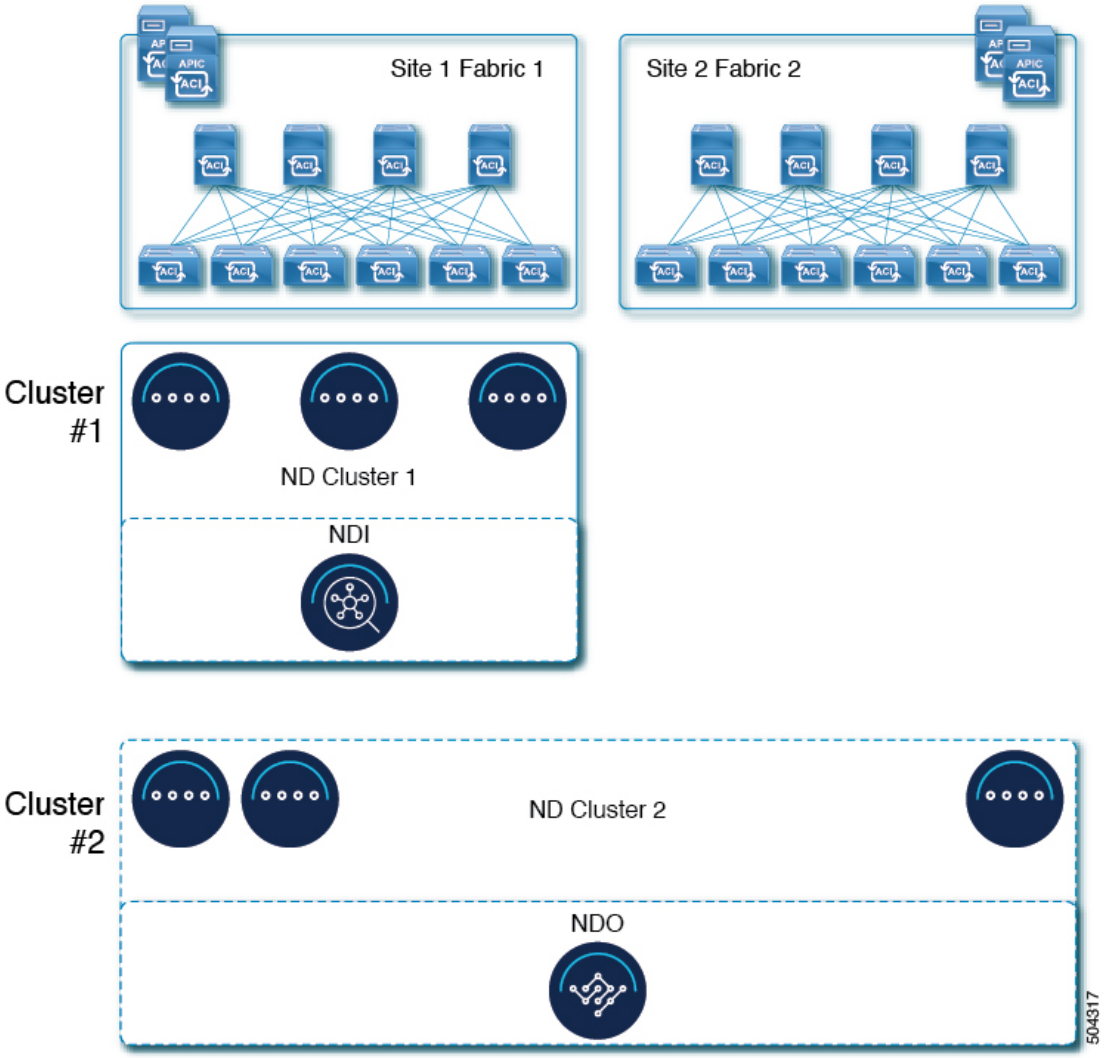
Figure 7: Multiple Sites, Single Cluster for Nexus Dashboard Insights and Orchestrator



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 8: 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 6: Cluster Details

Parameters	Example	Your Entry
Cluster **Name**	`nd-cluster`	
NTP Server	`170.78.48.55`	
DNS Provider	`170.71.68.83`	

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

Note Beginning with release 3.1(1), you can define all nodes during the initial cluster deployment, including the `secondary` and `standby` nodes. For simplicity, the following tables assumes a 3-node base cluster, but if you are deploying a larger cluster, you must also have the node details for all additional nodes.

Table 7: Node Details

Parameters	Example	Your Entry
For physical nodes, **CIMC** address and login information of the first node	`10.196.220.84/24` Username: `admin` Password: `Cisco1234`	
For physical nodes, **CIMC** address and login information of the second node	`10.196.220.85/24` Username: `admin` Password: `Cisco1234!`	
For physical nodes, **CIMC** address and login information of the third node	`10.196.220.86/24` 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	`192.168.11.172/24`	
Management Gateway of the first node.	`192.168.11.1`	
Data Network IP of the first node	`192.168.8.172/24`	
Data Network Gateway of the first node	`192.168.8.1`	
(Optional) **Data Network VLAN** of the first node	`101`	

Parameters	Example	Your Entry
If you enable BGP, **ASN** of the first node	`63331`	
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`	

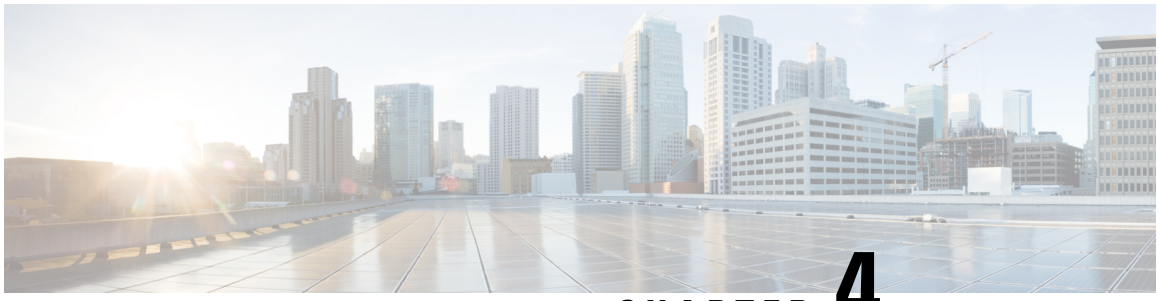
If you enable BGP, IP addresses of the second node's **BGP Peer**(s)	`200.12.12.2`
	or
	`200:12:12::2`	

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`	

Parameters	Example	Your Entry
Data Network Gateway of the third node	`192.168.6.1`	
(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 4

Prerequisites: Fabric Controller

- Requirements for Fabric Controller, on page 31
- Communication Ports for Fabric Controller, on page 33

Requirements for Fabric Controller

Overview

Nexus Dashboard Fabric Controller (NDFC) is the comprehensive management solution for all NX-OS deployments spanning LAN Fabric, SAN, and IP Fabric for Media (IPFM) networks in data centers powered by Cisco. Cisco Nexus Dashboard Fabric Controller also supports other devices, such as IOS-XE switches, IOS-XR routers, and non-Cisco devices. Being a multi-fabric controller, Cisco Nexus Dashboard Fabric Controller manages multiple deployment models like VXLAN EVPN, Classic 3-Tier, FabricPath, and Routed based fabrics for LAN while providing ready-to-use control, management, monitoring, and automation capabilities for all these environments. In addition, Cisco NDFC when enabled as a SAN Controller automates Cisco MDS Switches and Cisco Nexus Family infrastructure in NX-OS mode with a focus on storage-specific features and analytics capabilities.

NDFC primarily focuses on Control and Management for three primary market segments:

- LAN networking including VXLAN, Multi-Site, Classic Ethernet, and External Fabrics supporting Cisco Nexus switches running standalone NX-OS, with additional support for IOS-XR, IOS-XE, and adjacent Host, Compute, Virtual Machine, and Container Management systems.
- SAN networking for Cisco MDS and Cisco Nexus switches running standalone NX-OS, including support for integration with storage arrays and additionally Host, Compute, Virtual Machine, and Container Orchestration systems.
- Media Control for Multicast Video production networks running Cisco Nexus switches operated as standalone NX-OS, with additional integrations for 3rd party media control systems.

After you deploy Nexus Dashboard using a deployment mode that includes NDFC:

- **Fabric Discovery**—Discover, Monitor, and Visualize LAN Deployments.
- **Fabric Controller**—LAN Controller for Classic Ethernet (vPC), Routed, VXLAN, and IP Fabric for Media Deployments.
- **SAN Controller**—SAN Controller for MDS and Nexus switches. Enhanced SAN Analytics with streaming telemetry.

Network Requirements

Note This section describes *additional* requirements and guidelines if you plan to enable the Fabric Controller service. Ensure that you have already satisfied the platform-level requirements described in the Prerequisites and Guidelines, on page 9 section.

- Starting with Nexus Dashboard release 3.1.1, Cisco DC App Center connectivity has been removed from Nexus Dashboard because downloading the services separately is no longer required.

  To deploy Fabric Controller, download the unified installation image from the Software Download page; individual services' installation images are no longer available from the Cisco DC App Center.

- As mentioned in the previous section, all new Nexus Dashboard deployments must have the management network and data network in different subnets.

Note Only SAN Controller persona can be deployed in Nexus Dashboard using the same subnets for the data and management networks.

- Interfaces on both Data and Management networks can be either Layer 2 or Layer 3 adjacent.
- Connectivity between the nodes is required on both networks with the following additional round trip time (RTT) requirements:

Table 8: Fabric Controller RTT Requirements

Connectivity	Maximum RTT
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.

- You must allocate the following number of persistent IP addresses depending on your use case.

  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, one of the following:
    - 1 additional IP in the **management network** for telemetry for single node ND
    - 3 additional IPs in the **management network** for telemetry in a 3 node ND cluster

  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, one of the following:

    - 1 additional IP in the **data network** for telemetry for single node ND
    - 3 additional IPs in the **data network** for telemetry for multi-node ND cluster

- When operating in Layer 3 mode with LAN deployment type, **LAN Device Management Connectivity** must be set to `Data` and all persistent IPs must be part of a separate pool that must not overlap with the ND management or data subnets.

When operating in Layer 2 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

For an overview of Persistent IP functionality, see Prerequisites and Guidelines, on page 9. Allocating persistent IP addresses can be done during the initial cluster deployment or after the cluster is deployed using the External Service Pools configuration in the UI.

Communication Ports for Fabric Controller

In addition to the ports required by the Nexus Dashboard cluster nodes (listed in a previous section), the following ports are required by the Fabric Controller service.

- 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 9: 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.

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)
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. This applies to LAN deployments only. NX-API is an optional feature, used by limited set of NDFC functions, such as: • Endpoint locator (EPL), as described in **EPL Behavior with NX-API Configuration** • Layer 4 to Layer 7 services, as described in **Guidelines and Limitations for L4-L7 Services** You must enable the NX-API feature in the **Advanced** tab for the appropriate fabrics, such as the VXLAN, Enhanced Classic LAN, eBGP, and Campus fabrics.

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)
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
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.

- The following ports apply to the External Service IPs, also known as persistent IPs, used by some of the NDFC services:

  Note that these External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 10: 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.
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.

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)
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.
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.

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)
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
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

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)
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	
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.

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)	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.

- The following ports are required for NDFC SAN deployments on single-node clusters:

Table 11: 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.
SMTP	25	TCP	Out	SMTP port is configurable through NDFC's Server Settings menu. 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)
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.
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.

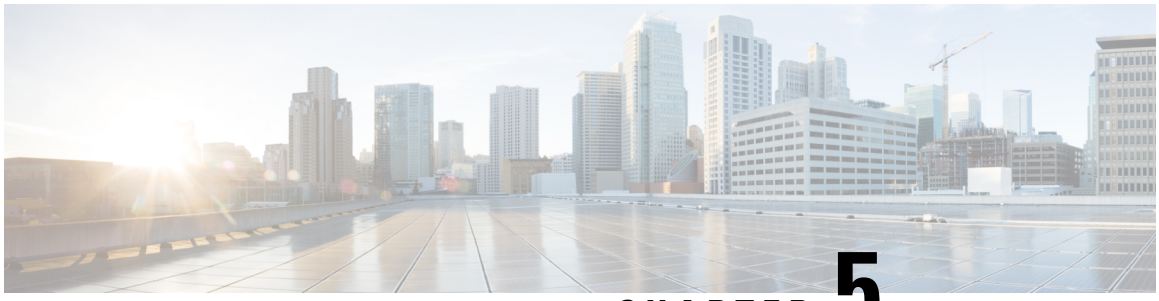
- The following ports apply to the External Service IPs, also known as Persistent IPs, used by some of the NDFC services:

  Note that these External Service IPs may come from the Nexus Dashboard management subnet pool or the data subnet pool depending on the configured settings.

Table 12: 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.



CHAPTER 5

Prerequisites: Orchestrator

- Requirements for Orchestrator, on page 45
- Communication Ports for Orchestrator, on page 46
- Fabric Requirements for Orchestrator, on page 46

Requirements for Orchestrator

Note This section describes *additional* requirements and guidelines if you plan to enable the Orchestrator service. Ensure that you have already satisfied the platform-level requirements described in the Prerequisites and Guidelines, on page 9 section.

- Starting with Nexus Dashboard release 3.1.1, Cisco DC App Center connectivity has been removed from Nexus Dashboard because downloading the services separately is no longer required.

  To deploy Orchestrator, download the unified installation image from the Software Download page; individual services' installation images are no longer available from the Cisco DC App Center.

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

  Note that if the fabric connectivity is from the Nexus Dashboard's management interface, you must configure specific static routes or ensure that the management interface is part of the same IP subnet of the APIC interfaces.

- If you are deploying Nexus Dashboard Orchestrator to manage Cisco NDFC fabrics, the data network must have in-band reachability for Cisco NDFC sites.

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

Table 13: Orchestrator RTT Requirements

Connectivity	Maximum RTT
To any managed APIC sites	500 ms

Connectivity	Maximum RTT
To any managed NDFC sites	150 ms

Communication Ports for Orchestrator

In addition to the ports required by the Nexus Dashboard cluster nodes (listed in a previous section), the following ports are required by the Orchestrator service.

Table 14: Nexus Dashboard Orchestrator Ports (Management Network)

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SCP or SFTP	22	TCP	In/Out	Remote servers for storing backups and downloading software upgrade images
HTTP	80	TCP	Out	Splunk or syslog server if external log streaming is enabled
HTTPS	443	TCP	In/Out	Splunk or syslog server if external log streaming is enabled

Table 15: Nexus Dashboard Orchestrator Ports (Data Network)

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

Fabric Requirements for Orchestrator

The following additional fabric-related guidelines apply to the Orchestrator service:

- Cisco Mini ACI fabrics are supported as typical on-premises sites without requiring any additional configuration.

  Detailed info on deploying and configuring this type of fabrics is available in *Cisco Mini ACI Fabric and Virtual APICs*.

- If you are managing ACI fabrics that contain Remote Leaf switches, the following restrictions apply:
  - Only physical Remote Leaf switches are supported.
  - Only -EX and -FX or later switches are supported as Remote Leaf switches.
  - Remote Leaf is not supported with back-to-back connected sites without IPN switches.
  - Remote Leaf switches in one site cannot use another site's L3Out.
  - Stretching a bridge domain between one site (local leaf or remote leaf) and a Remote Leaf in another site is not supported.

  You must also perform the following tasks before the site can be added to and managed by the Nexus Dashboard Orchestrator:

  - You must enable Remote Leaf direct communication directly in the site's APIC.

    To enable direct communication, log in to the site's APIC, navigate to **System** > **System Settings** > **Fabric Wide Setting** and **Enable Remote Leaf Direct Traffic Forwarding**.

Note You cannot disable this option after you enable it.

  - You must configure external TEP pools for remote leaf switches.

    To configure one or more external TEP pools, log in to the site's APIC and navigate to **Fabric** > **Inventory** > **Pod Fabric Setup Policy**. Then double-click the pod where you want to configure the subnets and click **+** in the **External TEP** area. Finally, enter the **IP** and **Reserve Address Count**, set the state to `Active` or `Inactive`, then click **Update** to save the subnet.

    When configuring external TEP pools, you must provide a netmask between `/22` and `/29`. Multiple, non-contiguous, external TEP pools can be configured, including at different points in time.

  - You must add the routable IP addresses of APIC nodes (assigned from the defined external TEP pool) in the DHCP-Relay configuration applied on the interfaces of the Layer 3 routers connecting to the Remote Leaf switches.

    The routable IP address of each APIC node is listed in the **Routable IP Address** field of the **System** > **Controllers** > **<controller-name>** screen of the APIC GUI.

- You must configure Pod Profile, Policy Group, and Fabric Access policies as described in the following sections.

Pod Profile and Policy Group

In each site's APIC, you must have one Pod profile with a Pod policy group. If your site does not have a Pod policy group you must create one. Typically, these settings will already exist as you will have configured them when you first deployed the fabric.

Procedure

Step 1 Log in to the site's APIC GUI.

Step 2 Check that the Pod profile contains a Pod policy group.

Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**.

Step 3 If necessary, create a Pod policy group.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Policy Groups**.
b) Right-click **Policy Groups** and select **Create Pod Policy Group**.
c) Enter the appropriate information and click **Submit**.

Step 4 Assign the new Pod policy group to the default Pod profile.

a) Navigate to **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**
b) Select the default profile.
c) Choose the new pod policy group and click **Update**.

Configuring Fabric Access Global Policies

This section describes the global fabric access policy configurations that must be created for each APIC site before it can be on-boarded on the Nexus Dashboard cluster and managed by the Nexus Dashboard Orchestrator.

Procedure

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

You must configure a number of fabric policies before the site can be managed by the Nexus Dashboard Orchestrator. From the APIC's perspective, this is something you do just like you would if you were connecting a bare-metal host, where you would configure domains, AEPs, policy groups, and interface selectors; you must configure the same options for connecting the spine switch interfaces to the inter-site network for all the sites that will be part of the same Multi-Site domain.

Step 3 Specify the VLAN pool.

The first thing you configure is the VLAN pool. We use Layer 3 sub-interfaces tagging traffic with VLAN-4 to connect the spine switches to the inter-site network.

a) In the left navigation tree, browse to **Pools** > **VLAN**.
b) Right-click the **VLAN** category and choose **Create VLAN Pool**.

In the **Create VLAN Pool** window, specify the following:

- For the **Name** field, specify the name for the VLAN pool, for example `msite`.
- For **Allocation Mode**, specify `Static Allocation`.

- And for the **Encap Blocks**, specify just the single VLAN 4. You can specify a single VLAN by entering the same number in both **Range** fields.

Step 4 Configure Attachable Access Entity Profiles (AEP).

a) In the left navigation tree, browse to **Global Policies** > **Attachable Access Entity Profiles**.

b) Right-click the **Attachable Access Entity Profiles** category and choose **Create Attachable Access Entity Profiles**.

In the **Create Attachable Access Entity Profiles** window, specify the name for the AEP, for example `msite-aep`.

c) Click **Next** and **Submit**

No additional changes, such as interfaces, are required.

Step 5 Configure the external routed domain.

The domain you configure is what you will select from the Nexus Dashboard Orchestrator when adding this site.

a) In the left navigation tree, browse to **Physical and External Domains** > **External Routed Domains**.

b) Right-click the **External Routed Domains** category and choose **Create Layer 3 Domain**.

In the **Create Layer 3 Domain** window, specify the following:

- For the **Name** field, specify the name the domain, for example `msite-l3`.
- For **Associated Attachable Entity Profile**, select the AEP you created in Step 4.
- For the **VLAN Pool**, select the VLAN pool you created in Step 3.

c) Click **Submit**.

No additional changes, such as security domains, are required.

What to do next

After you have configured the global access policies, you must still add interfaces policies as described in Configuring Fabric Access Interface Policies, on page 49.

Configuring Fabric Access Interface Policies

This section describes the fabric access interface configurations that must be done for the Nexus Dashboard Orchestrator on each APIC site.

Before you begin

You must have configured the global fabric access policies, such as VLAN Pool, AEP, and domain, in the site's APIC, as described in Configuring Fabric Access Global Policies, on page 48.

Procedure

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

In addition to the VLAN, AEP, and domain you have configured in previous section, you must also create the interface policies for the fabric's spine switch interfaces that connect to the Inter-Site Network (ISN).

Step 3 Configure a spine policy group.

a) In the left navigation tree, browse to **Interface Policies** > **Policy Groups** > **Spine Policy Groups**.

This is similar to how you would add a bare-metal server, except instead of a Leaf Policy Group, you are creating a Spine Policy Group.

b) Right-click the **Spine Policy Groups** category and choose **Create Spine Access Port Policy Group**.

In the **Create Spine Access Port Policy Group** window, specify the following:

- For the **Name** field, specify the name for the policy group, for example `Spine1-PolGrp`.
- For the **Link Level Policy** field, specify the link policy used between your spine switch and the ISN.
- For **CDP Policy**, choose whether you want to enable CDP.
- For the **Attached Entity Profile**, select the AEP you have configured in previous section, for example `msite-aep`.

c) Click **Submit**.

No additional changes, such as security domains, are required.

Step 4 Configure a spine profile.

a) In the left navigation tree, browse to **Interface Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Interface Profile**.

In the **Create Spine Interface Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1-ISN`.
- For **Interface Selectors**, click the **+** sign to add the port on the spine switch that connects to the ISN. Then in the **Create Spine Access Port Selector** window, provide the following:
  - For the **Name** field, specify the name for the port selector, for example `Spine1-ISN`.
  - For the **Interface IDs**, specify the switch port that connects to the ISN, for example `5/32`.
  - For the **Interface Policy Group**, choose the policy group you created in the previous step, for example `Spine1-PolGrp`.

  Then click **OK** to save the port selector.

c) Click **Submit** to save the spine interface profile.

Step 5 Configure a spine switch selector policy.

a) In the left navigation tree, browse to **Switch Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Profile**.

In the **Create Spine Profile** window, specify the following:

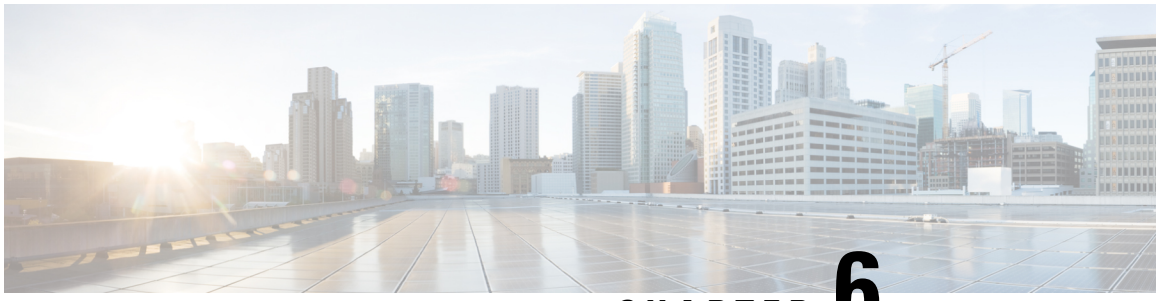
- For the **Name** field, specify the name for the profile, for example `Spine1`.
- For **Spine Selectors**, click the **+** to add the spine and provide the following:
  - For the **Name** field, specify the name for the selector, for example `Spine1`.

- For the **Blocks** field, specify the spine node, for example `201`.

c) Click **Update** to save the selector.
d) Click **Next** to proceed to the next screen.
e) Select the interface profile you have created in the previous step

   For example `Spine1-ISN`.

f) Click **Finish** to save the spine profile.



CHAPTER 6

Prerequisites: Insights

- Requirements for Insights, on page 53
- Communication Ports for Insights, on page 54
- Fabric Requirements for Insights, on page 55

Requirements for Insights

Note This section describes *additional* requirements and guidelines if you plan to enable the Insights service. Ensure that you have already satisfied the platform-level requirements described in the Prerequisites and Guidelines, on page 9 section.

- Starting with Nexus Dashboard release 3.1.1, Cisco DC App Center connectivity has been removed from Nexus Dashboard because downloading the services separately is no longer required.

  To deploy Insights, download the unified installation image from the Software Download page; individual services' installation images are no longer available from the Cisco DC App Center.

- For Nexus Dashboard Insights service, the data network must provide IP reachability to the following:
  - The in-band network of each fabric and of the APIC.
  - The DNS server.
  - For Panduit PDU integration, to the Panduit PDU server.
  - For External Kafka integration, to the External Kafka server (consumer).
  - For SysLog integration, to the SysLog server.
  - For Network-Attached Storage integration, to the Network-Attached Storage server.
  - For vCenter integration, to vCenter.
  - For AppDynamics integration, to the AppDynamics controller.
- If you are using the Insights service with NDFC fabrics or have SFLOW/NetFlow enabled, the Data network interfaces must be Layer 2 adjacent.
- You must allocate the following number of persistent IP addresses depending on your use case.

For an overview of Persistent IP functionality, see Prerequisites and Guidelines, on page 9.

For ACI Fabrics:

- Nexus Dashboard Insights without Netflow or Panduit PDU integration: 0 IP needed in data network.
- Nexus Dashboard Insights with Panduit PDU integration: 1 IP (if using IPv4) and the integration is not supported with pure IPv6 stack.
- Nexus Dashboard Insights with Netflow and Panduit PDU integration : 8 IPs (if using IPv4) and 6 IPs (if using IPv6) in data network.
- Nexus Dashboard Insights with Netflow and without Panduit PDU integration : 8 IPs (if using IPv4) and 6 IPs (if using IPv6) in data network.

For NDFC fabrics:

- 8 IPs (if using IPv4) and 6 IPs (if using IPv6) in data network.

For Standalone NX-OS switches:

- 10 IPs (if using IPv4) and 8 IPs (if using IPv6) in data network.

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*.

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

Table 16: Insights RTT Requirements

Connectivity	Maximum RTT
To switches	150 ms

Communication Ports for Insights

In addition to the ports required by the Nexus Dashboard cluster nodes (listed in a previous section), the following ports are required by the Insights service.

Note By default, Insights requires connectivity only between data interfaces of Nexus Dashboard cluster nodes and in-band IP of the switches. However, if a switch becomes unavailable, then Insights will attempt to connect to the OOB IP of the switches using the cluster nodes' management or data interface (depending on the route settings).

Table 17: 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

Fabric Requirements for Insights

Additional Prerequisites for ACI Fabrics

If you plan to use the Insights service with ACI fabrics, ensure that:

- You can on-board only 1 type of sites (ACI, NDFC, or Standalone NX-OS) within the same cluster.

  Onboarding a mix of ACI and NDFC, ACI and NX-OS, or NDFC and NX-OS within the same cluster is not supported.

- You have configured NTP settings on Cisco APIC.

  For more information, see Configure NTP in ACI Fabric Solution.

- If you plan to use the flow telemetry functions in Nexus Dashboard Insights, Telemetry Priority must be selected in the ACI fabric node control policy.

  In Cisco APIC, choose **Fabric** > **Fabric Policies** > **Policies** > **Monitoring** > **Fabric Node Controls** > **<*policy-name*>** > **Feature Selection** to select Telemetry Priority. Monitoring <*policy-name*> should be attached to **Fabric** > **Fabric Policies** > **Switches** > **Leaf/Spine Switches** > **Profiles** > .

- If you plan to use the flow telemetry functions in Nexus Dashboard Insights, Precision Time Protocol (PTP) must be enabled on Cisco APIC so that Nexus Dashboard Insights can correlate flows from multiple switches accordingly

  In Cisco APIC, choose **System** > **System Settings** > **PTP and Latency Measurement** > **Admin State** to enable PTP.

  The quality of the time synchronization via PTP depends on the accuracy of the PTP Grandmaster (GM) clock which is the source of the clock, and the accuracy and the number of PTP devices such as ACI switches and IPN devices in between.

  Although a PTP GM device is generally equipped with a GNSS/GPS source to achieve the nanosecond accuracy which is the standard requirement of PTP, microsecond accuracy is sufficient for Nexus Dashboard Insights and its flow telemetry, hence a GNSS/GPS source is typically not required.

  For a single-pod ACI fabric, you can connect your PTP GM via leaf switches. Otherwise, one of the spine switches will be elected as a GM. For a multi-pod ACI fabric, you can connect your PTP GM via leaf switches or via IPN devices. Your IPN devices should be PTP boundary clocks or PTP transparent clocks so that ACI switch nodes can synchronize their clock across pods. To maintain the same degree of accuracy across pods, it is recommended to connect your PTP GM via IPN devices.

  See section "Precision Time Protocol" in Cisco APIC System Management Configuration Guide for details about PTP connectivity options.

- You have configured in-band management as described in Cisco APIC and Static Management Access.

- If one or more DNS Domains are set under DNS Profiles, it is mandatory to set one DNS Domain as default.

  In Cisco APIC, choose **Fabric** > **Fabric Policies** > **Policies** > **Global** > **DNS Profile** > **default** > **DNS Domains** and set one as default.

  Failure to do so will result in the same switch appearing multiple times in the Nexus Dashboard Insights Flow map.

- Deploy ACI in-band network by configuring EPG using the following:
  - Tenant = `mgmt`
  - VRF = `inb`
  - BD = `inb`
  - Node Management EPG = `default/<any_epg_name>`

- Nexus Dashboard's data-network IP address and ACI fabric's in-band IP address must be in different subnets.

Additional Prerequisites for NDFC Fabrics or Standalone NX-OS Switches

If you plan to use the Insights service with NDFC fabrics or Standalone NX-OS switches, ensure that:

- You can on-board only 1 type of sites (ACI, NDFC, or Standalone NX-OS) within the same cluster.

  Onboarding a mix of ACI and NDFC, ACI and NX-OS, or NDFC and NX-OS within the same cluster is not supported.

- The data network must have IP reachability to the fabrics' in-band IP addresses.

- To enable Flow Telemetry or Traffic Analytics, Precision Time Protocol (PTP) must be configured on all nodes you want to support with Nexus Dashboard Insights.

  In both managed and monitor site mode, you must ensure PTP is correctly configured on all nodes in the site. You can enable PTP in NDFC easy site setup's **Advanced** tab by checking the **Enable Precision Time Protocol (PTP)** option.

  The PTP Grandmaster Clock should be provided by a device that is external to the network site. Using Cisco Nexus 9000 series switches as PTP Grandmaster is not supported.



Note N9k-C93180YC-FX3 switch in the fabric can be used as a PTP GM.

  The quality of the time synchronization via PTP depends on the accuracy of the PTP Grandmaster (GM) clock which is the source of the clock, and the accuracy and the number of PTP devices along the network path. Although a PTP GM device is generally equipped with a GNSS/GPS source to achieve the nanosecond accuracy which is the standard requirement of PTP, microsecond accuracy is sufficient for Nexus Dashboard Insights and its flow telemetry, hence a GNSS/GPS source is typically not required.

  For details about configuring Precision Time Protocol on Nexus switches using NDFC, *Cisco NDFC LAN Fabric Configuration Guide*.

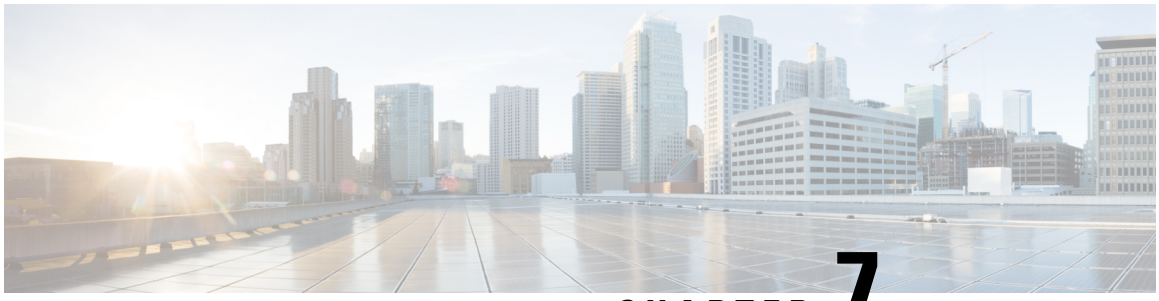
  For details about manually configuring Precision Time Protocol on Nexus switches, see *Cisco Nexus 9000 Series NX-OS System Management Configuration Guide*.



PART II

Deploying the Cluster

- Deploying as Physical Appliance, on page 61
- Deploying in VMware ESX, on page 81
- Deploying in Linux KVM, on page 119
- Deploying in Amazon Web Services, on page 137
- Deploying in Microsoft Azure, on page 151
- Onboarding Fabrics, on page 165



CHAPTER 7

Deploying as Physical Appliance

- Prerequisites and Guidelines, on page 61
- Physical Node Cabling, on page 64
- Deploying Nexus Dashboard as Physical Appliance, on page 66

Prerequisites and Guidelines

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

- Review and complete the general and service-specific prerequisites described in Prerequisites: Nexus Dashboard, on page 9.
- 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* specific to the model of server you have.

  The physical appliance form factor is supported on the UCS-C220-M5 (`SE-NODE-G2`) and UCS-C225-M6 (`ND-NODE-L4`) original Cisco Nexus Dashboard platform hardware only.

Note A known issue exists if you are performing a fresh virtual media installation of the 3.1.1k software on a cluster with UCS-C225-M6 (ND-NODE-L4) nodes and ACI sites, where onboarding the site to NDI or NDO will fail. The workaround for this issue is to perform a fresh installation of the **3.1.1l** version of the software instead. Note that upgrading from release 3.1.1k to 3.1.1l will not resolve the issue; you must perform a fresh installation of the 3.1.1l software to resolve the issue.

The following table lists the PIDs and specifications of the physical appliance servers:

Table 18: 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 19: 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)
	• One of the following:
	  • 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
```

- 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 2.3.2 image pre-installed, but you want to deploy Release 3.1.1 instead, we recommend:

  1. First, bring up the Release 2.3.2 cluster, as described in the deployment guide for that release.
  2. Then upgrade to Release 3.1.1, as described in Upgrading Existing ND Cluster to This Release, on page 173.

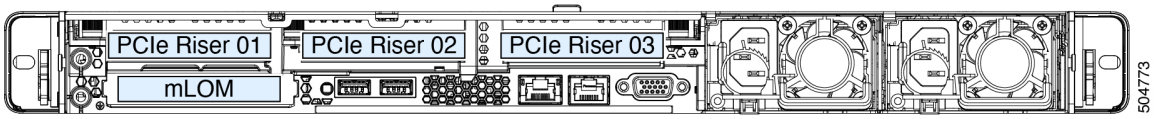
Note For brand new deployments, you can also choose to simply re-image the nodes with the latest version of the Cisco Nexus Dashboard (for example, if the hardware came with an image which does not support a direct upgrade to this release through the GUI workflow) before returning to this document for 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 secondary nodes can be added for horizontal scaling if required by the enter and number of services you deploy. For the maximum number of `secondary` and `standby` nodes in a single cluster, see the *Release Notes* for your release.

Physical Node Cabling

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 9: 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.

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 Nexus Dashboard cluster.

Before you begin

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

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, as well as network connectivity between the nodes is established on the data network.

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

```
C220-WZP23150D4C# connect host
CISCO Serial Over LAN:
Press Ctrl+x to Exit the session
```

After connecting to the host, press **Enter** to continue.

b) After you see the Nexus Dashboard setup utility prompt, press **Enter**.

```
Starting Nexus Dashboard setup utility
Welcome to Nexus Dashboard 3.1.1k
Press Enter to manually bootstrap your first master node...
```

c) 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:
```

d) 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.

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, enter the capital letter `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 of the first node, the initial setup configures the networking and brings up the UI, which you will use to add two and configure other nodes and complete the cluster deployment.

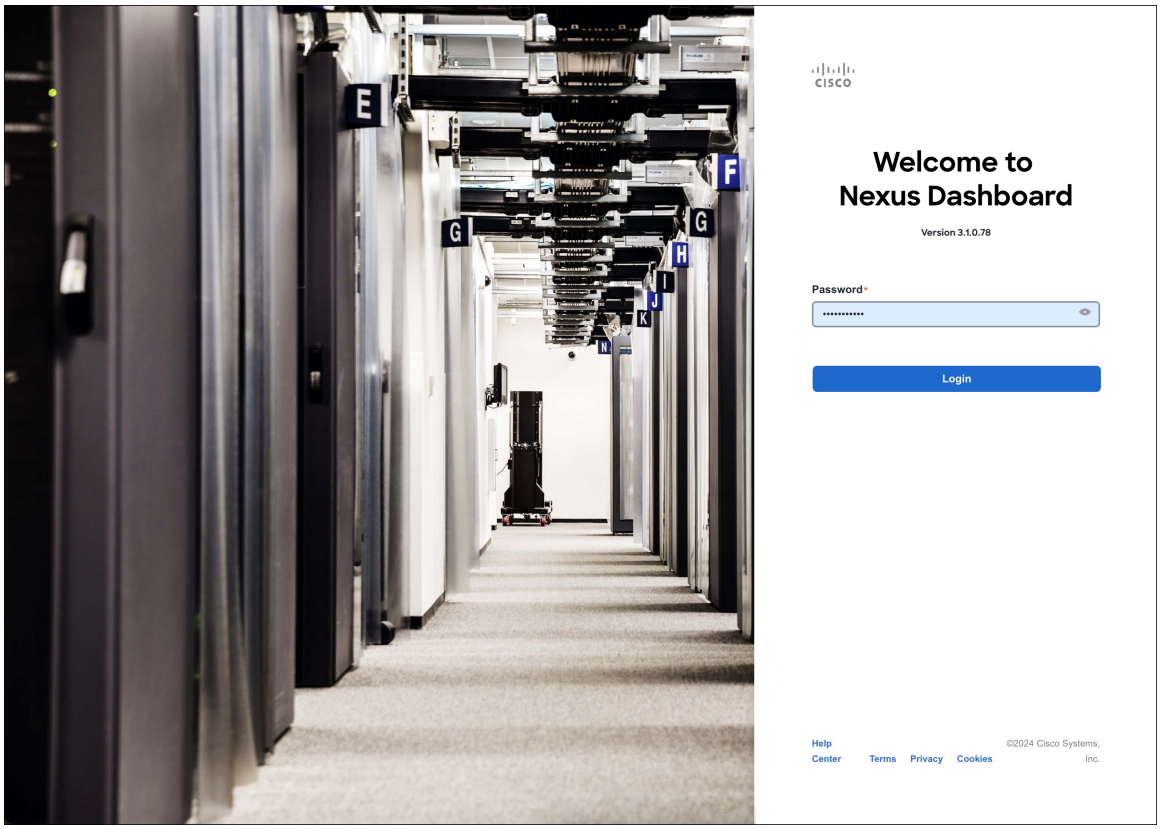
```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

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.

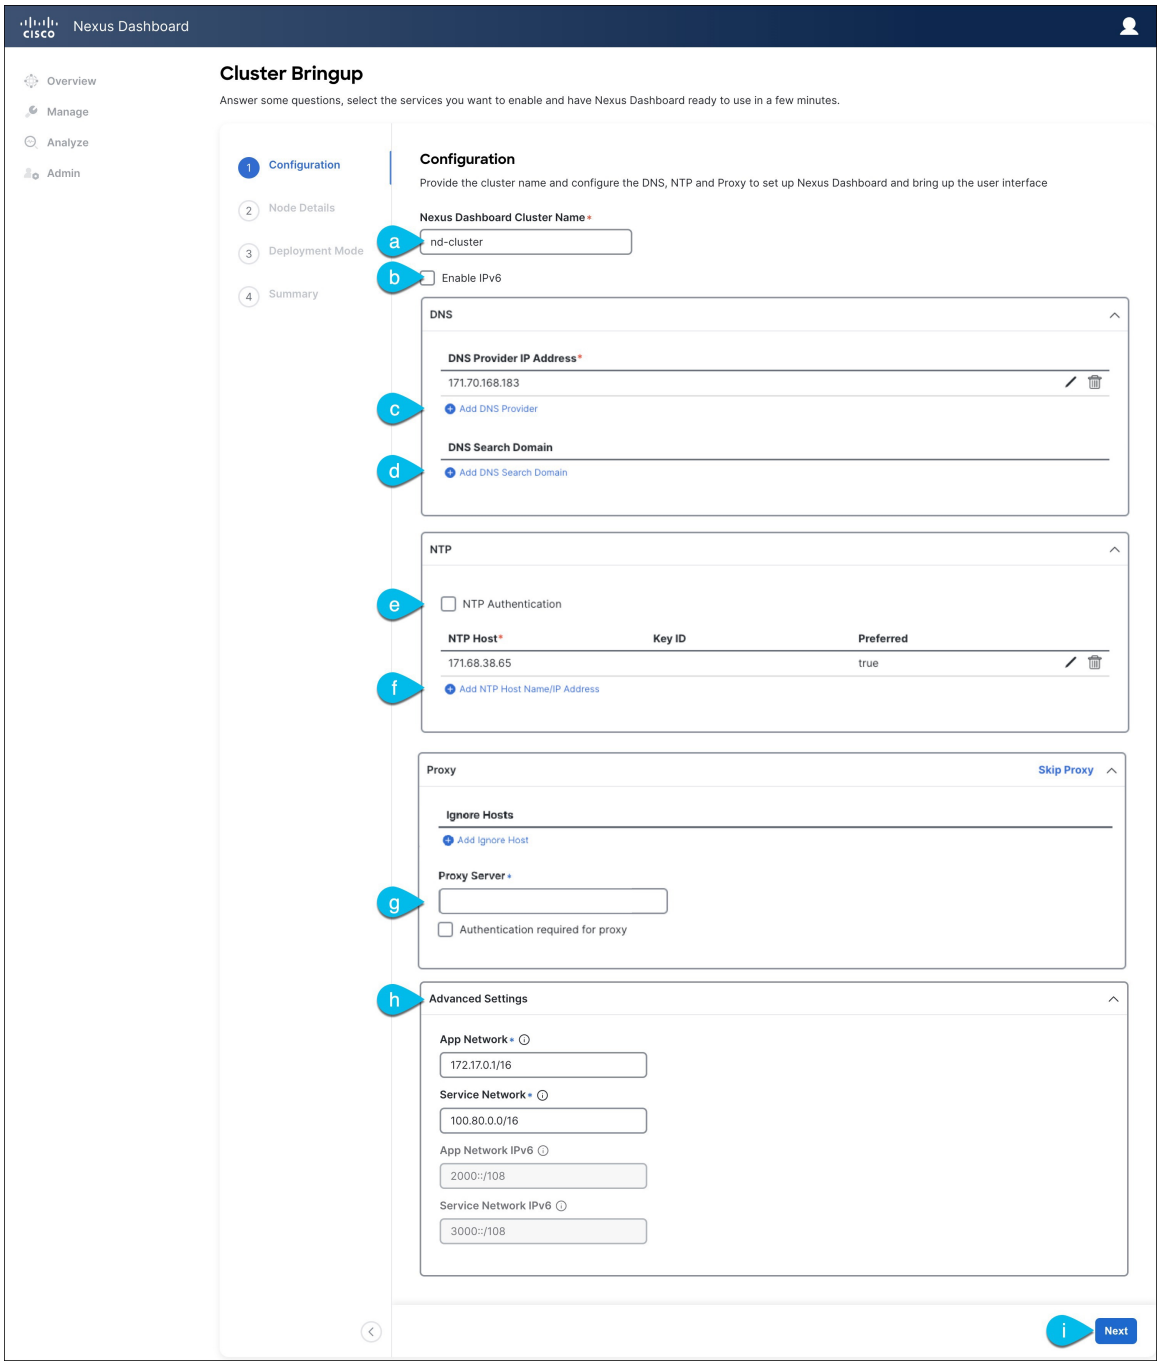
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

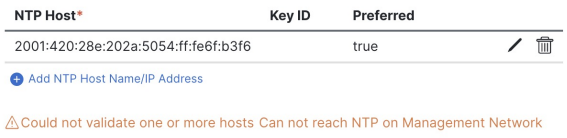
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

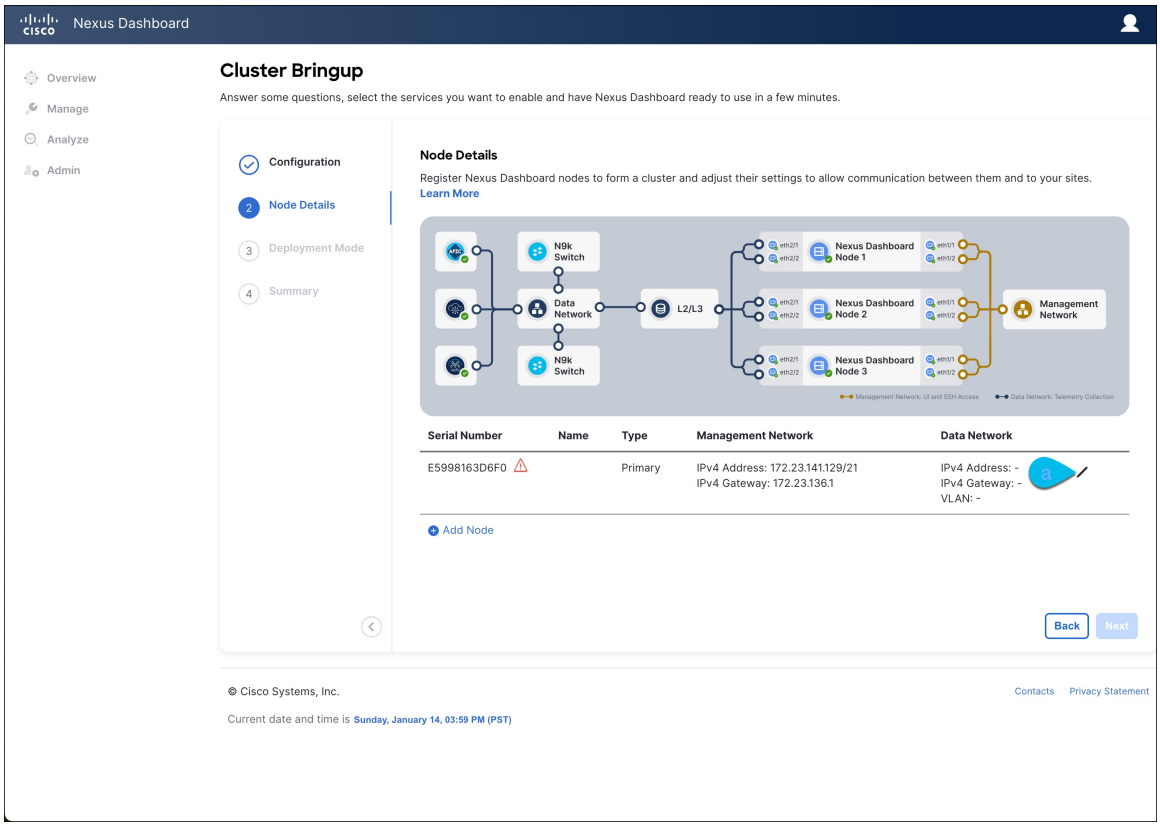
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

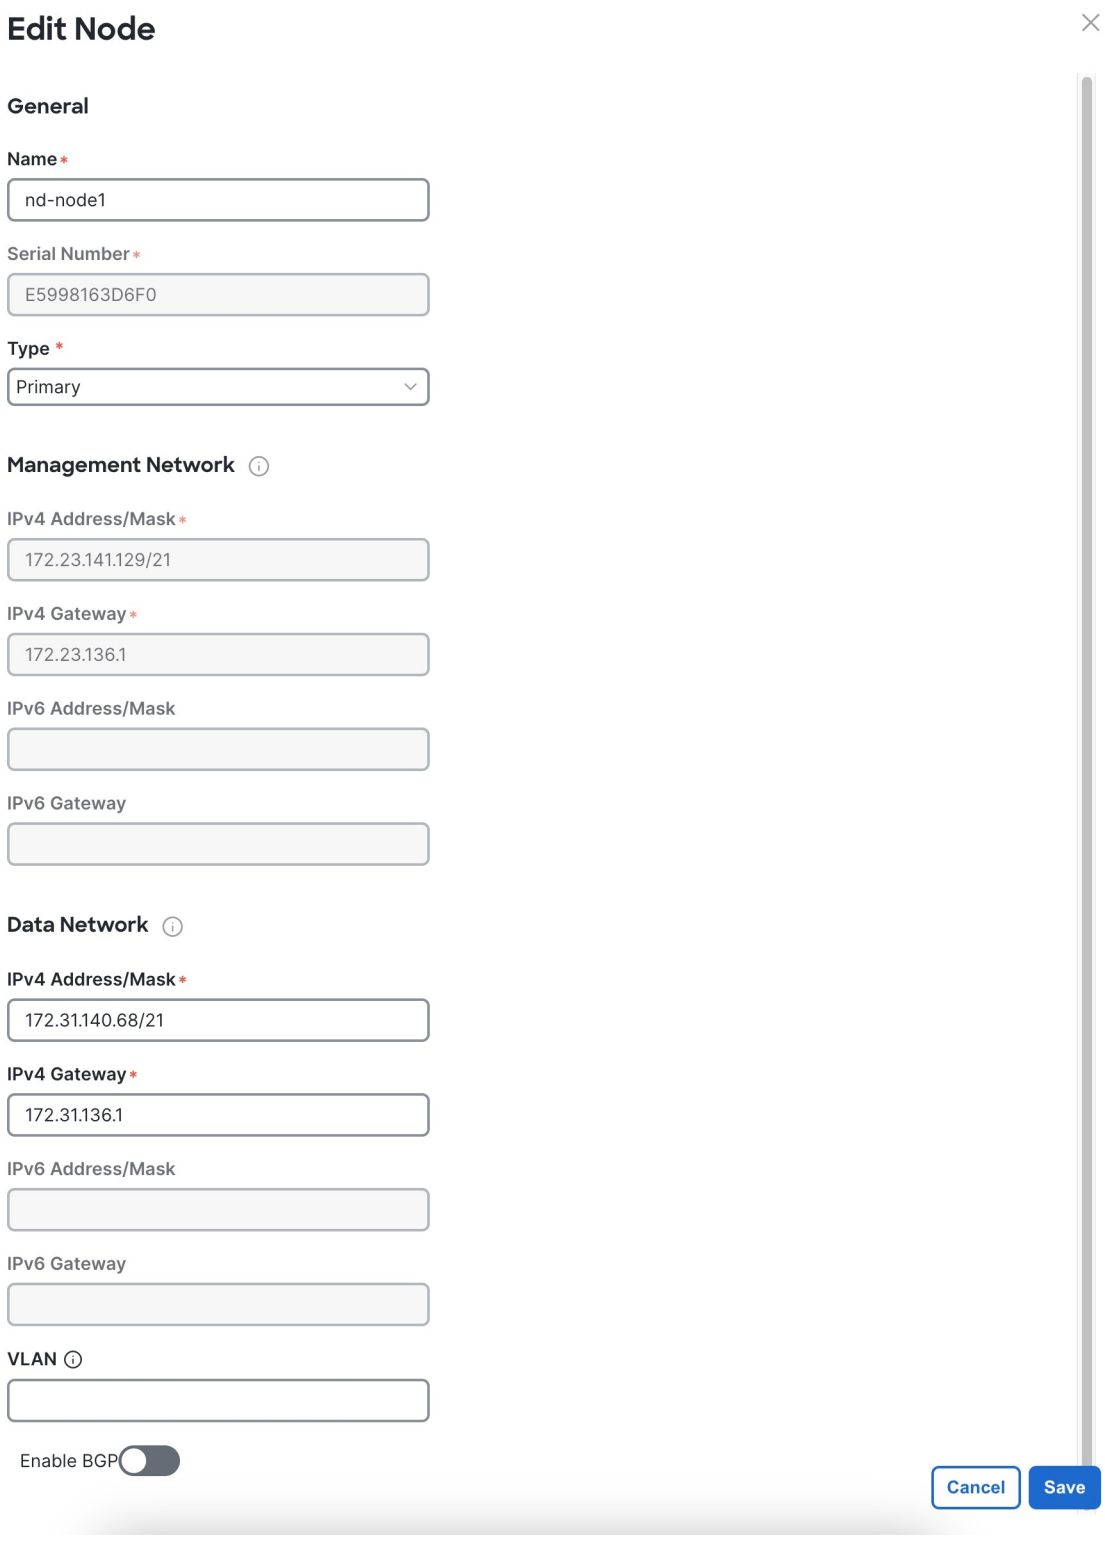
  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) Click **Next** to continue.

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

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


Nexus Dashboard
Overview
Manage
Analyze
Admin
Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.
Configuration
2 Node Details
3 Deployment Mode
4 Summary
Node Details
Register Nexus Dashboard nodes to form a cluster and adjust their settings to allow communication between them and to your sites.
Learn More
N9k Switch
Data Network
N9k Switch
L2/L3
Nexus Dashboard Node 1
Nexus Dashboard Node 2
Nexus Dashboard Node 3
Management Network
Serial Number
Name
Type
Management Network
Data Network
E5998163D6F0
Primary
IPv4 Address: 172.23.141.129/21
IPv4 Gateway: 172.23.136.1
IPv4 Address: -
IPv4 Gateway: -
VLAN: -
Add Node
Back
Next
© Cisco Systems, Inc.
Contacts
Privacy Statement
Current date and time is Sunday, January 14, 03:59 PM (PST)



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

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated but you must provide other information.

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

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

c) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

e) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

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

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

f) Click **Save** 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.

Note
For the **Username** field for the second node, enter the admin user ID (the ID 1 of the local user management).

b) Click **Validate** to verify connectivity to the node.

The node's **Serial Number** is automatically populated after CIMC connectivity is validated.

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

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

d) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

e) In the **Management Network** area, provide the node's **Management Network** information.

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

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

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

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

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

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

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

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

i) Repeat this step for the final (third) primary node of the cluster.

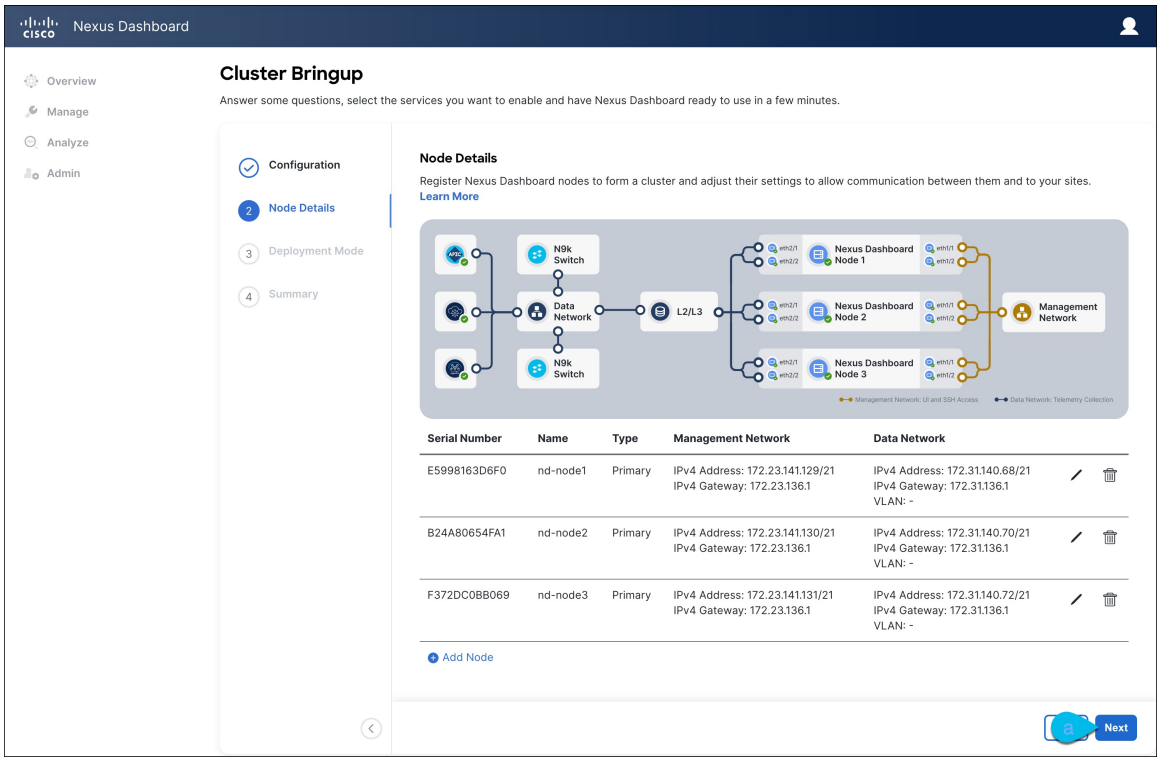
Step 7 (Optional) Repeat the previous step to provide information about any additional secondary or standby nodes.

Note

In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.

You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 8 In the **Node Details** page, verify the provided information and click **Next** to continue.



Step 9 Choose the **Deployment Mode** for the cluster.

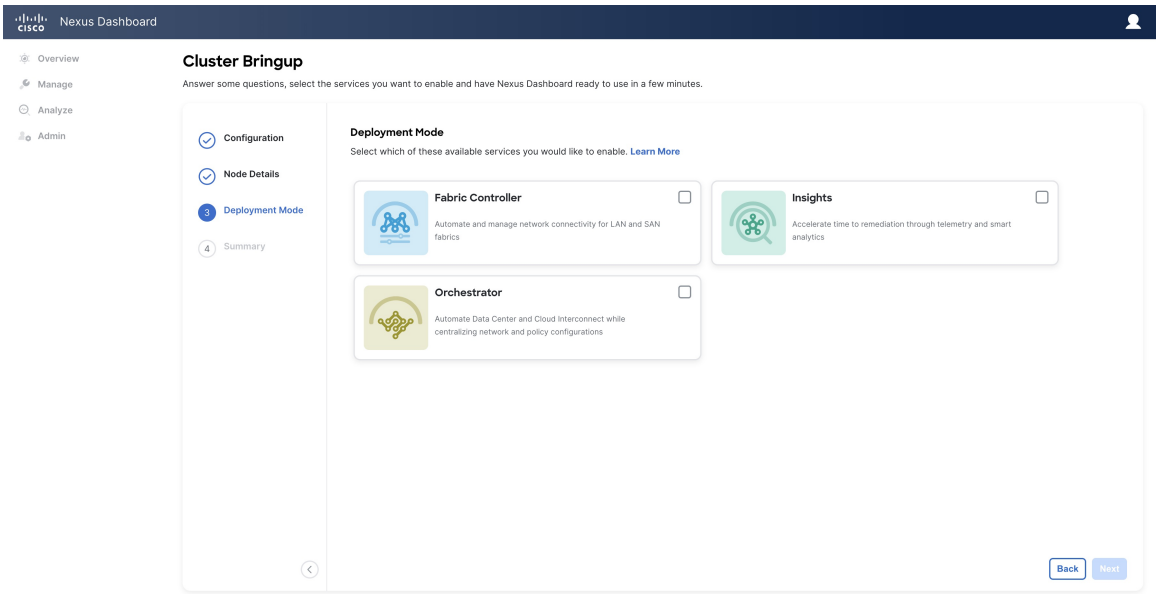
a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

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

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

- Nexus Dashboard Fabric Controller prerequisites
- Nexus Dashboard Orchestrator prerequisites
- Nexus Dashboard Insights prerequisites



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

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section and the service-specific requirements chapters.

c) Click **Next** to proceed.

Step 10 In the **Summary** screen, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

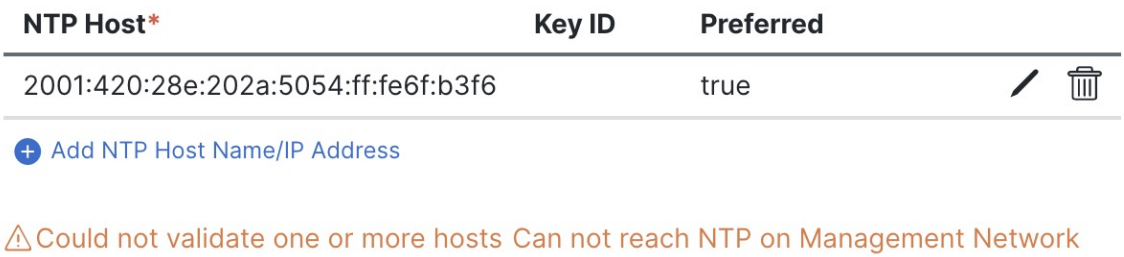
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

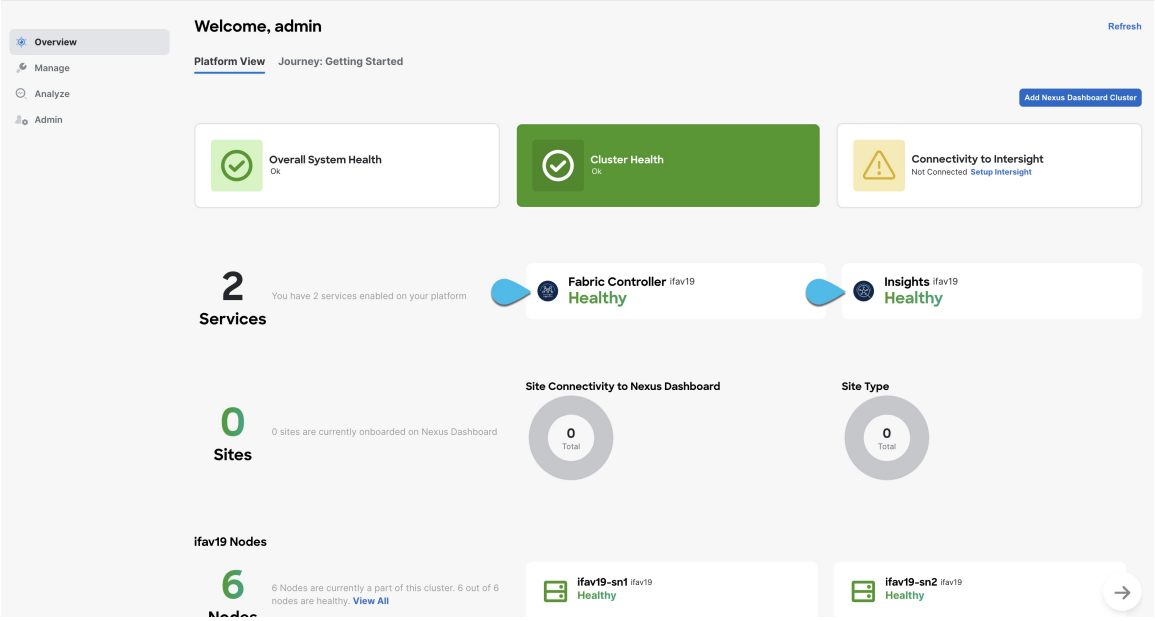
Step 11 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

Note
In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

```
deploy base system services
```

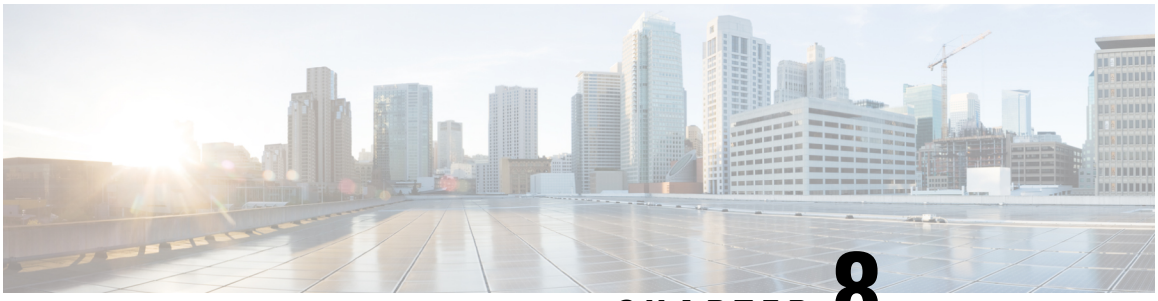
This is due to an issue with `etcd` on the node after a reboot of the pND (Physical Nexus Dashboard) cluster.

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

Step 12 After you have deployed your Nexus Dashboard and services, you can configure each service as described in its configuration and operations articles.

- For Fabric Controller, see the *NDFC persona configuration* white paper and documentation library.
- For Orchestrator, see the documentation page.

- For Insights, see the documentation library.



CHAPTER 8

Deploying in VMware ESX

- Prerequisites and Guidelines, on page 81
- Deploying Nexus Dashboard Using VMware vCenter, on page 84
- Deploying Nexus Dashboard Directly in VMware ESXi, on page 103

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

- Review and complete the general prerequisites described in Prerequisites: Nexus Dashboard, on page 9.

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

- Review and complete any additional prerequisites described in the *Release Notes* for the services you plan to deploy.
- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.
- 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 secondary nodes. If you plan to add secondary nodes to your NDFC cluster, you can deploy all nodes (the initial 3-node cluster and the additional secondary nodes) using the OVA-App profile. Detailed scale information is available in the *Verified Scalability Guide for Cisco Nexus Dashboard Fabric Controller* for your release.

Ensure you have enough system resources:

Table 20: Deployment Requirements

Data Node Requirements	App Node Requirements

• VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0.2
• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0.2 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, 8.0.2
	• VMware vCenter 7.0.1, 7.0.2, 7.0.3, 8.0.2 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 vMotion is not supported for Nexus Dashboard cluster nodes.

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

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

- Deploying via content library is not supported.
- VMware snapshots are only supported for Nexus Dashboard VMs that are powered off and must be done for all Nexus Dashboard VMs belonging to the same cluster.

  Snapshots of powered on VMs are not supported.
- 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 84.

  If you want to deploy directly in ESXi, following the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 103.



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 tool.

Deploying Nexus Dashboard Using VMware vCenter

This section describes how to deploy Cisco Nexus Dashboard cluster using VMware vCenter. If you prefer to deploy directly in ESXi, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 103 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 81.

Procedure

Step 1 Obtain the Cisco Nexus Dashboard OVA image.

a) Browse to the Software Download page.

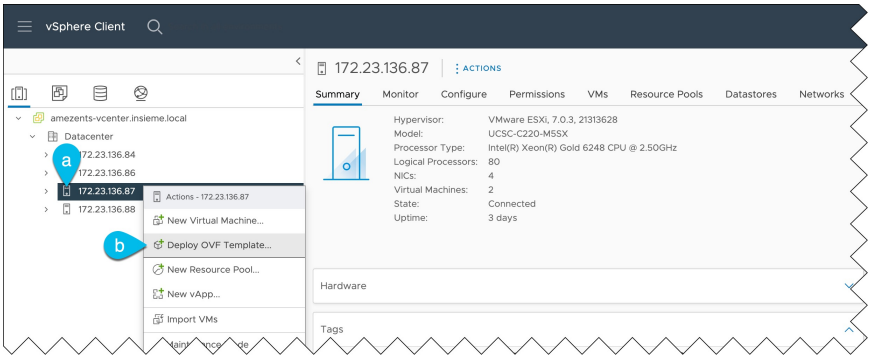
https://software.cisco.com/download/home/286327743/type/286328258/

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.

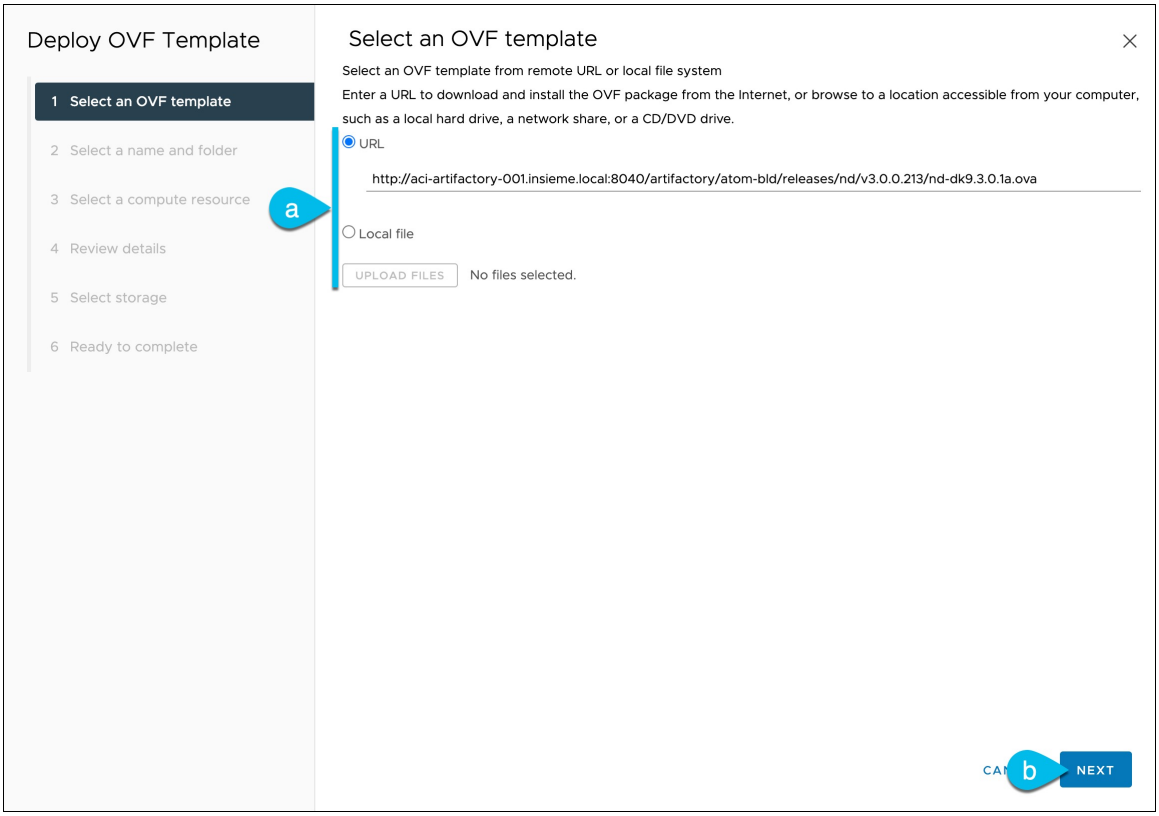
Step 3 Start the new VM deployment.



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.



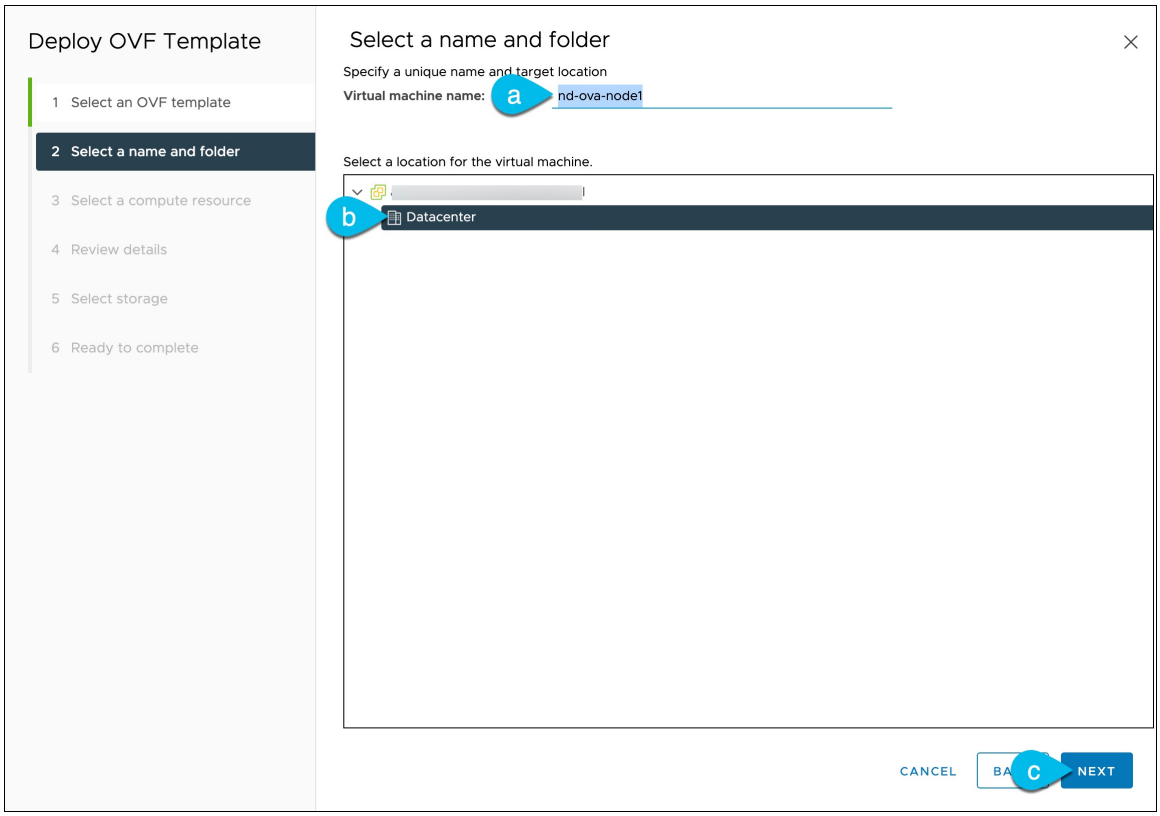
a) Provide the location of the image.

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.



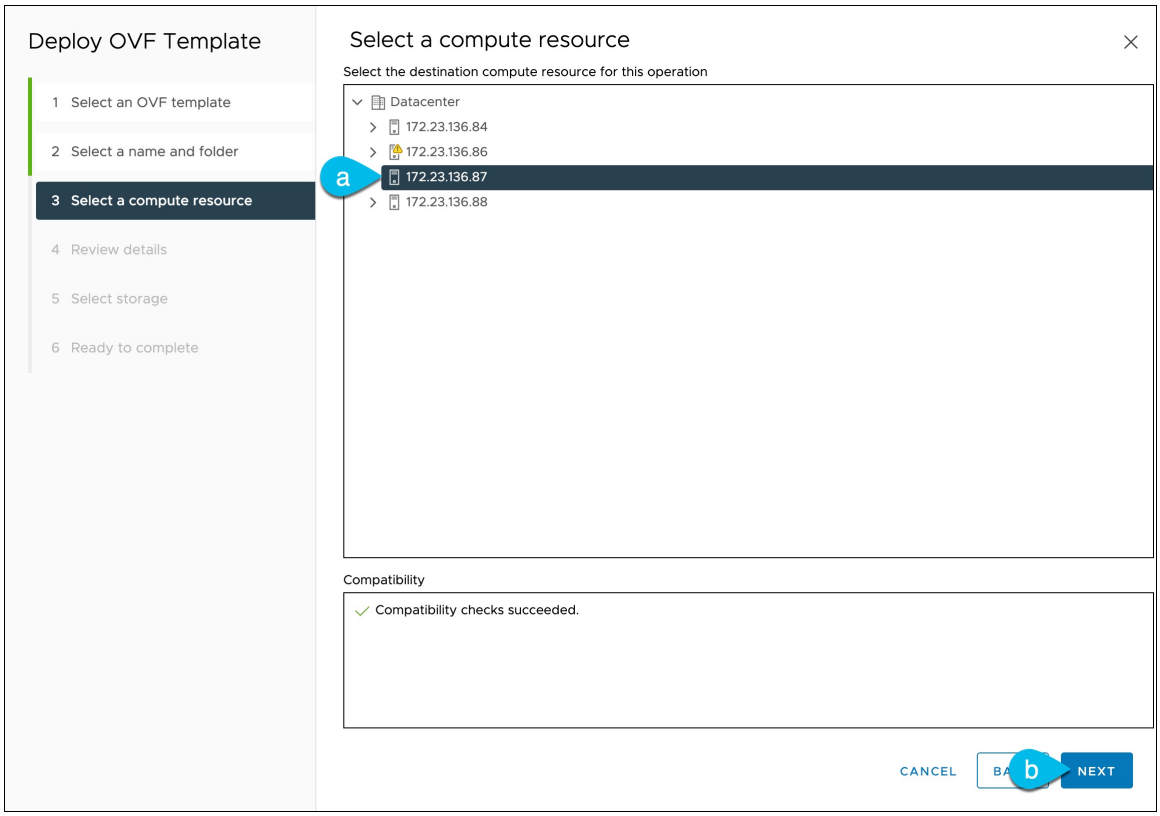
a) Provide the name for the virtual machine.

For example, `nd-ova-node1`.

b) Select the location for the virtual machine.

c) Click **Next** to continue

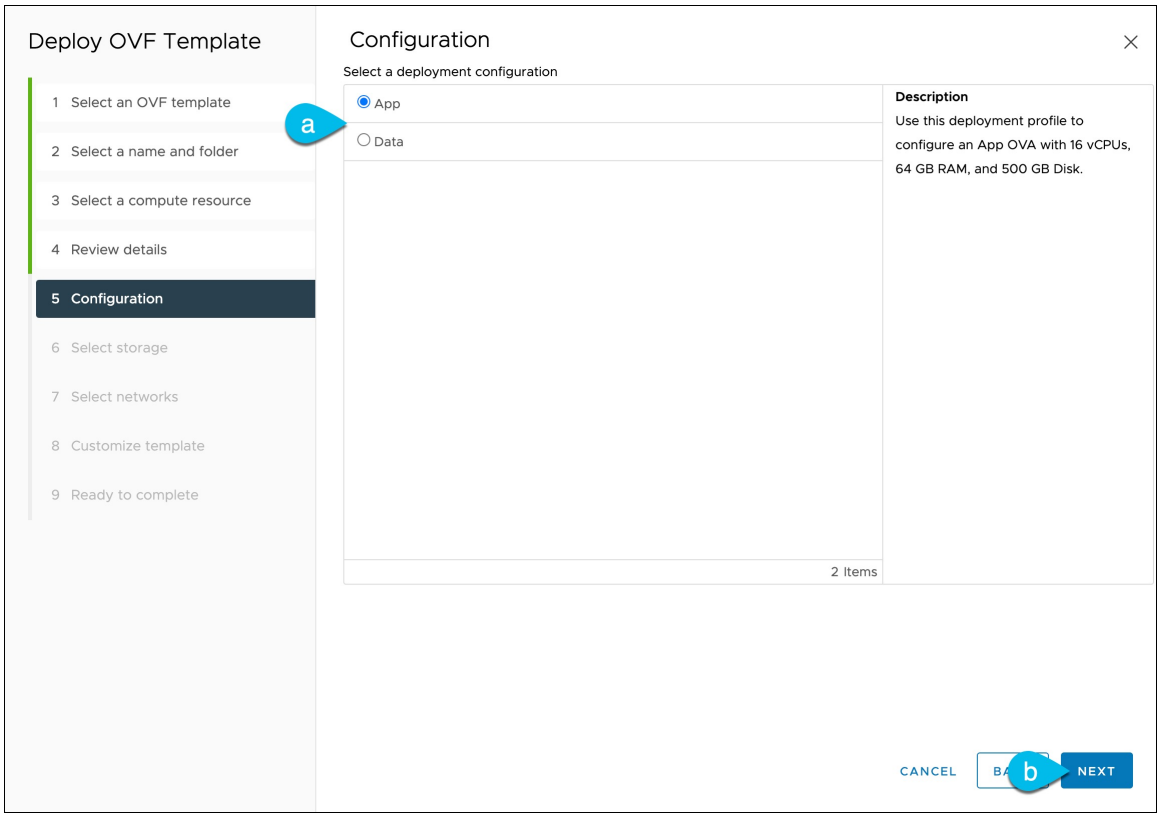
Step 6 In the **Select a compute resource** screen, select the ESX host.



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.

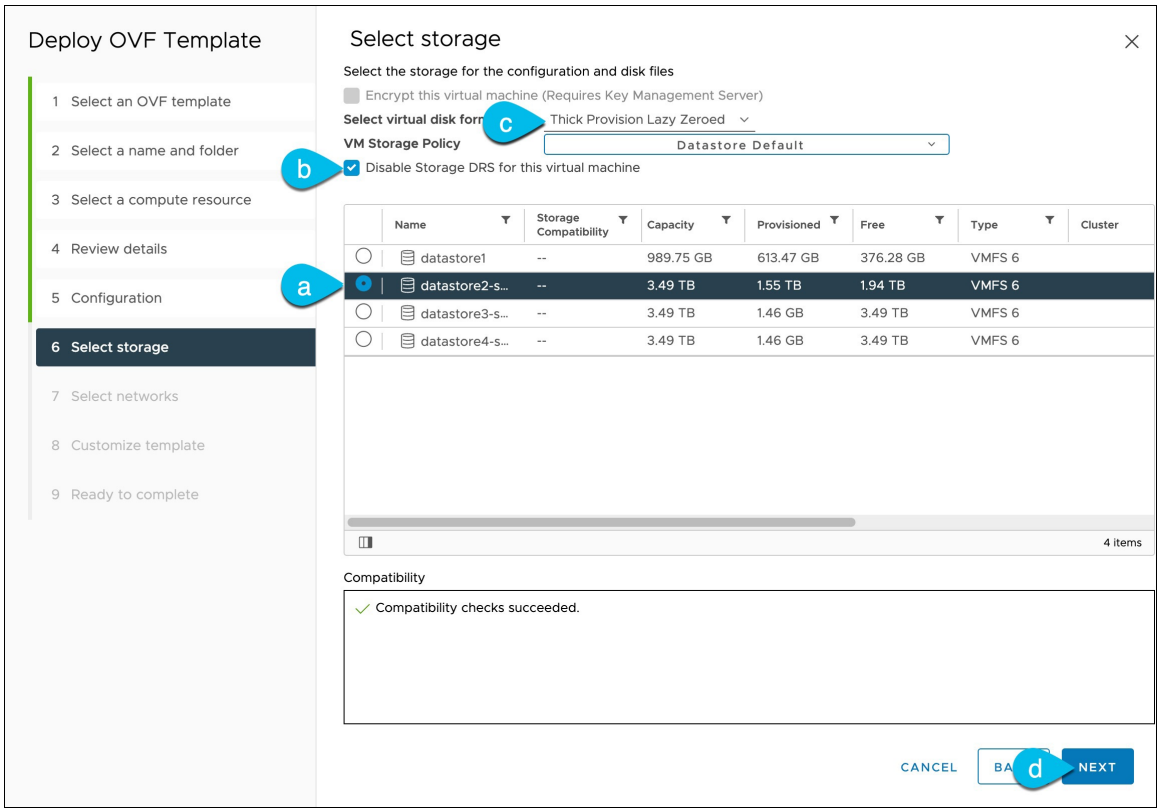


a) Select either `App` or `Data` node profile based on your use case requirements.

For more information about the node profiles, see Prerequisites and Guidelines, on page 81.

b) Click **Next** to continue

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



a) Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

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

Nexus Dashboard does not support VMware DRS.

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

d) Click **Next** to continue

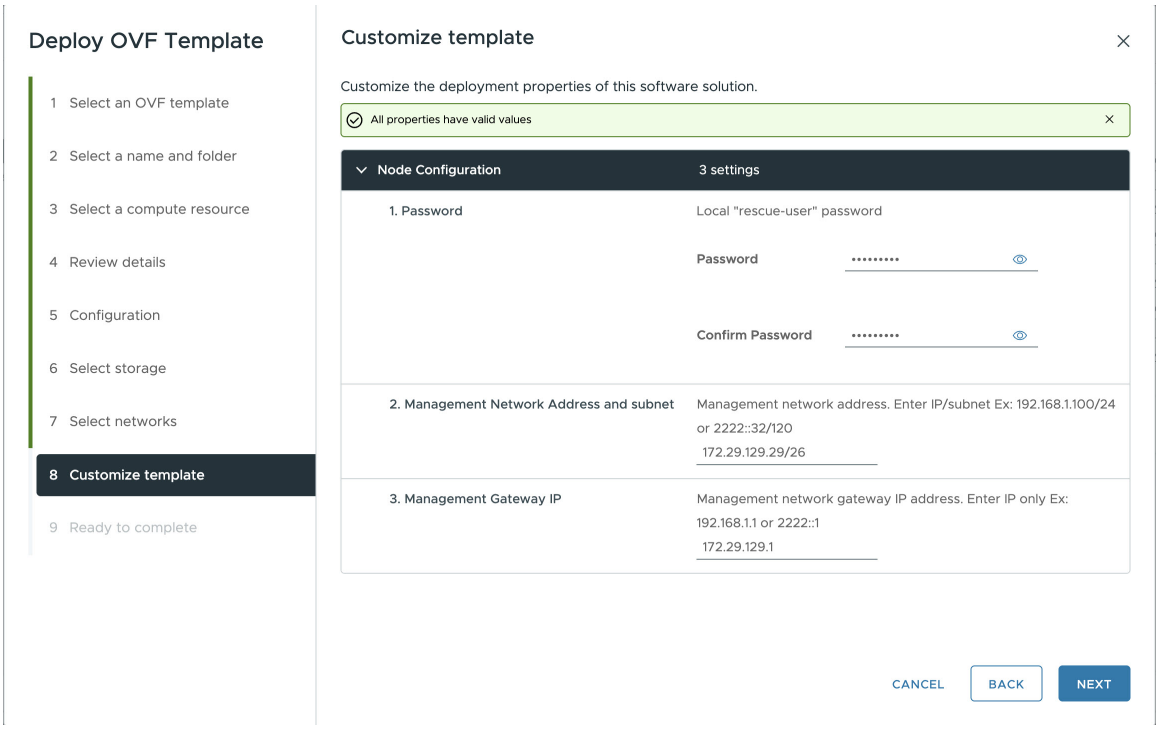
Step 10 In the **Select networks** screen, choose the VM network for the Nexus Dashboard's Management and Data networks and click **Next** to continue.

There are two networks required by the Nexus Dashboard cluster:

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

For more information about these networks, see Prerequisites and Guidelines, on page 9 in the "Deployment Overview and Requirements" chapter.

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



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

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

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

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

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

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

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

c) Provide the **Management Network** IP address and netmask.
d) Provide the **Management Network** IP gateway.
e) Click **Next** to continue.

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

Step 13 Repeat previous steps to deploy the additional nodes.

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

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

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

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

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

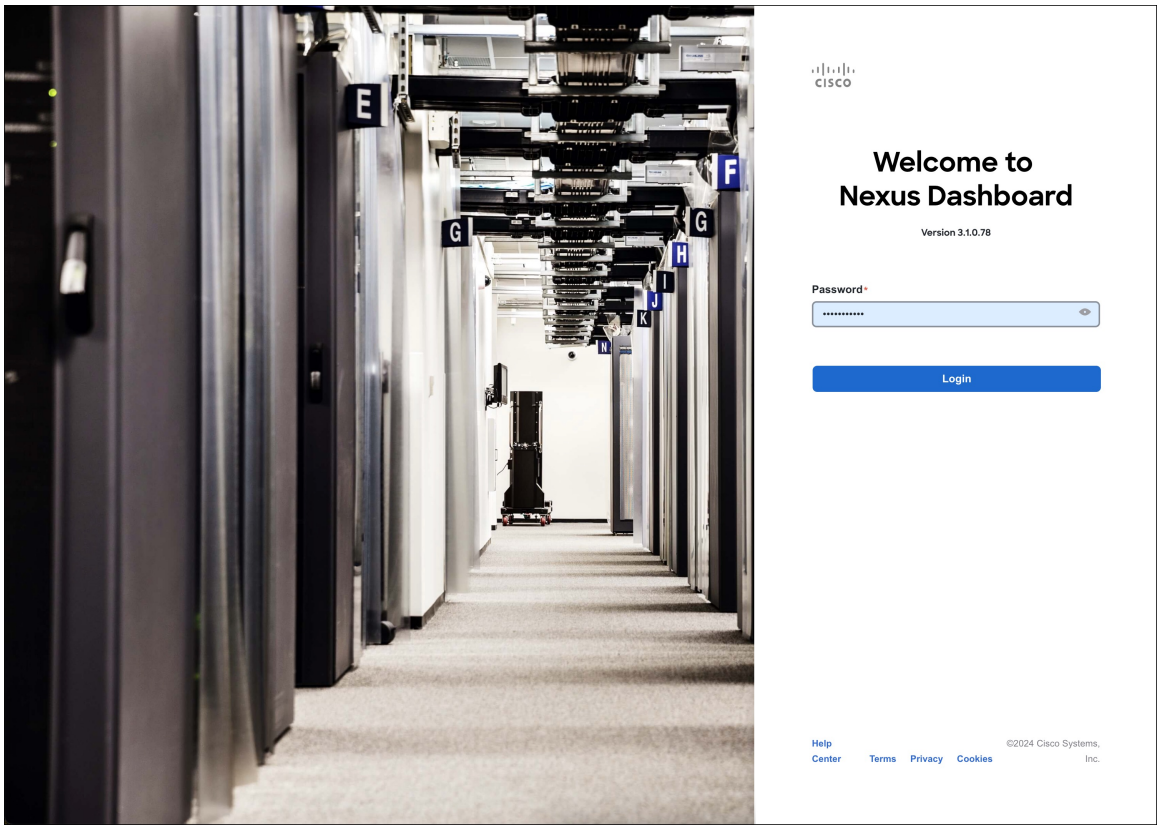
To disable time synchronization:

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

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

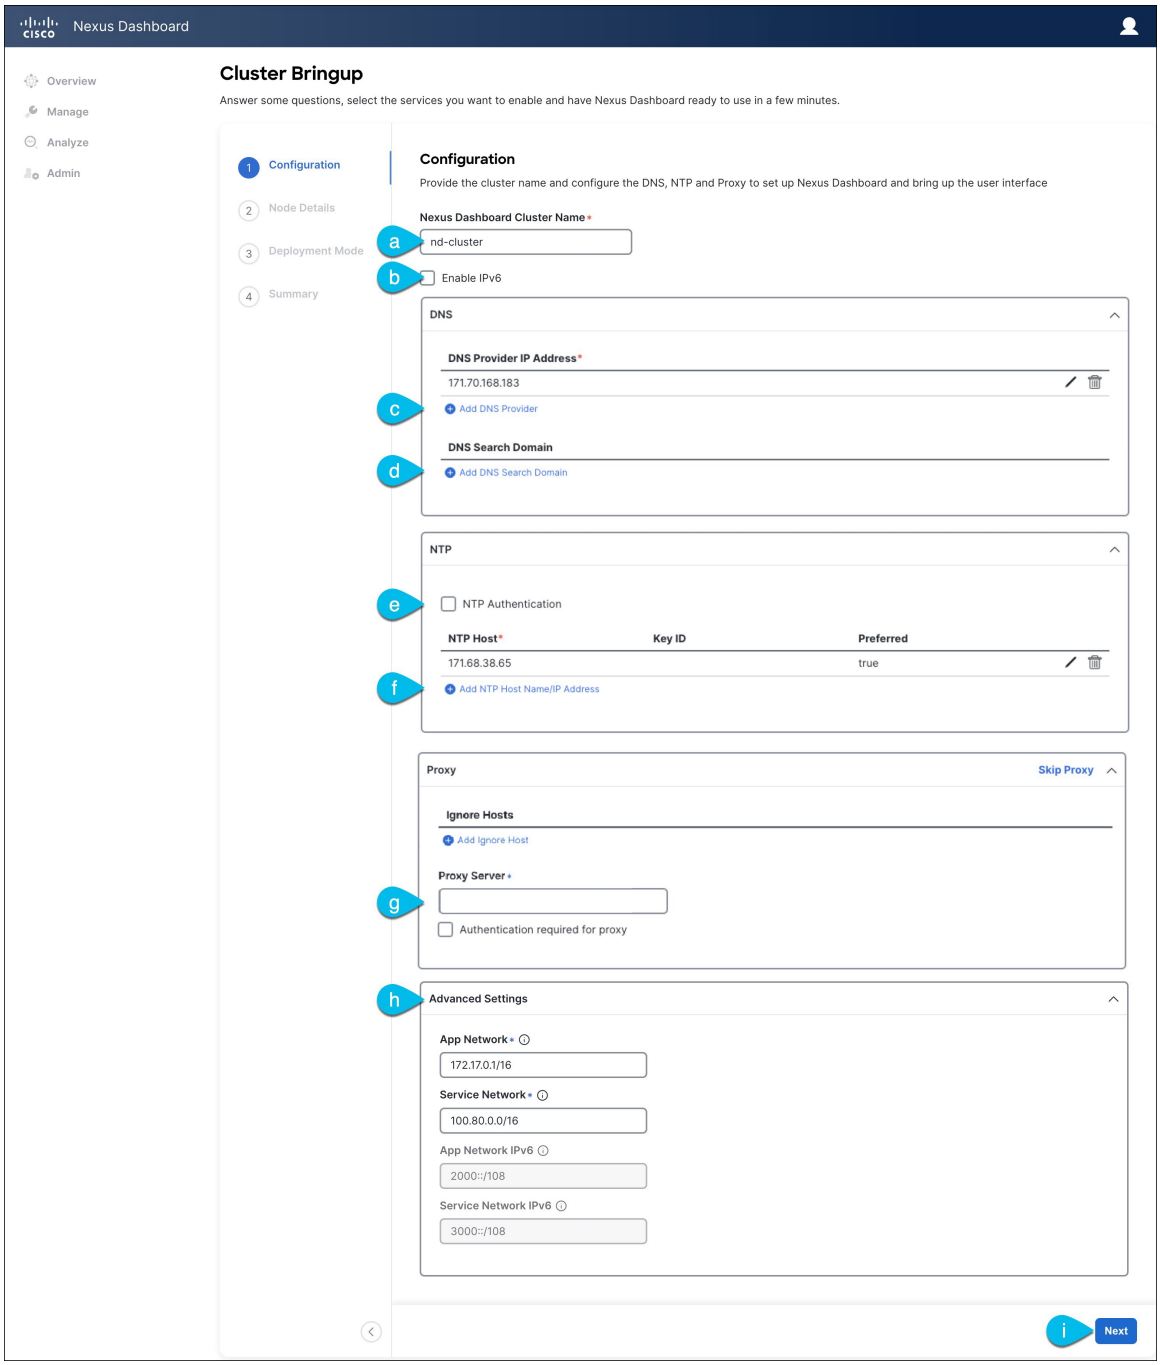
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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



Step 17 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

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

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

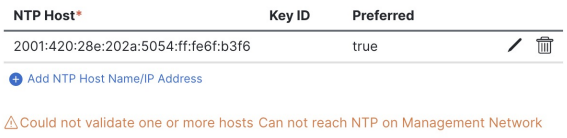
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

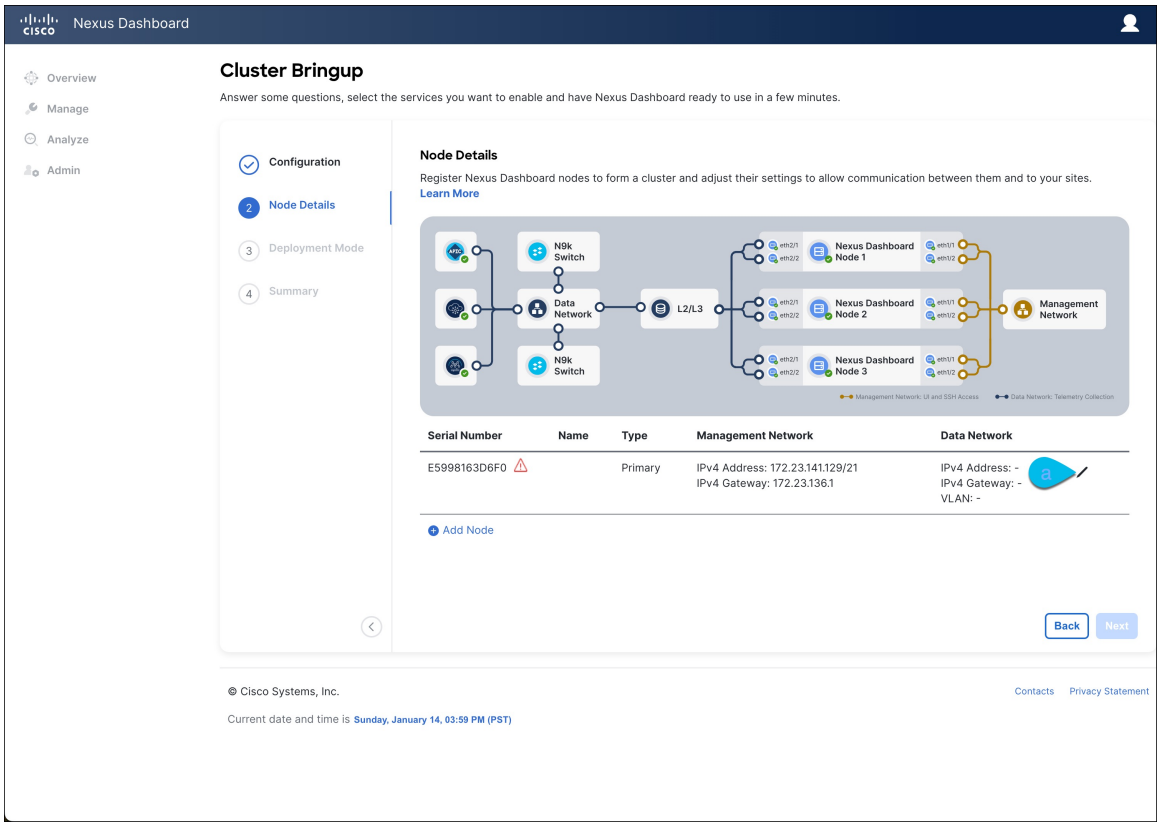
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

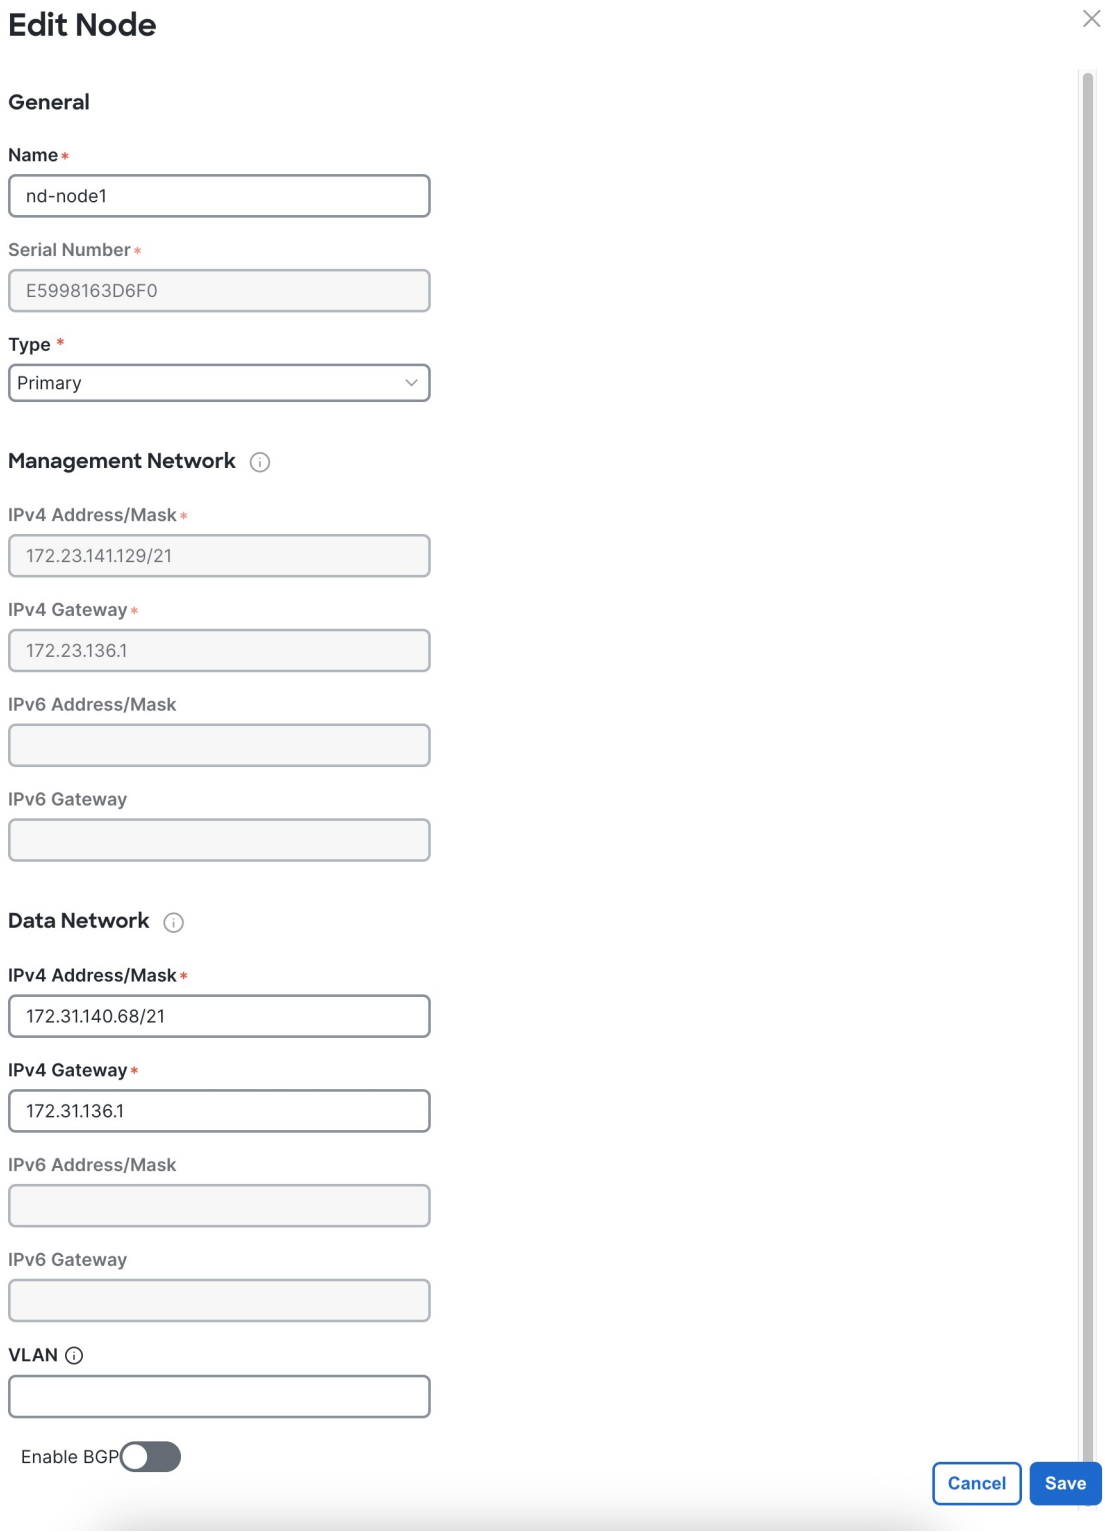
  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) Click **Next** to continue.

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

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


Nexus Dashboard
Overview
Manage
Analyze
Admin
Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.
Configuration
2 Node Details
3 Deployment Mode
4 Summary
Node Details
Register Nexus Dashboard nodes to form a cluster and adjust their settings to allow communication between them and to your sites.
Learn More
N9k Switch
Data Network
N9k Switch
L2/L3
Nexus Dashboard Node 1
Nexus Dashboard Node 2
Nexus Dashboard Node 3
Management Network
Serial Number
Name
Type
Management Network
Data Network
E5998163D6F0
Primary
IPv4 Address: 172.23.141.129/21
IPv4 Gateway: 172.23.136.1
IPv4 Address: -
IPv4 Gateway: -
VLAN: -
Add Node
Back
Next
© Cisco Systems, Inc.
Contacts
Privacy Statement
Current date and time is Sunday, January 14, 03:59 PM (PST)



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

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated but you must provide other information.

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

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

c) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

e) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

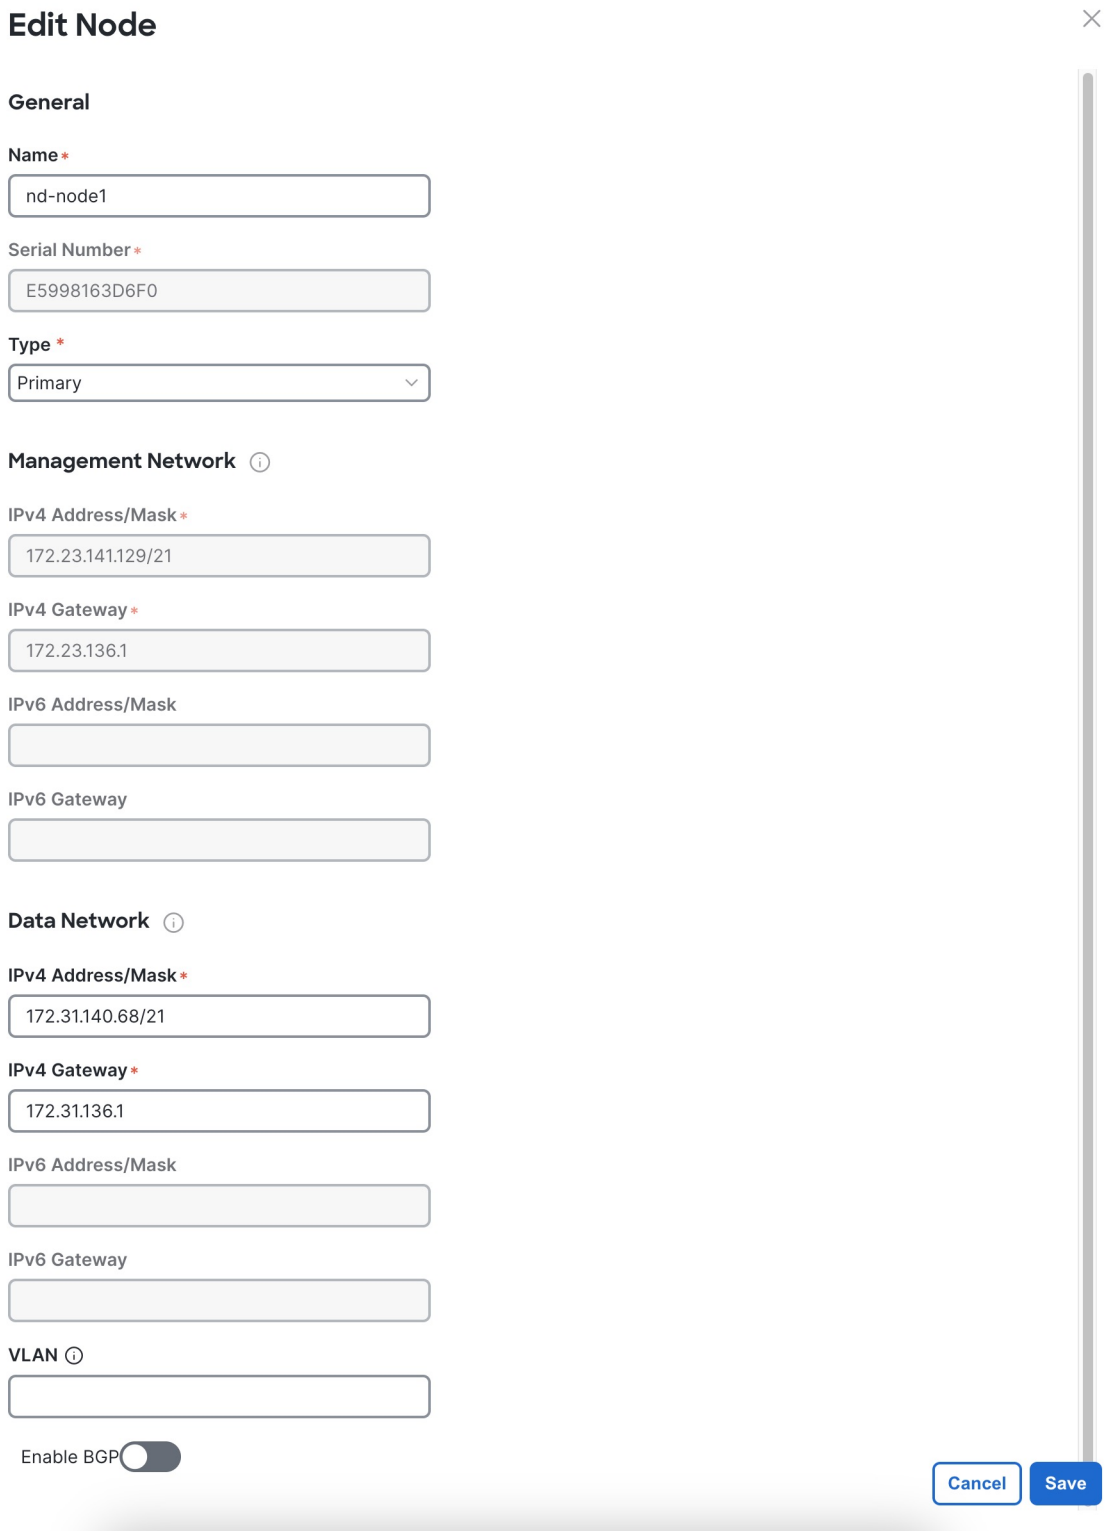
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

Step 19 In the **Node Details** screen, click **Add Node** to add the second node to the cluster.

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



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

You defined the management network information and the password during the initial node configuration steps.

b) Click **Validate** to verify connectivity to the node.

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

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

d) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

f) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

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

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

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

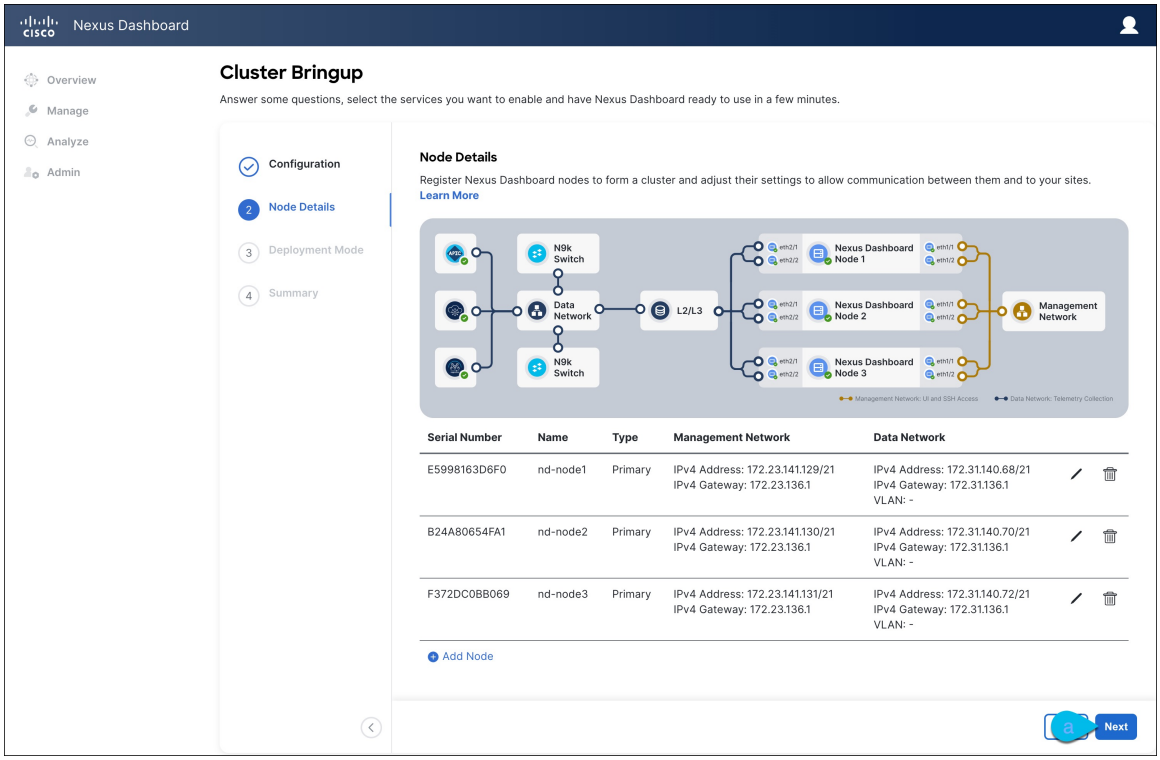
h) Repeat this step for the final (third) primary node of the cluster.

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

Note
In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.

You can choose to add the standby nodes now or at a later time after the cluster is deployed.

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



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

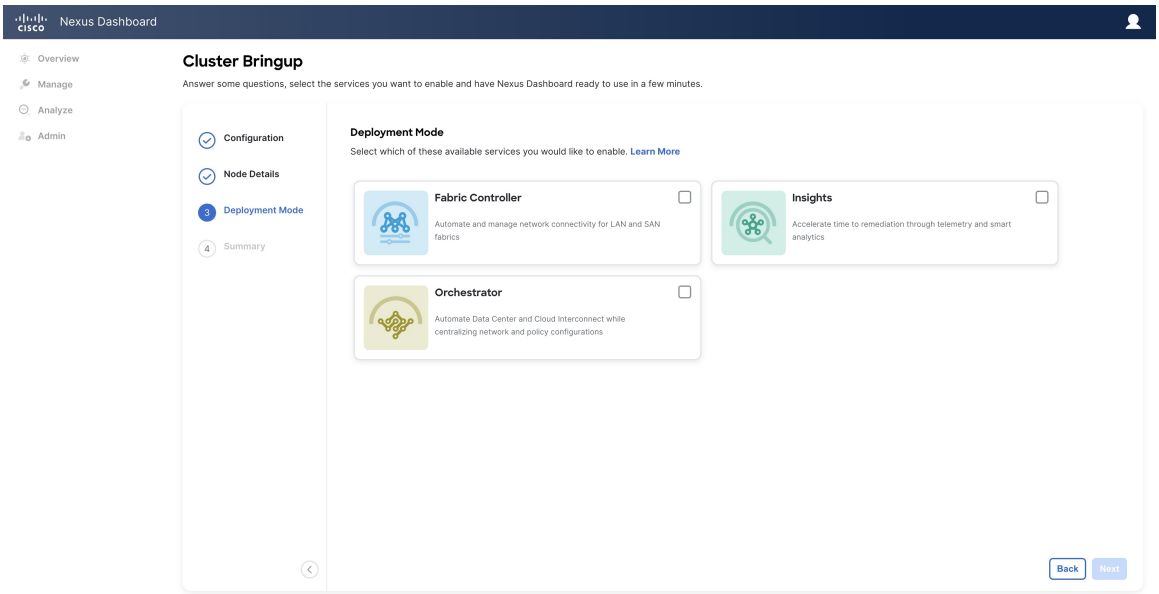
a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

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

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

- Nexus Dashboard Fabric Controller prerequisites
- Nexus Dashboard Orchestrator prerequisites
- Nexus Dashboard Insights prerequisites



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

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section and the service-specific requirements chapters.

c) Click **Next** to proceed.

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

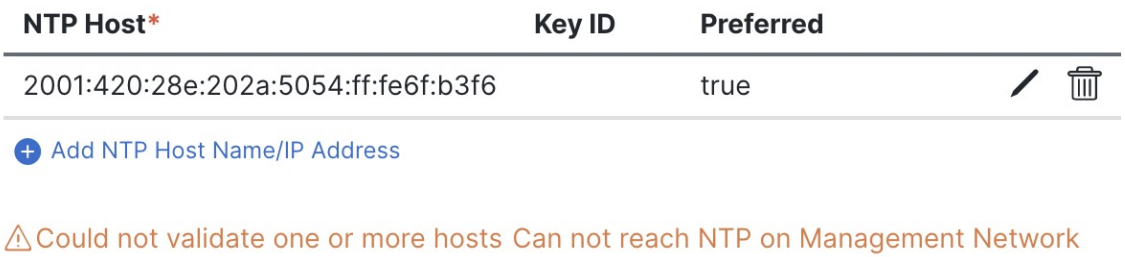
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

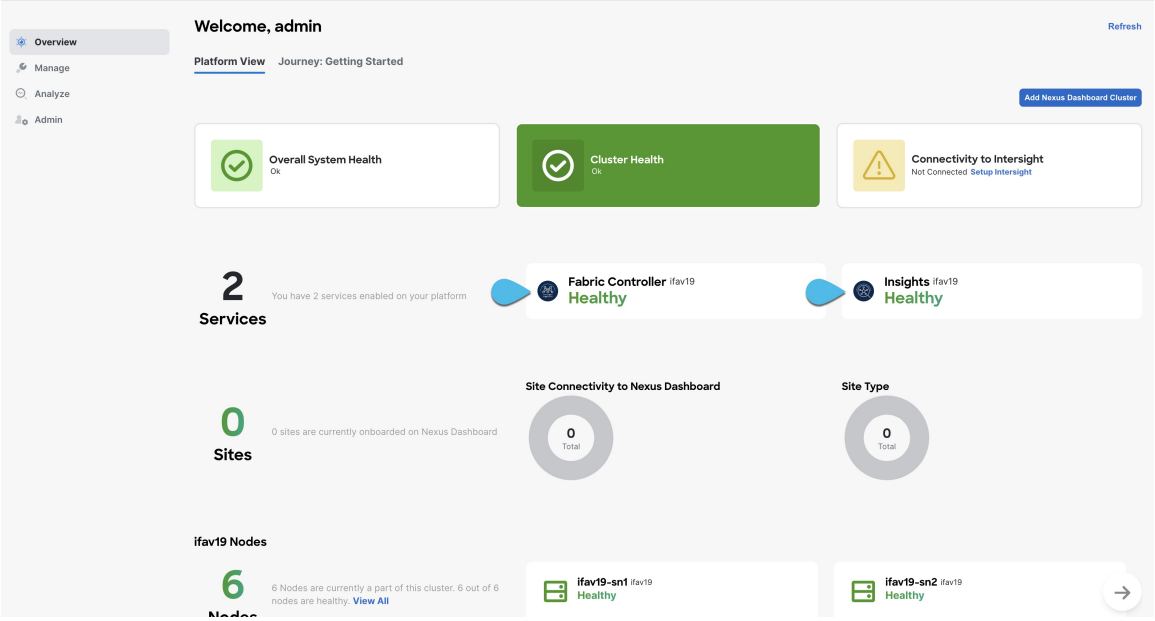
Step 24 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

Note
In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

```
deploy base system services
```

This is due to an issue with `etcd` on the node after a reboot of the pND (Physical Nexus Dashboard) cluster.

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

Step 25 After you have deployed your Nexus Dashboard and services, you can configure each service as described in its configuration and operations articles.

- For Fabric Controller, see the *NDFC persona configuration* white paper and documentation library.
- For Orchestrator, see the documentation page.

- For Insights, see the documentation library.

Deploying Nexus Dashboard Directly in VMware ESXi

This section describes how to deploy Cisco Nexus Dashboard cluster directly in VMware ESXi. If you prefer to deploy using vCenter, follow the steps described in Deploying Nexus Dashboard Directly in VMware ESXi, on page 103 instead.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 81.

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/

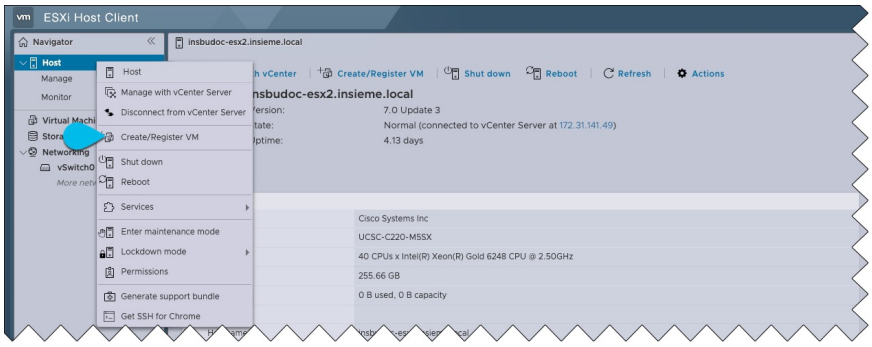
b) Choose the Nexus Dashboard release version you want to download.

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

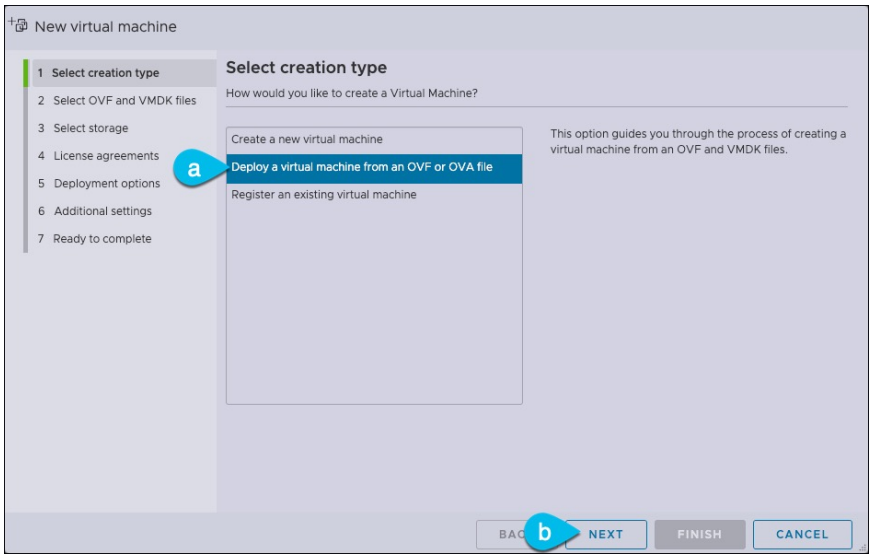
Step 2 Log in to your VMware ESXi.

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 Prerequisites: Nexus Dashboard, on page 9.

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

  Node profiles are described in Prerequisites and Guidelines, on page 81.

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

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

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

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

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

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

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

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

To disable time synchronization:

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

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

a) Begin initial setup.

You will be prompted to run the first-time setup utility:

```
[ OK ] Started atomix-boot-setup.
       Starting Initial cloud-init job (pre-networking)...
       Starting logrotate...
       Starting logwatch...
       Starting keyhole...
[ OK ] Started keyhole.
[ OK ] Started logrotate.
[ OK ] Started logwatch.

Press any key to run first-boot setup on this console...
```

b) Enter and confirm the `admin` password

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

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

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

c) Enter the management network information.

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

d) For the first node only, designate it as the "Cluster Leader".

You will log into the cluster leader node to finish configuration and complete cluster creation.

```
Is this the cluster leader?: y
```

e) Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, choose `n` to proceed. If you want to change any of the entered information, enter `y` to re-start the basic configuration script.

```
Please review the config
Management network:
  Gateway: 192.168.9.1
  IP Address/Mask: 192.168.9.172/24
Cluster leader: no

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

Step 14 Repeat previous steps to deploy the additional nodes.

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

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

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

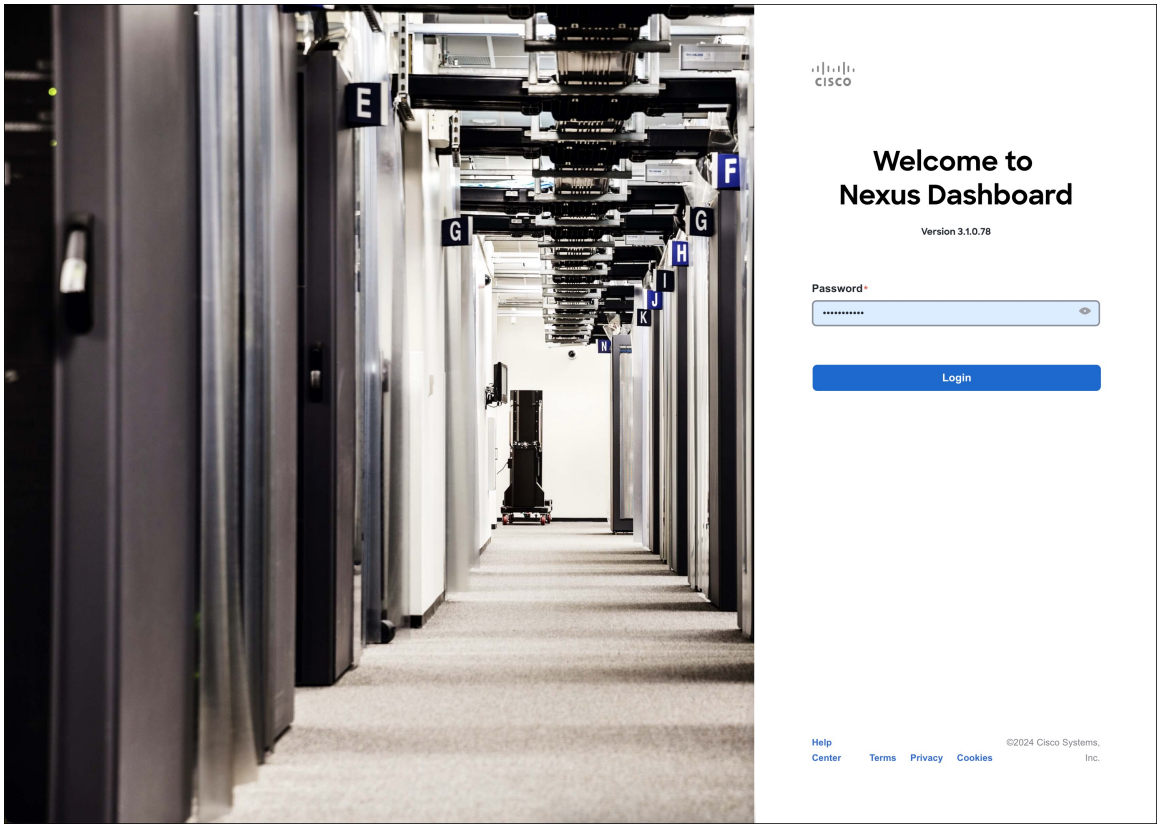
Note
You must provide the same password for all nodes or the cluster creation will fail.

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

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

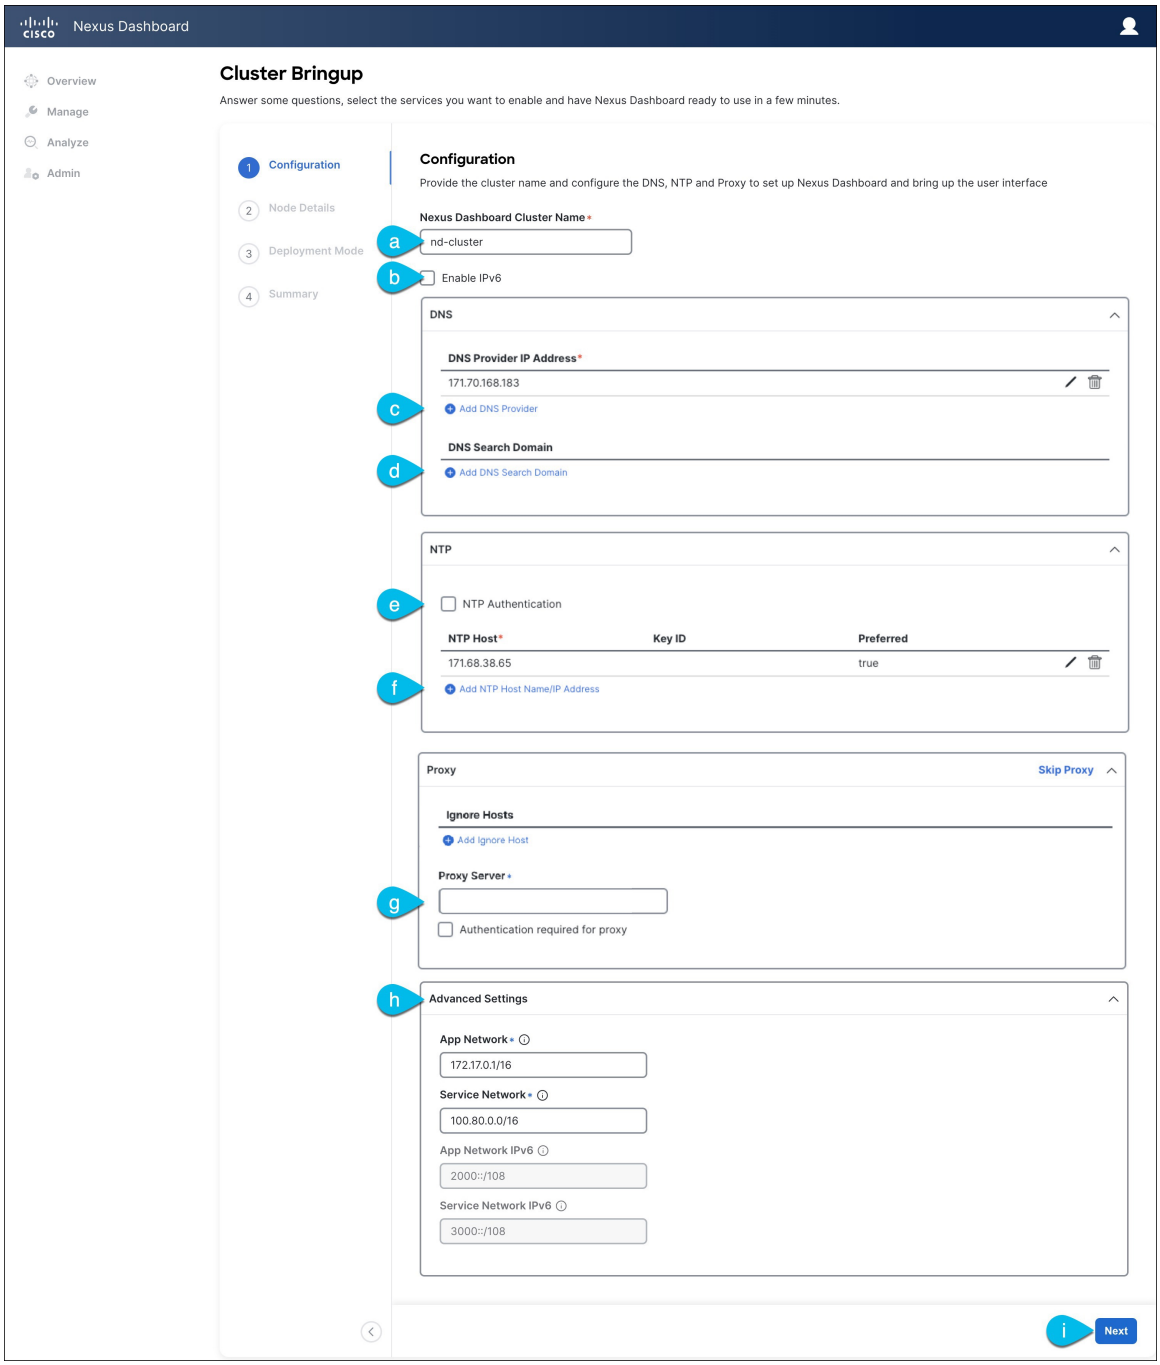
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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



Step 16 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

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

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

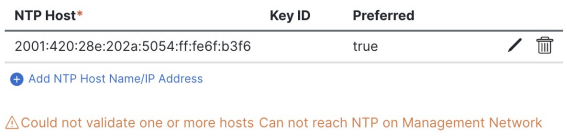
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

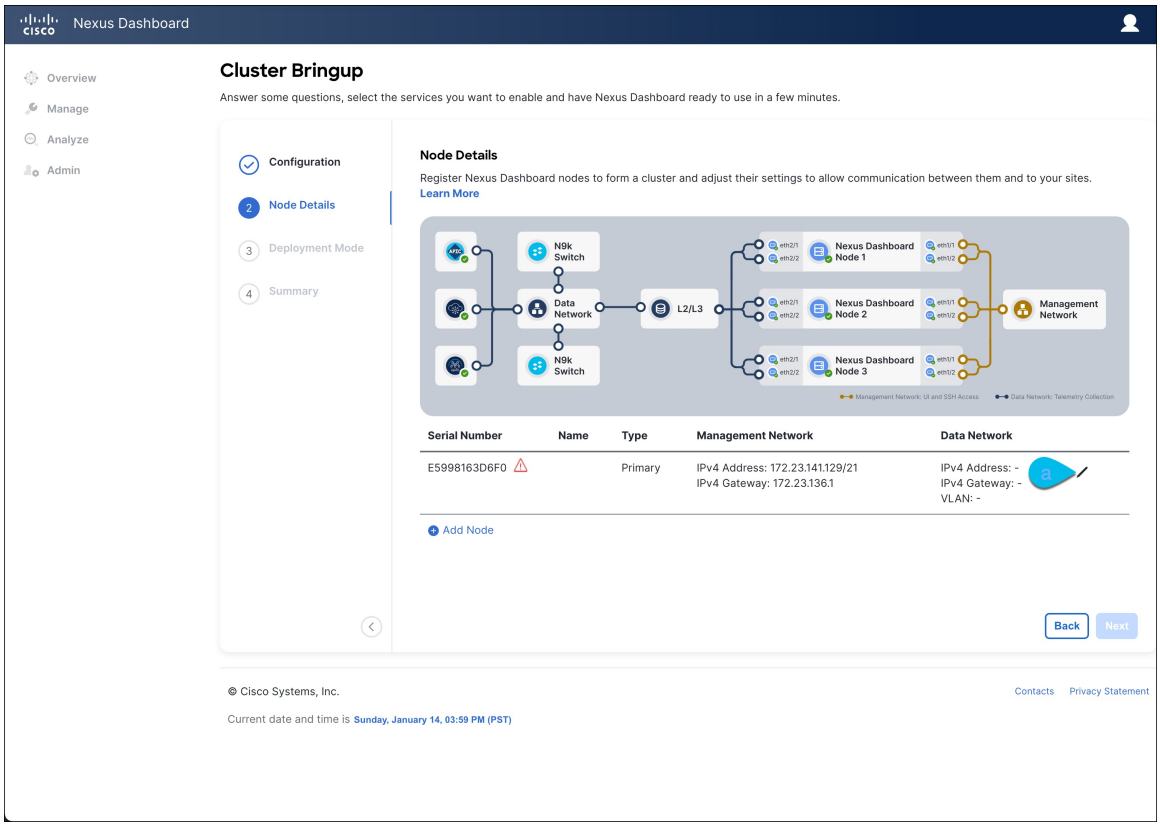
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

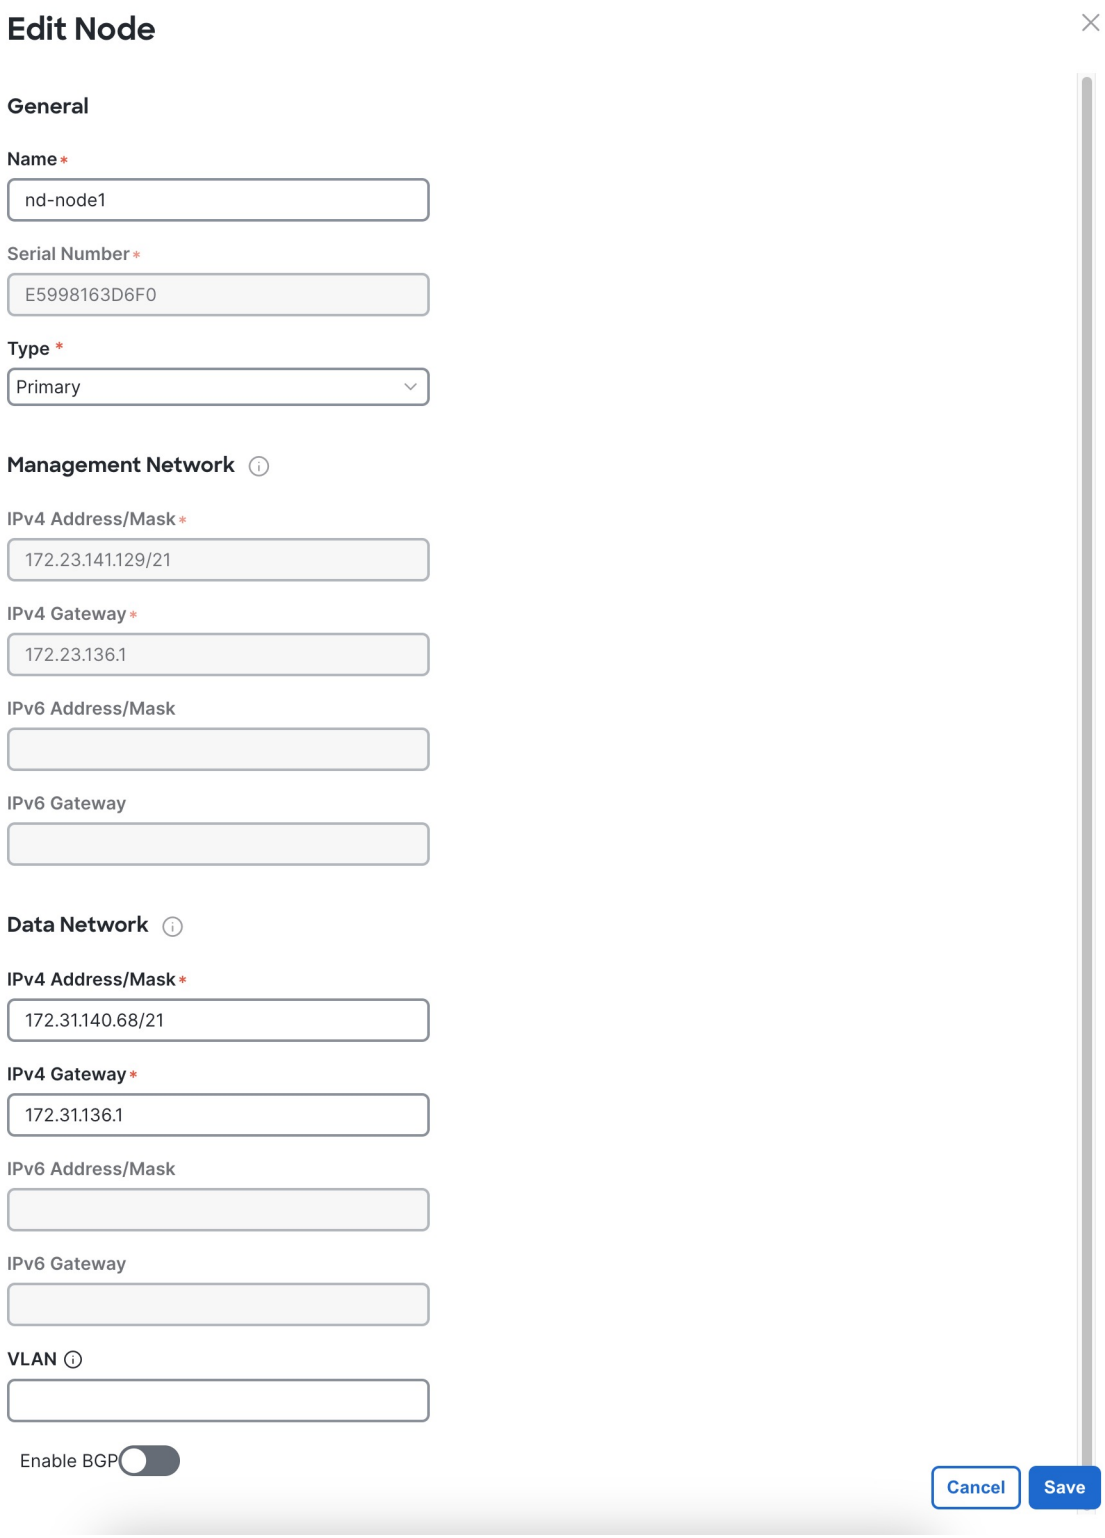
  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) Click **Next** to continue.

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

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


Nexus Dashboard
Overview
Manage
Analyze
Admin
Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.
Configuration
2 Node Details
3 Deployment Mode
4 Summary
Node Details
Register Nexus Dashboard nodes to form a cluster and adjust their settings to allow communication between them and to your sites.
Learn More
N9k Switch
Data Network
N9k Switch
L2/L3
Nexus Dashboard Node 1
Nexus Dashboard Node 2
Nexus Dashboard Node 3
Management Network
Serial Number
Name
Type
Management Network
Data Network
E5998163D6F0
Primary
IPv4 Address: 172.23.141.129/21
IPv4 Gateway: 172.23.136.1
IPv4 Address: -
IPv4 Gateway: -
VLAN: -
Add Node
Back
Next
© Cisco Systems, Inc.
Contacts
Privacy Statement
Current date and time is Sunday, January 14, 03:59 PM (PST)



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

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated but you must provide other information.

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

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

c) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

e) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

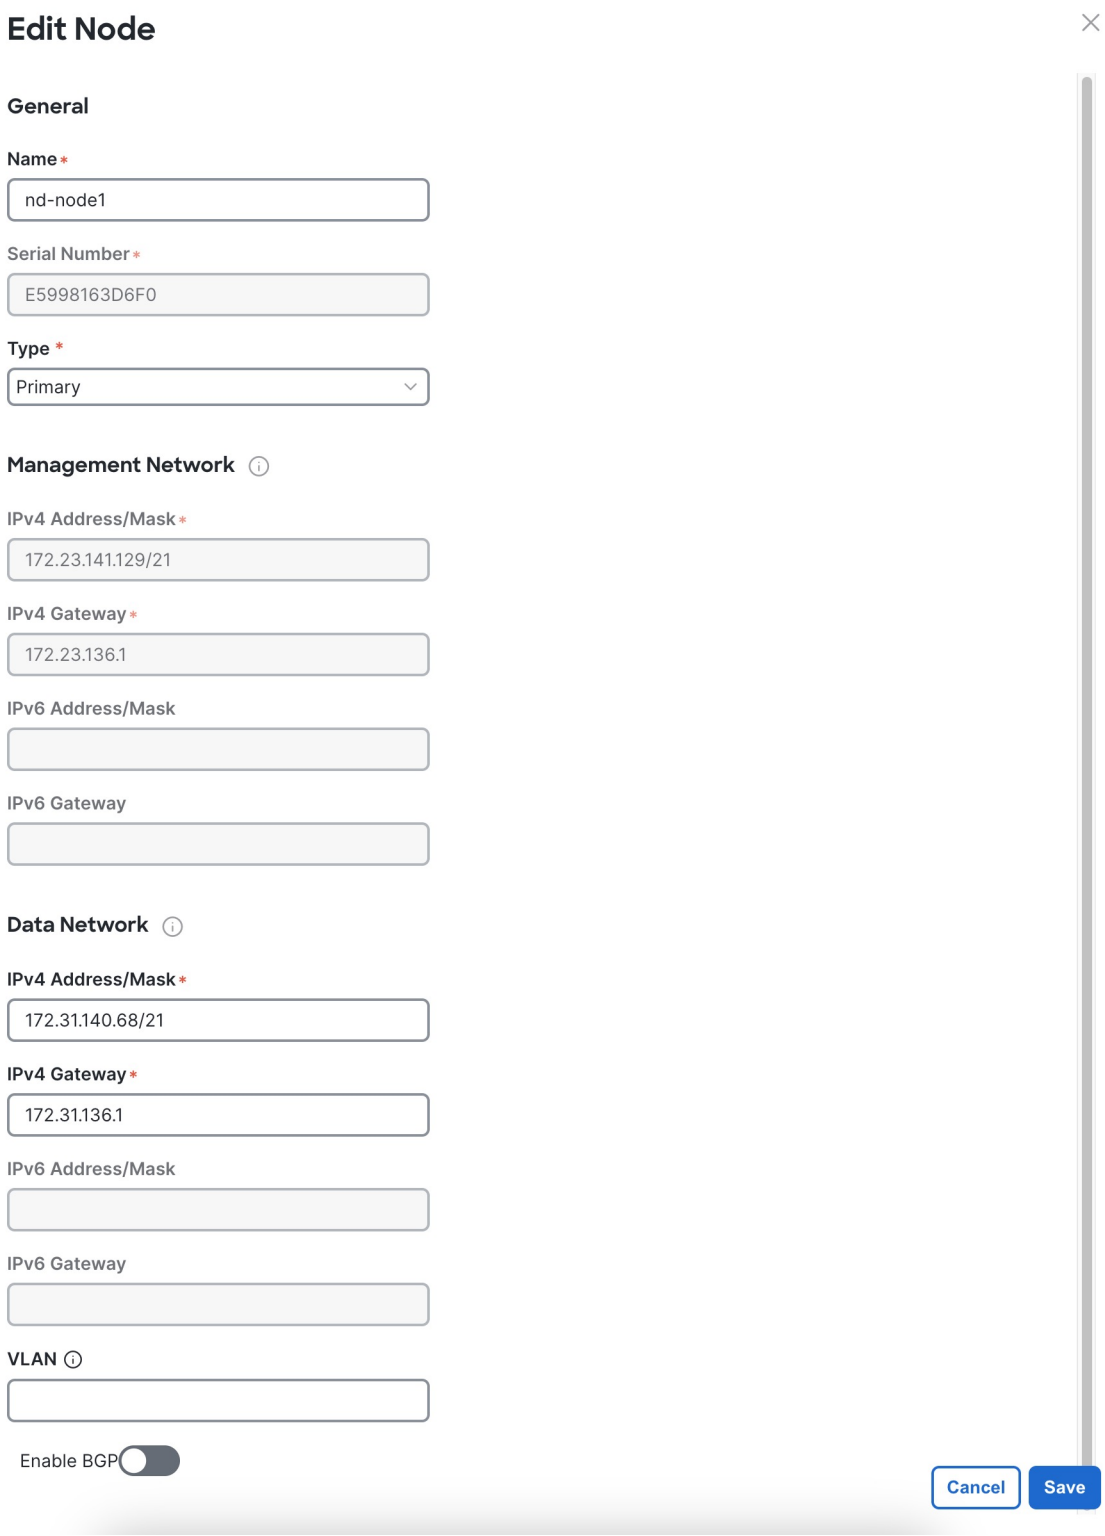
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

Step 18 In the **Node Details** screen, click **Add Node** to add the second node to the cluster.

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



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

You defined the management network information and the password during the initial node configuration steps.

b) Click **Validate** to verify connectivity to the node.

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

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

d) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

f) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

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

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

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

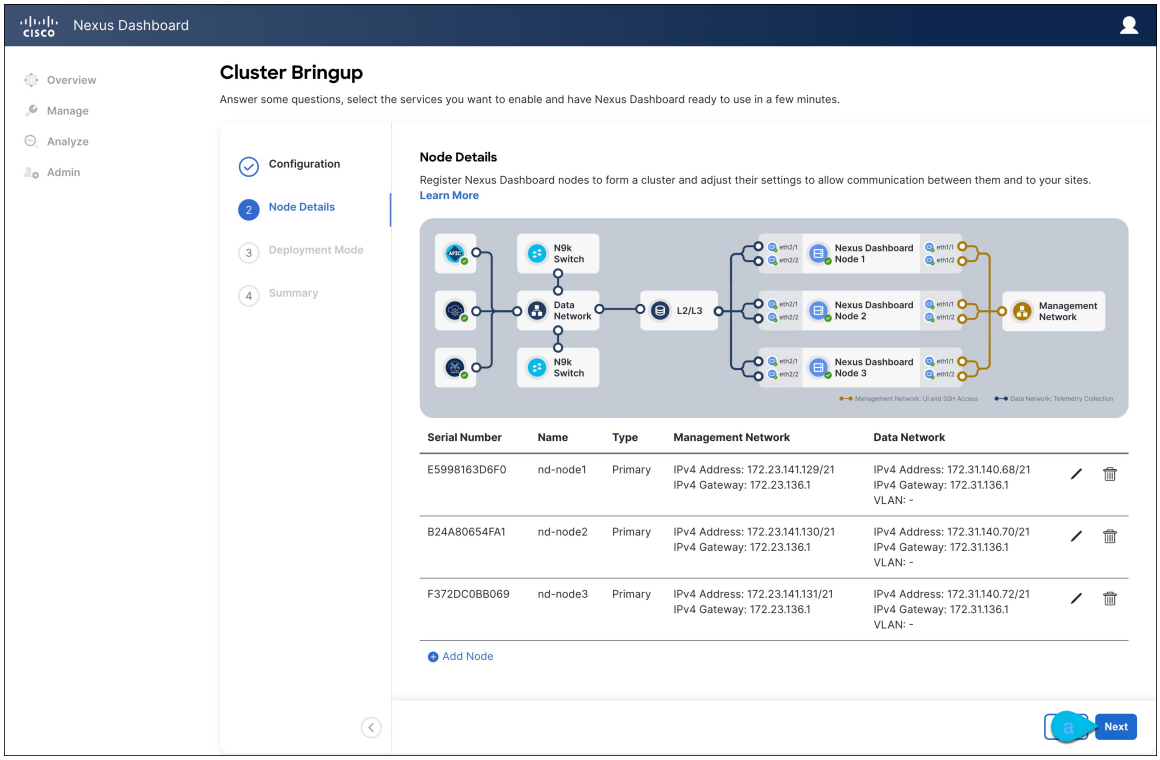
h) Repeat this step for the final (third) primary node of the cluster.

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

Note
In order to enable multiple services concurrently in your cluster or to support higher scale, you must provide sufficient number of secondary nodes during deployment. Refer to the Nexus Dashboard Cluster Sizing tool for exact number of additional secondary nodes required for your specific use case.

You can choose to add the standby nodes now or at a later time after the cluster is deployed.

Step 20 In the **Node Details** page, verify the provided information and click **Next** to continue.



Step 21 Choose the **Deployment Mode** for the cluster.

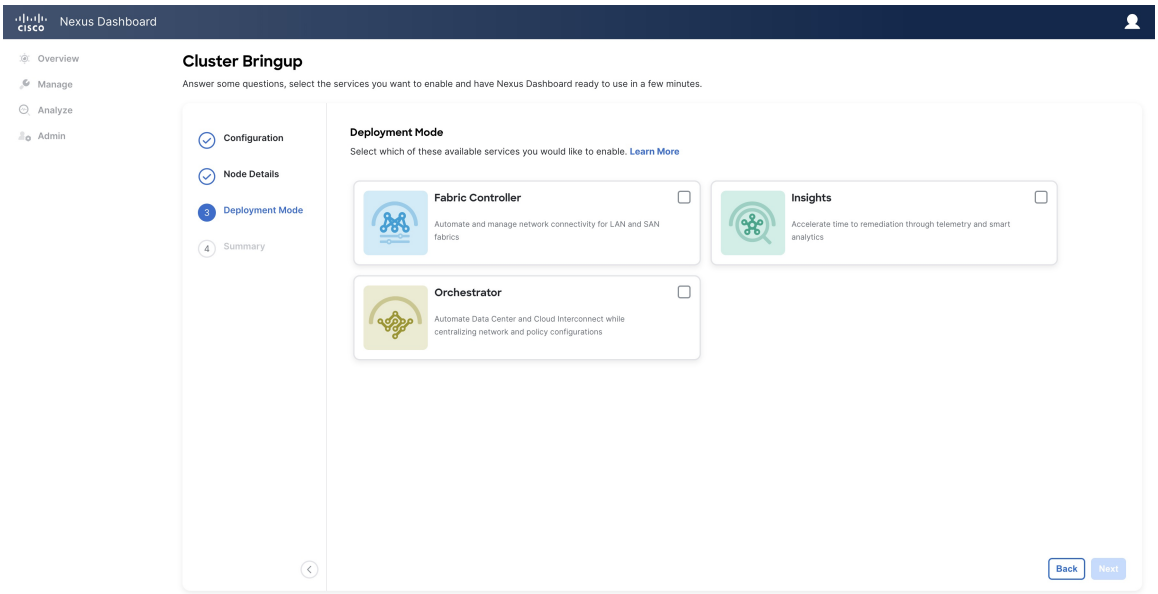
a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

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

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

- Nexus Dashboard Fabric Controller prerequisites
- Nexus Dashboard Orchestrator prerequisites
- Nexus Dashboard Insights prerequisites



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

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section and the service-specific requirements chapters.

c) Click **Next** to proceed.

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

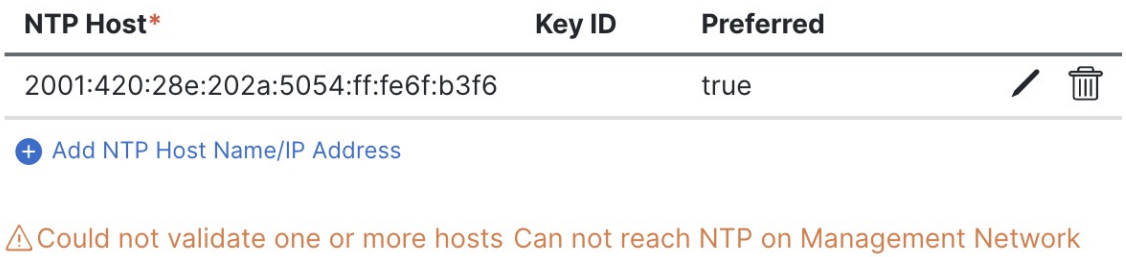
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

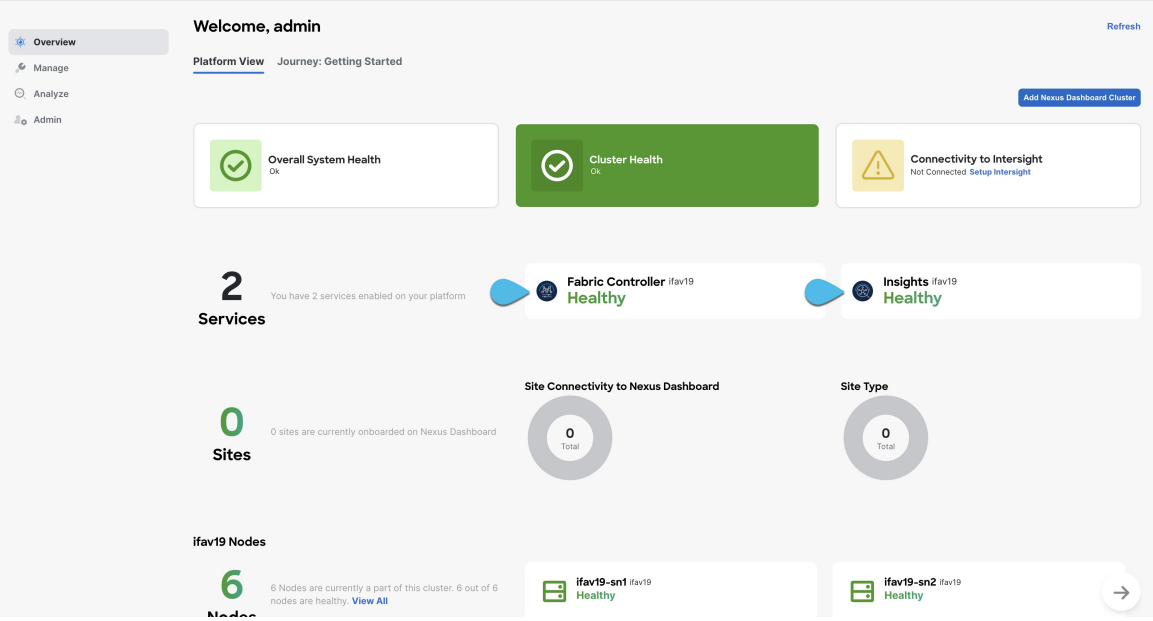
Step 23 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

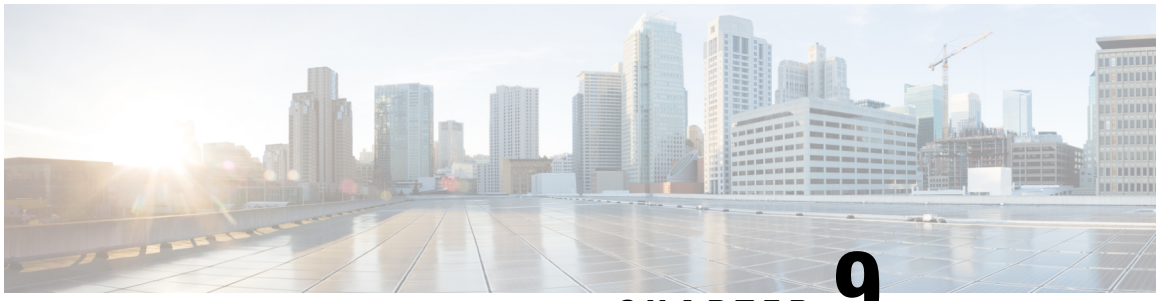
Note
There may be an issue during the bootstrap process on 3-node vND (ESX) clusters which can cause the 'acs health' command to show the following error: `'k8s: services not in desired state - aaamgr,cisco-intersightdc,eventmonitoring,infra-kafka,kafka,mongodb,sm,statscollect'`

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 After you have deployed your Nexus Dashboard and services, you can configure each service as described in its configuration and operations articles.

- For Fabric Controller, see the *NDFC persona configuration* white paper and documentation library.
- For Orchestrator, see the documentation page.

- For Insights, see the documentation library.



CHAPTER 9

Deploying in Linux KVM

- Prerequisites and Guidelines, on page 119
- Deploying Nexus Dashboard in Linux KVM, on page 120

Prerequisites and Guidelines

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

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

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

- Review and complete the general prerequisites described in Prerequisites: Nexus Dashboard, on page 9.

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

- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.

- 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.

- Ensure you have enough system resources:

Table 21: Deployment Requirements

Requirements

• KVM deployments are supported for Nexus Dashboard Fabric Controller services only.
• You must deploy in CentOS 7.9 or Red Hat Enterprise Linux 8.6
• You must have the supported versions of Kernel and KVM:
  • For CentOS 7.9, Kernel version `3.10.0-957.el7.x86_64` and KVM version `libvirt-4.5.0-23.el7_7.1.x86_64`
  • For RHEL 8.6, Kernel version `4.18.0-372.9.1.el8.x86_64` and KVM version `libvert 8.0.0`
• 16 vCPUs
• 64 GB of RAM
• 550 GB disk
Each node requires a dedicated disk partition
• The disk must have I/O latency of 20ms or less.
To verify the I/O latency:
1. Create a test directory.
For example, `test-data`.
2. Run the following command:
`# fio --rw=write --ioengine=sync --fdatasync=1 --directory=test-data --size=22m --bs=2300 --name=mytest`
3. After the command is executed, confirm that the `99.00th=[<value>]` in the `fsync/fdatasync/sync_file_range` section is under 20ms.
• We recommend that each Nexus Dashboard node is deployed in a different KVM hypervisor.

Deploying Nexus Dashboard in Linux KVM

This section describes how to deploy Cisco Nexus Dashboard cluster in Linux KVM.

Before you begin

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

Procedure

Step 1 Download the Cisco Nexus Dashboard image.

a) Browse to the Software Download page.

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

b) Click **Nexus Dashboard Software**.

c) From the left sidebar, choose the Nexus Dashboard version you want to download.

d) Download the Cisco Nexus Dashboard image for Linux KVM (`nd-dk9.<version>.qcow2`).

Step 2 Copy the image to the Linux KVM servers where you will host the nodes.

You can use `scp` to copy the image, for example:

```
# scp nd-dk9.<version>.qcow2 root@<kvm-host-ip>:/home/nd-base
```

The following steps assume you copied the image into the `/home/nd-base` directory.

Step 3 Create the required disk images for the first node.

You will create a snapshot of the base `qcow2` image you downloaded and use the snapshots as the disk images for the nodes' VMs. You will also need to create a second disk image for each node.

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

b) Create a directory for the node's snapshot.

The following steps assume you create the snapshot in the `/home/nd-node1` directory.

```
# mkdir -p /home/nd-node1/
# cd /home/nd-node1
```

c) Create the snapshot.

In the following command, replace `/home/nd-base/nd-dk9.<version>.qcow2` with the location of the base image you created in the previous step.

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.<version>.qcow2
/home/nd-node1/nd-node1-disk1.qcow2
```

Note
If you are deploying in RHEL 8.6, you may need to provide an additional parameter to define the destination snapshot's format as well. In that case, update the above command to the following:

```
# qemu-img create -f qcow2 -b /home/nd-base/nd-dk9.2.1.1a.qcow2 /home/nd-node1/nd-node1-disk1.qcow2
-F qcow2
```

d) Create the additional disk image for the node.

Each node requires two disks: a snapshot of the base Nexus Dashboard `qcow2` image and a second 500GB disk.

```
# qemu-img create -f qcow2 /home/nd-node1/nd-node1-disk2.qcow2 500G
```

Step 4 Repeat the previous step to create the disk images for the second and third nodes.

Before you proceed to the next step, you should have the following:

- For the first node, `/home/nd-node1/` directory with two disk images:

- `/home/nd-node1/nd-node1-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
- `/home/nd-node1/nd-node1-disk2.qcow2`, which is a new 500GB disk you created.

- For the second node, `/home/nd-node2/` directory with two disk images:
  - `/home/nd-node2/nd-node2-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
  - `/home/nd-node2/nd-node2-disk2.qcow2`, which is a new 500GB disk you created.

- For the third node, `/home/nd-node3/` directory with two disk images:
  - `/home/nd-node1/nd-node3-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
  - `/home/nd-node1/nd-node3-disk2.qcow2`, which is a new 500GB disk you created.

Step 5 Create the first node's VM.

a) Open the KVM console and click **New Virtual Machine**.

You can open the KVM console from the command line using the `virt-manager` command.

If your Linux KVM environment does not have a desktop GUI, run the following command instead and proceed to step 6.

```
virt-install --import --name <node-name> --memory 65536 --vcpus 16 --os-type generic --disk
path=/path/to/disk1/nd-node1-d1.qcow2,format=qcow2,bus=virtio --disk
path=/path/to/disk2/nd-node1-d2.qcow2,format=qcow2,bus=virtio --network
bridge=<mgmt-bridge-name>,model=virtio --network bridge=<data-bridge-name>,model=virtio --console
pty,target_type=serial --noautoconsole --autostart
```

b) In the **New VM** screen, choose **Import existing disk image option** and click **Forward**.

c) In the **Provide existing storage path** field, click **Browse** and select the `nd-node1-disk1.qcow2` file.

We recommend that each node's disk image is stored on its own disk partition.

d) Choose `Generic` for the **OS type** and **Version**, then click **Forward**.

e) Specify 64GB memory and 16 CPUs, then click **Forward**.

f) Enter the **Name** of the virtual machine, for example `nd-node1` and check the **Customize configuration before install** option. Then click **Finish**.

Note
You must select the **Customize configuration before install** checkbox to be able to make the disk and network card customizations required for the node.

The VM details window will open.

In the VM details window, change the NIC's device model:

a) Select **NIC <mac>**.

b) For **Device model**, choose `e1000`.

c) For **Network Source**, choose the bridge device and provide the name of the "mgmt" bridge.

Note

Creating bridge devices is outside the scope of this guide and depends on the distribution and version of the operating system. Consult the operating system's documentation, such as Red Hat's *Configuring a network bridge*, for more information.

In the VM details window, add a second NIC:

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

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

a) Click **Add Hardware**.
b) In the **Add New Virtual Hardware** screen, select **Storage**.
c) For the disk's bus driver, choose `IDE`.
d) Select **Select or create custom storage**, click **Manage**, and select the `nd-node1-disk2.qcow2` file you created.
e) Click **Finish** to add the second disk.

Note
Ensure that you enable the `Copy host CPU configuration` option in the Virtual Machine Manager UI.

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

Step 6 Repeat previous steps to deploy the second and third nodes, then start all VMs.

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

Step 7 Open one of the node's console and configure the node's basic information. If your Linux KVM environment does not have a desktop GUI, run the `virsh console <node-name>` command to access the console of the node.

a) Press any key to begin initial setup.

You will be prompted to run the first-time setup utility:

```
[ OK ] Started atomix-boot-setup.
       Starting Initial cloud-init job (pre-networking)...
       Starting logrotate...
       Starting logwatch...
       Starting keyhole...
[ OK ] Started keyhole.
[ OK ] Started logrotate.
[ OK ] Started logwatch.

Press any key to run first-boot setup on this console...
```

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: yes

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

Step 8 Repeat previous step to configure the initial information for the second and third nodes.

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

Note
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 9 Wait for the initial bootstrap process to complete on all nodes.

After you provide and confirm management network information, the initial setup on the first node (`Cluster Leader`) configures the networking and brings up the UI, which you will use to add two other nodes and complete the cluster deployment.

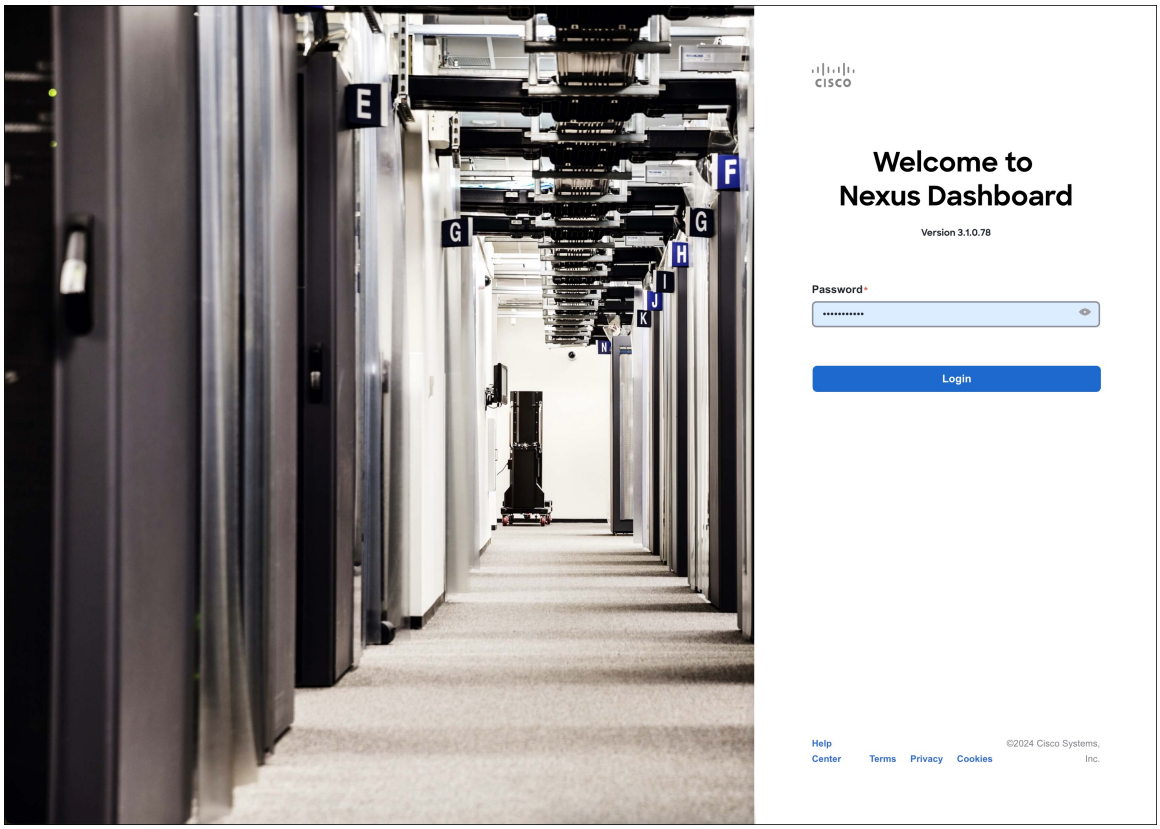
```
Please wait for system to boot: [#########################] 100%
System up, please wait for UI to be online.

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

Step 10 Open your browser and navigate to `https://<node-mgmt-ip>` to open the GUI.

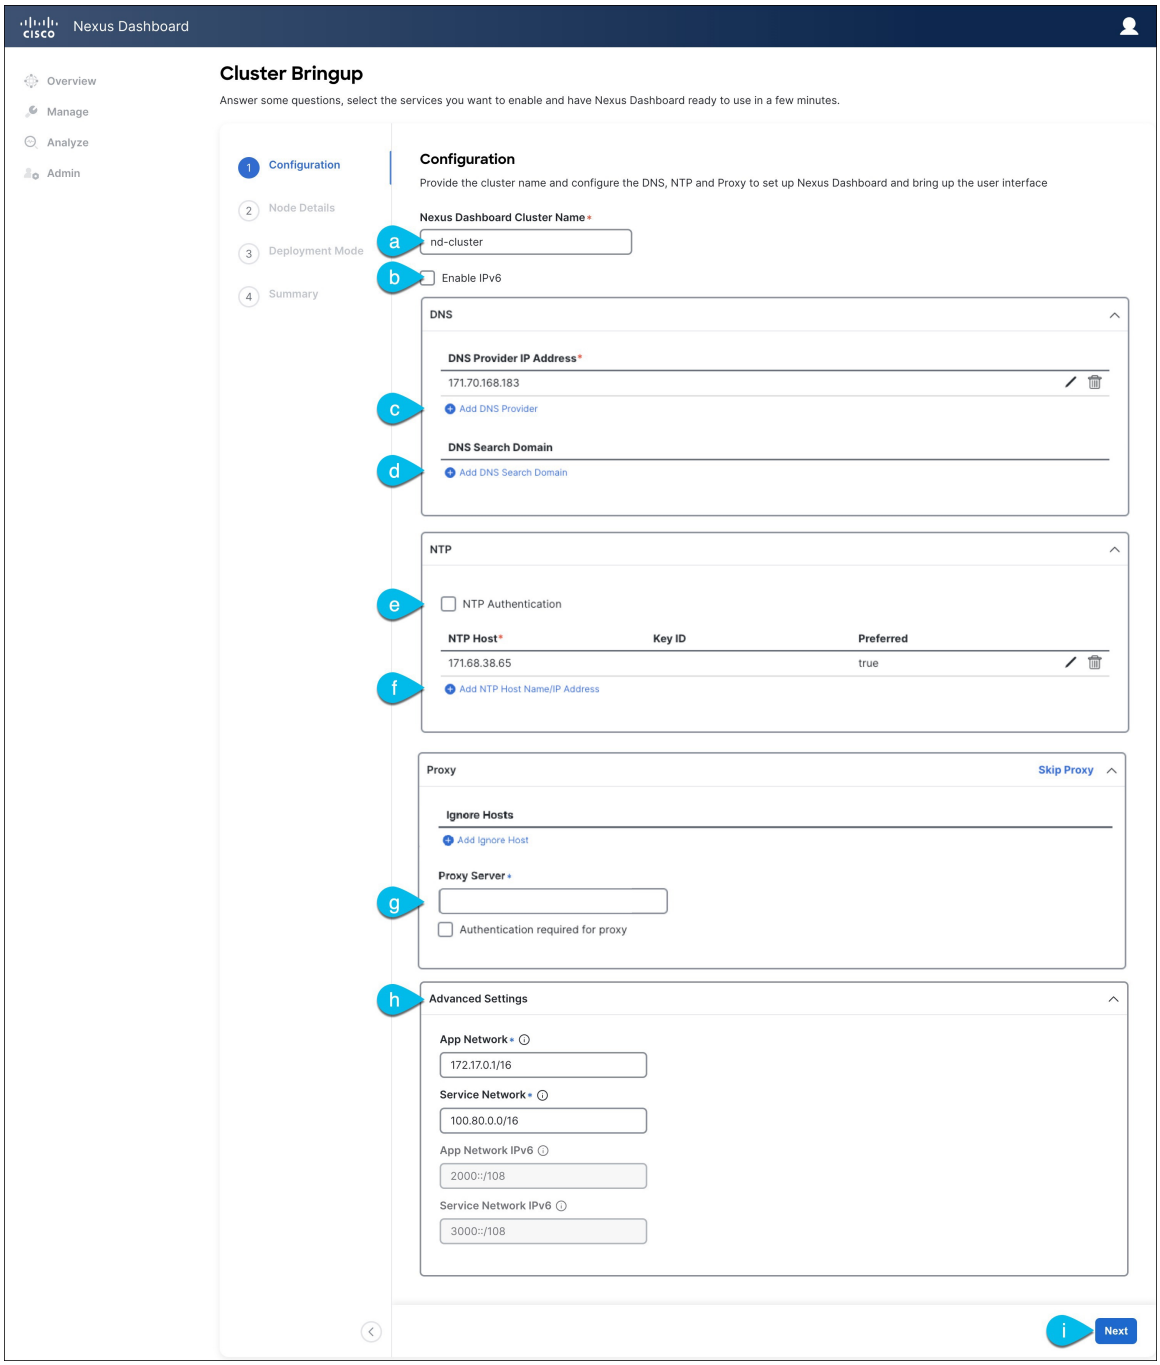
The rest of the configuration workflow takes place from one of the node's GUI. You can choose any one of the nodes you deployed to begin the bootstrap process and you do not need to log in to or configure the other two nodes directly.

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



Step 11 Provide the **Cluster Details**.

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



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

The cluster name must follow the RFC-1123 requirements.

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

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

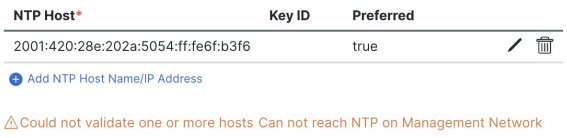
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

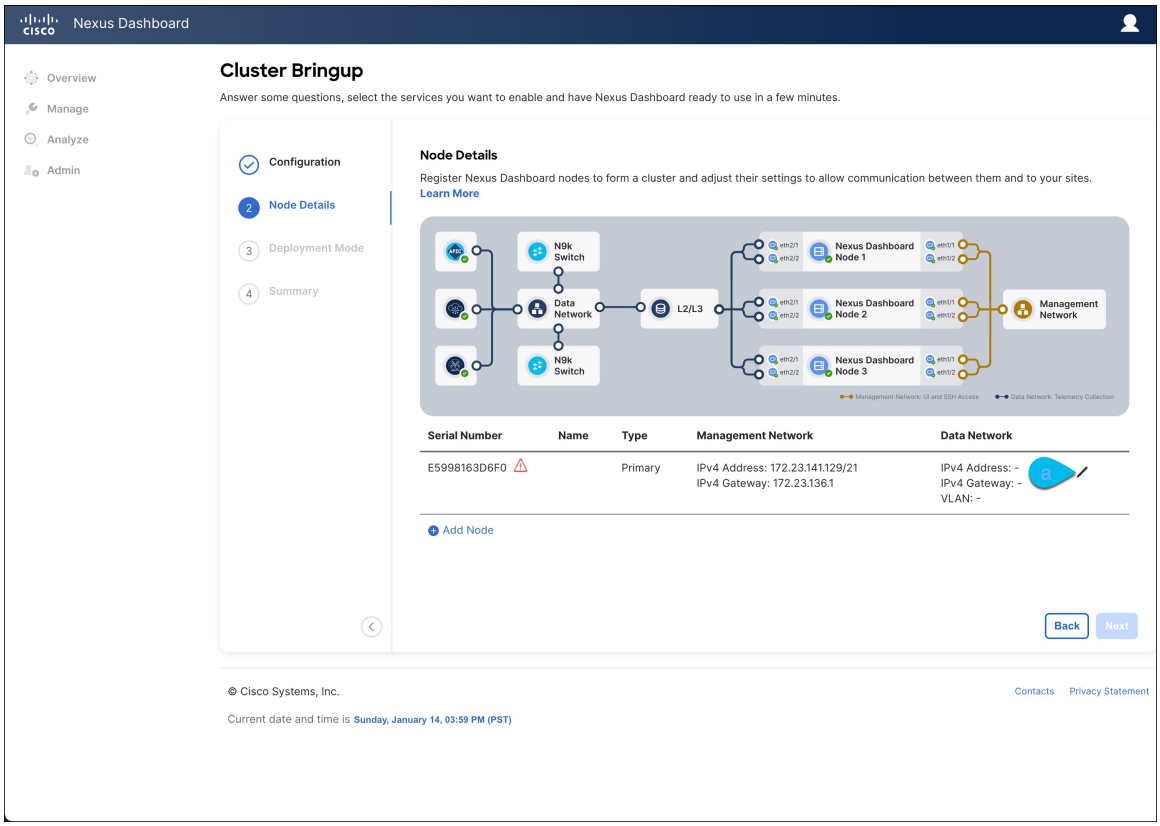
  If you have checked the **Enable IPv6** option earlier, you can also define the IPv6 subnets for the App and Service networks.

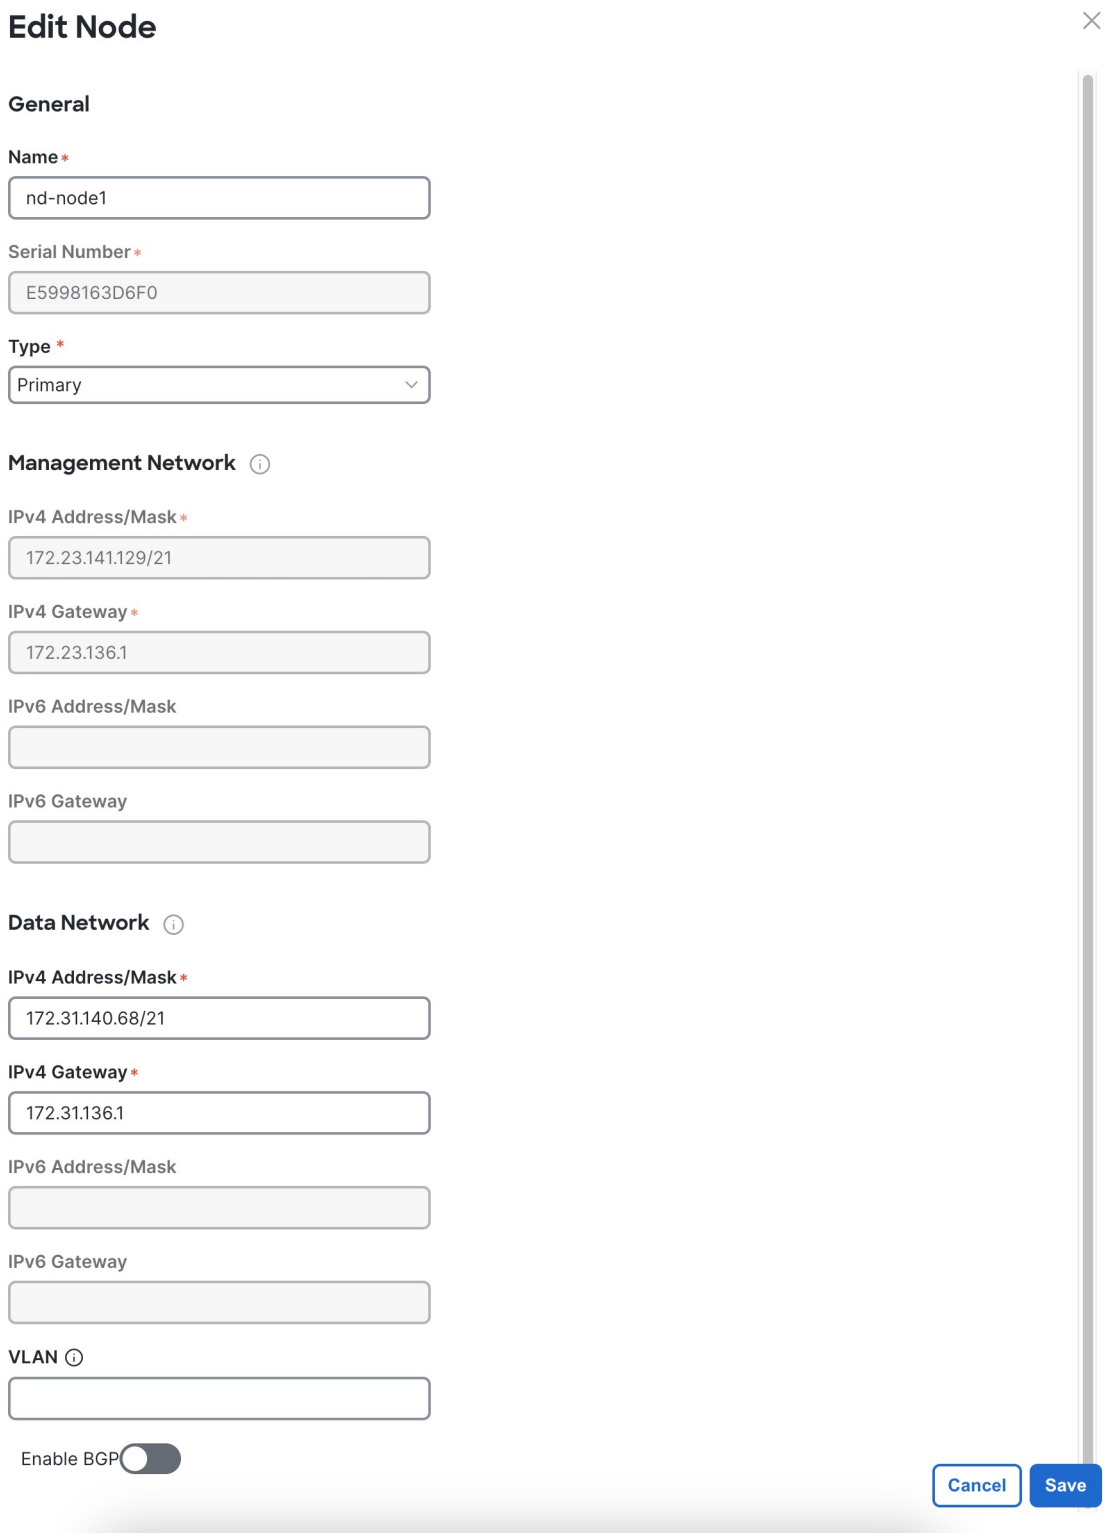
  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) Click **Next** to continue.

Step 12 In the **Node Details** screen, update the first node's information.

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


Nexus Dashboard
Overview
Manage
Analyze
Admin
Cluster Bringup
Answer some questions, select the services you want to enable and have Nexus Dashboard ready to use in a few minutes.
Configuration
2 Node Details
3 Deployment Mode
4 Summary
Node Details
Register Nexus Dashboard nodes to form a cluster and adjust their settings to allow communication between them and to your sites.
Learn More
N9k Switch
Data Network
N9k Switch
L2/L3
Nexus Dashboard Node 1
Nexus Dashboard Node 2
Nexus Dashboard Node 3
Management Network
Serial Number
Name
Type
Management Network
Data Network
E5998163D6F0
Primary
IPv4 Address: 172.23.141.129/21
IPv4 Gateway: 172.23.136.1
IPv4 Address: -
IPv4 Gateway: -
VLAN: -
Add Node
Back
Next
© Cisco Systems, Inc.
Contacts
Privacy Statement
Current date and time is Sunday, January 14, 03:59 PM (PST)



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

The node's **Serial Number**, **Management Network** information, and **Type** are automatically populated but you must provide other information.

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

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

c) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

e) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

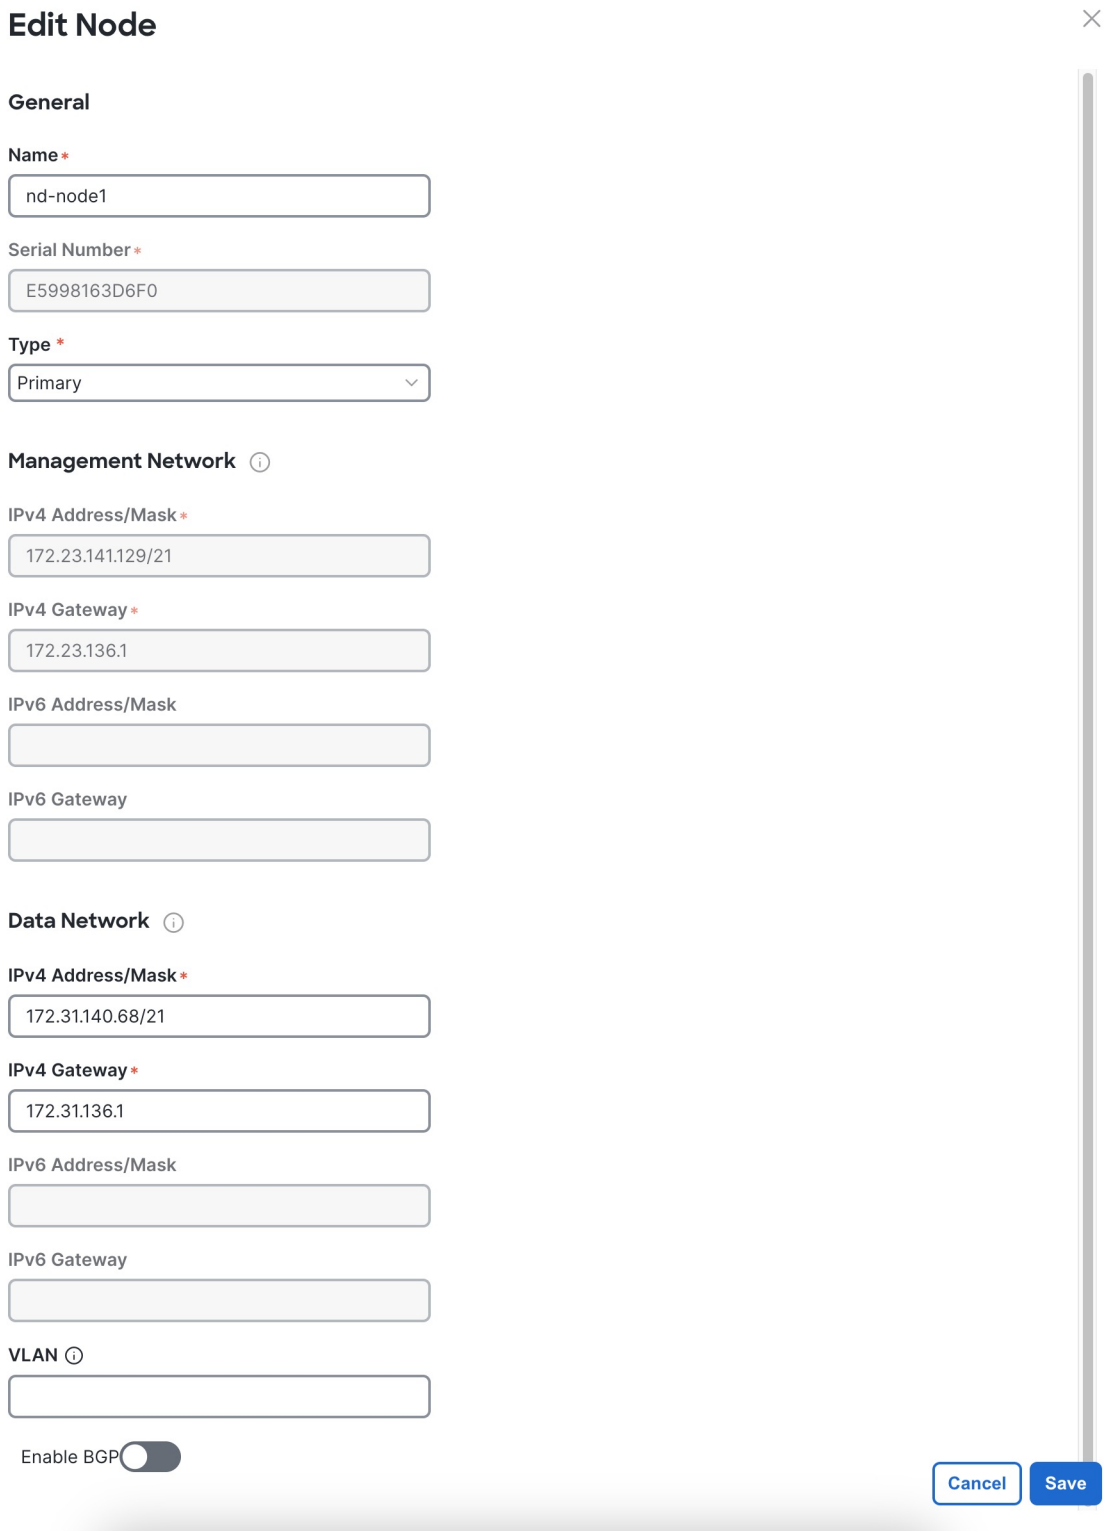
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

Step 13 In the **Node Details** screen, click **Add Node** to add the second node to the cluster.

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



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

You defined the management network information and the password during the initial node configuration steps.

b) Click **Validate** to verify connectivity to the node.

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

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

d) From the **Type** dropdown, select `Primary`.

The first 3 nodes of the cluster must be set to `Primary`. You will add the secondary nodes in a later step if require to enable cohosting of services and higher scale.

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

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

If you had enabled IPv6 functionality in a previous screen, you must also provide the IPv6 address, netmask, and gateway.

Note
If you want to provide IPv6 information, you must do it during cluster bootstrap process. To change IP configuration later, you would need to redeploy the cluster.

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

f) (Optional) If your cluster is deployed in L3 HA mode, **Enable BGP** for the data network.

BGP configuration is required for the Persistent IPs feature used by some services, such as Insights and Fabric Controller. This feature is described in more detail in Prerequisites and Guidelines, on page 9 and the "Persistent IP Addresses" sections of the *Cisco Nexus Dashboard User Guide*.

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

If you choose to enable BGP, you must also provide the following information:

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

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

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

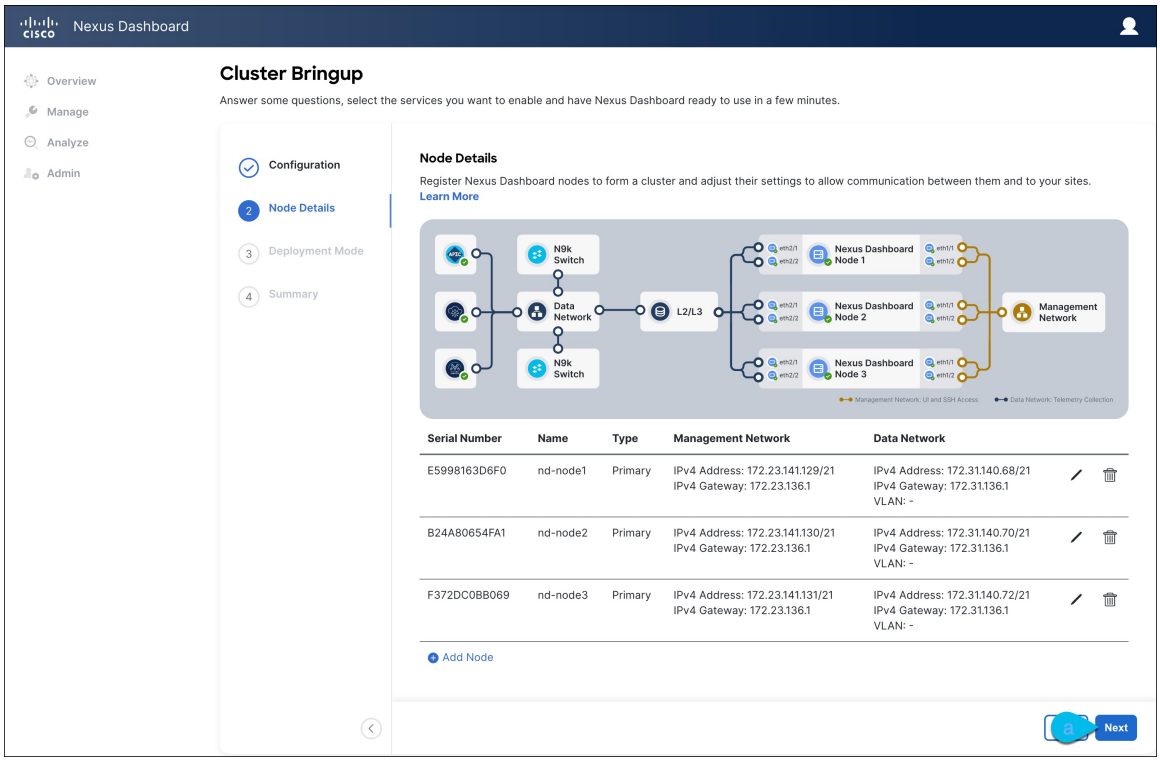
  The router ID must be an IPv4 address, for example `1.1.1.1`

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

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

h) Repeat this step for the final (third) primary node of the cluster.

Step 14 In the **Node Details** page, verify the provided information and click **Next** to continue.



Step 15 Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

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

b) Click **Add Persistent Service IPs/Pools** to provide one or more persistent IPs required by Insights or Fabric Controller services.

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section.

c) Click **Next** to proceed.

Step 16 In the **Summary** screen, review and verify the configuration information and click **Save** to build the cluster.

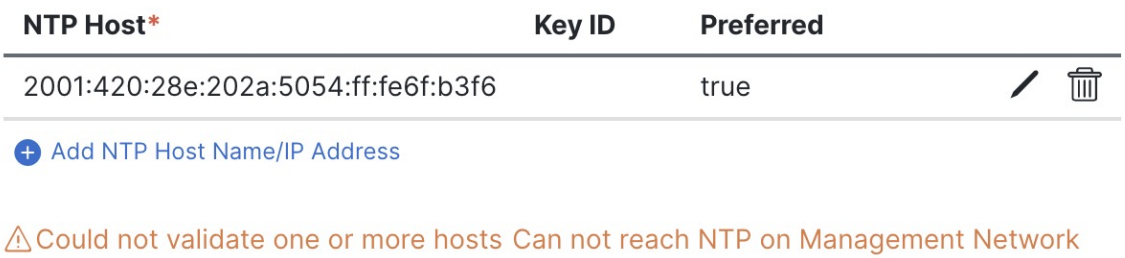
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

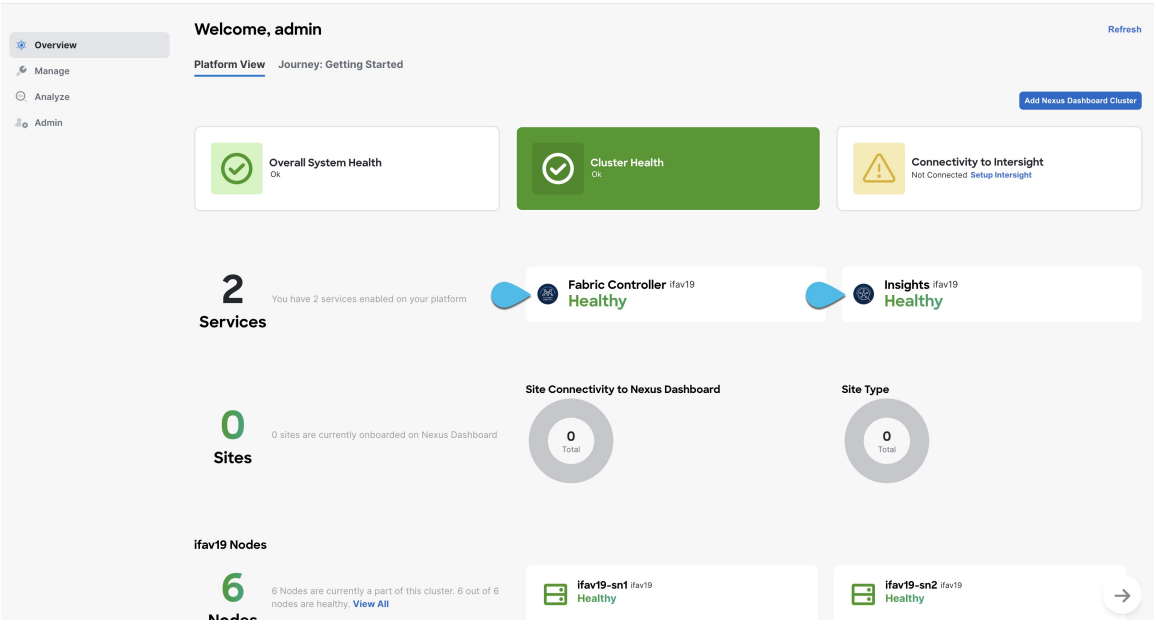
Step 17 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

Note
In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

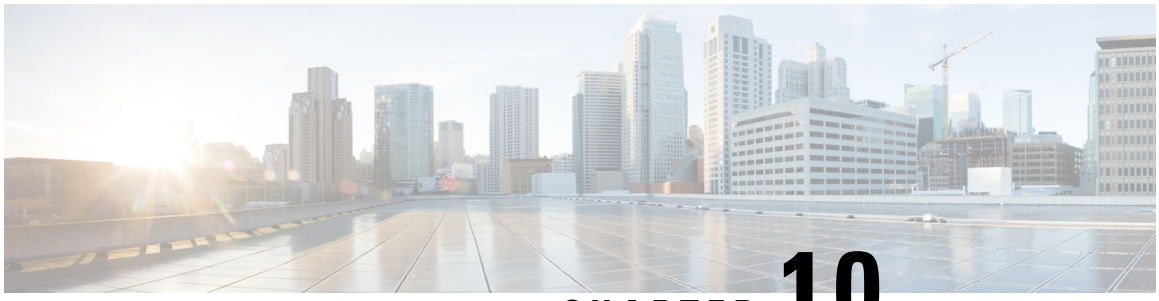
```
deploy base system services
```

This is due to an issue with `etcd` on the node after a reboot of the pND (Physical Nexus Dashboard) cluster.

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

Step 18 After you have deployed your Nexus Dashboard and services, you can configure each service as described in its configuration and operations articles.

- For Fabric Controller, see the *NDFC persona configuration* white paper and documentation library.
- For Orchestrator, see the documentation page.
- For Insights, see the documentation library.



CHAPTER 10

Deploying in Amazon Web Services

- Prerequisites and Guidelines, on page 137
- Deploying Nexus Dashboard in AWS, on page 139

Prerequisites and Guidelines

Note Only the Nexus Dashboard Orchestrator service can be deployed on a cloud-hosted form factor.

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.

- Review and complete the general prerequisites described in the Deployment Overview, on page 5.
- 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.

- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.
- 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 137.

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 **Fulfillment Option** dropdown, select `Nexus Dashboard - Cloud Deployment`
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.

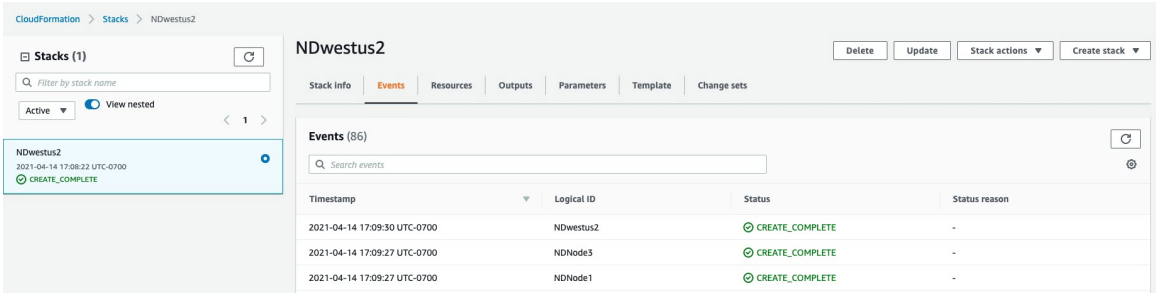
Step 6 In the **Advanced options** screen, simply click **Next**.

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

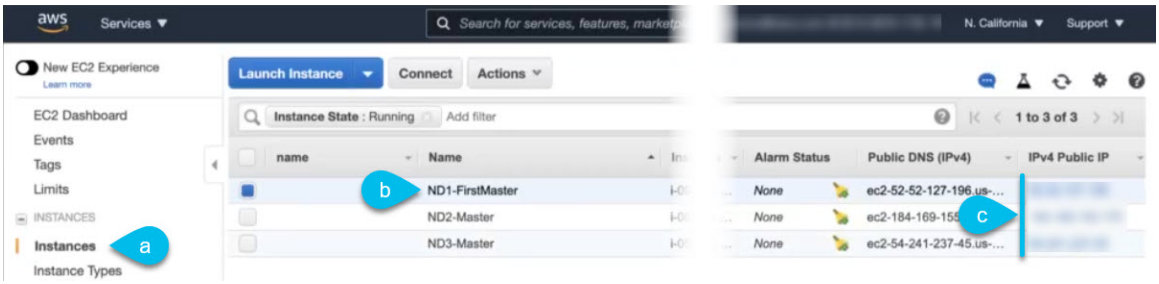
Step 8 Wait for the deployment to complete, then start the VMs.

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 `ND1-Master`.

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 137.

```
# 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
```

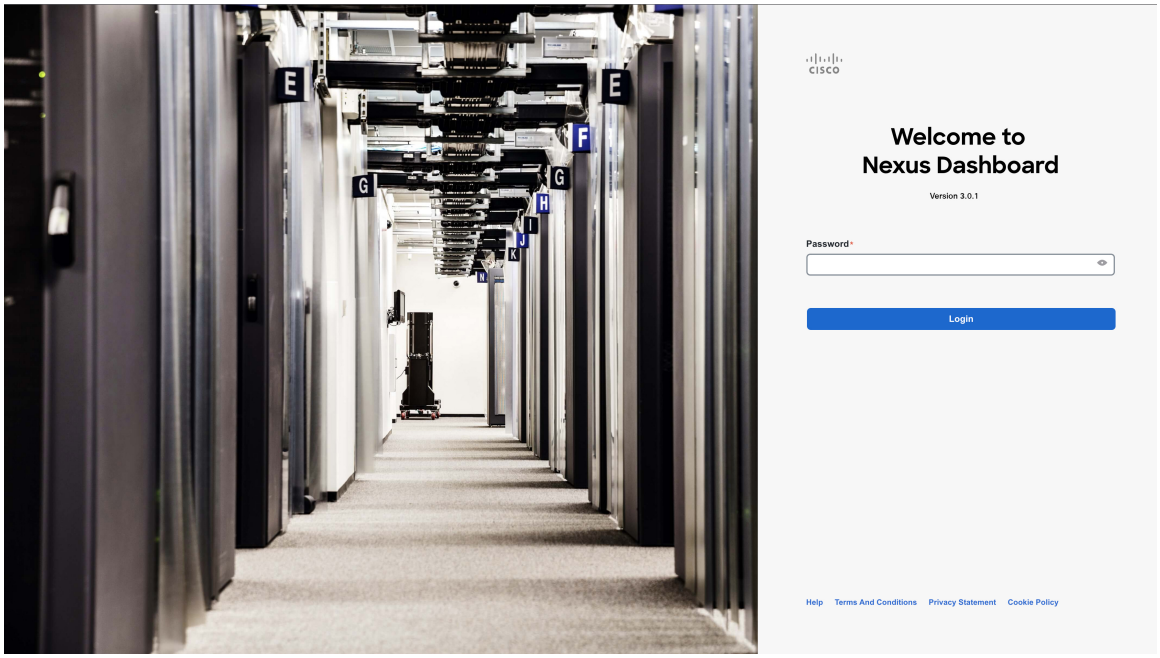
c) Repeat this step for the other two instances.

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 (`ND1-Master`) or cluster configuration cannot be completed.

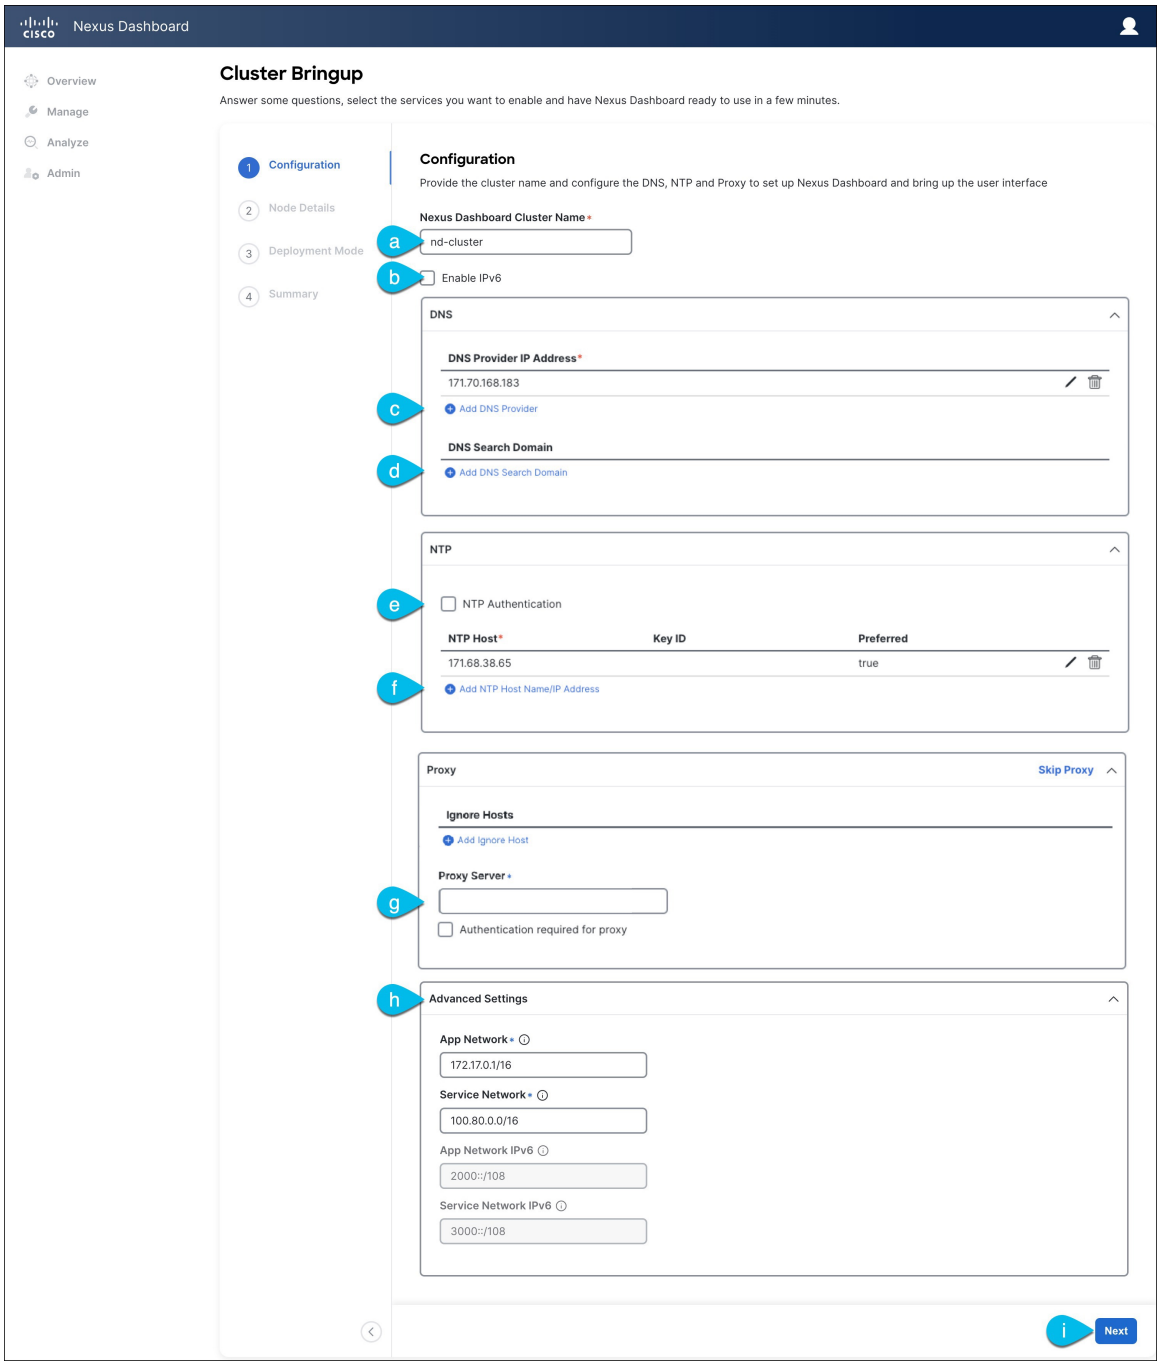
The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

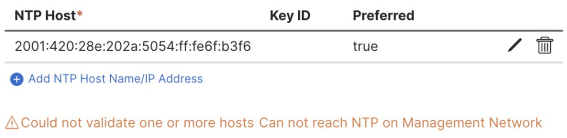
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

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

  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) Click **Next** to continue.

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) Ensure that you select `Primary` for the node type.

Only 3-node clusters are supported for cloud deployments, so all nodes must be `Primary`.

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

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

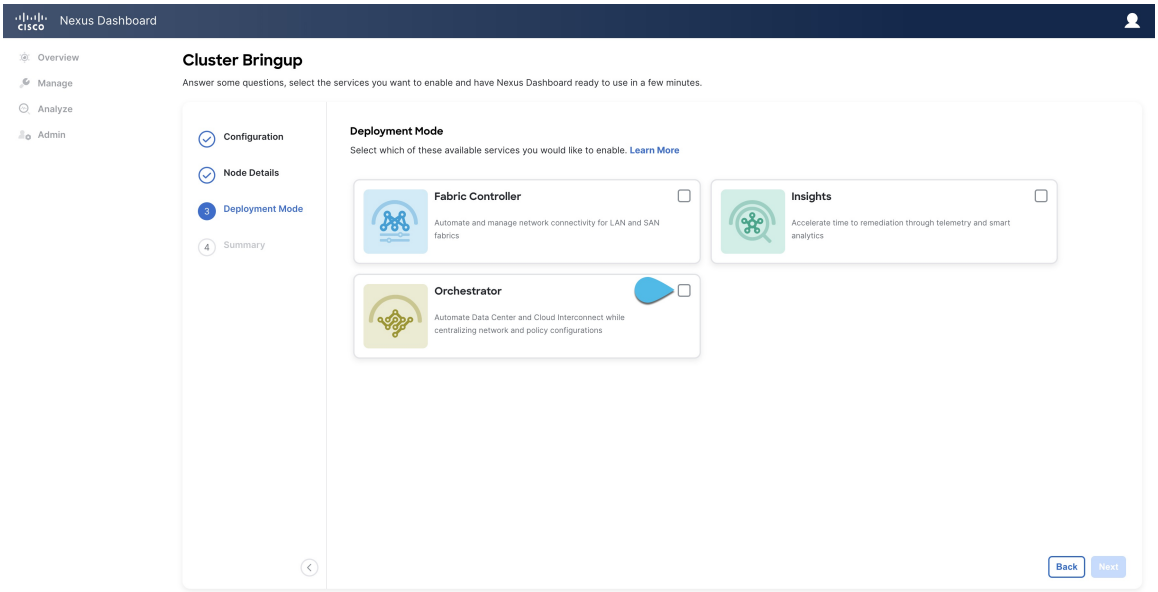
Step 16 In the **Node Details** page, click **Next** to continue.

Step 17 Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

Note
Cloud deployments support the Orchestrator service, so you must not select any other deployment modes.



b) Click **Add Persistent Service IPs/Pools** to provide one or more persistent IPs required by Insights or Fabric Controller services.

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section.

c) Click **Next** to proceed.

Step 18 In the **Summary** screen, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

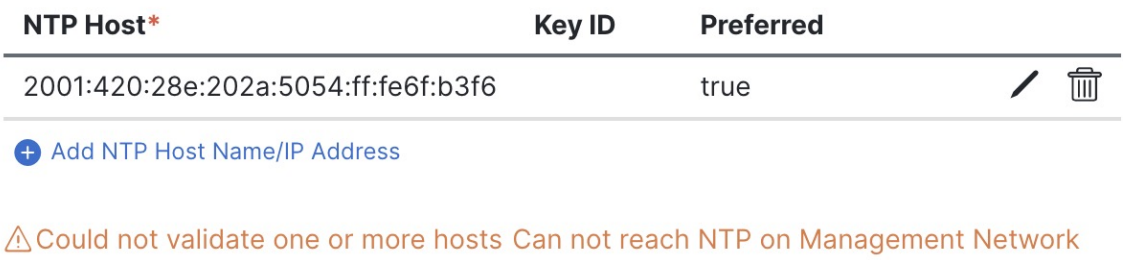
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

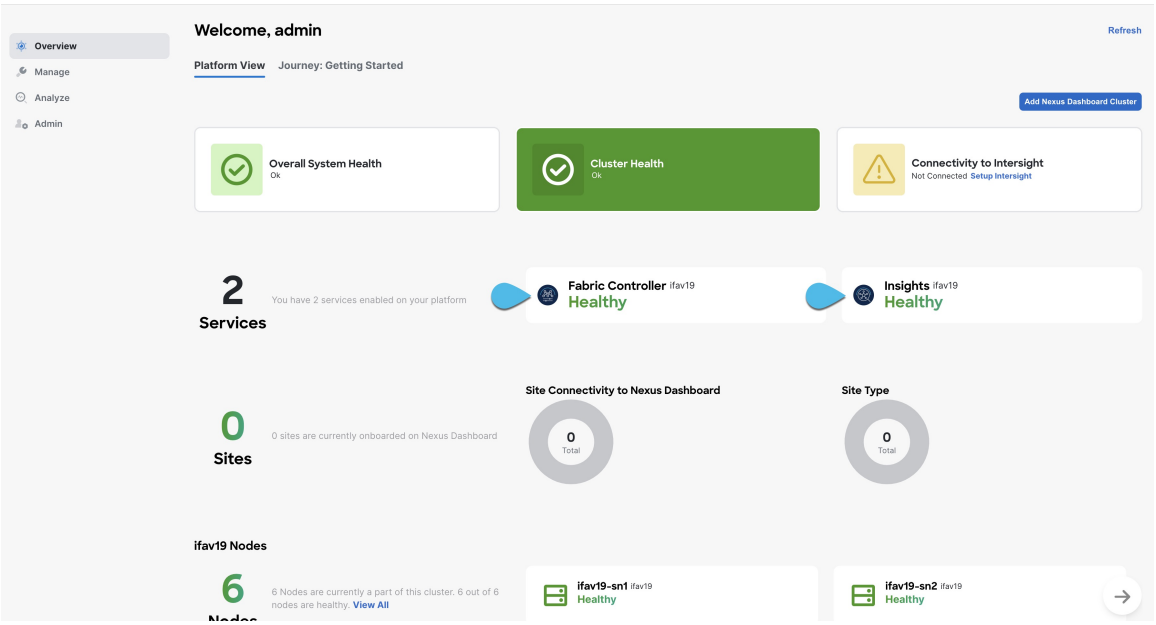
Step 19 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress
```

```
$ acs health
k8s services not in desired state - [...]
```

```
$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

Note
In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

`deploy base system services`

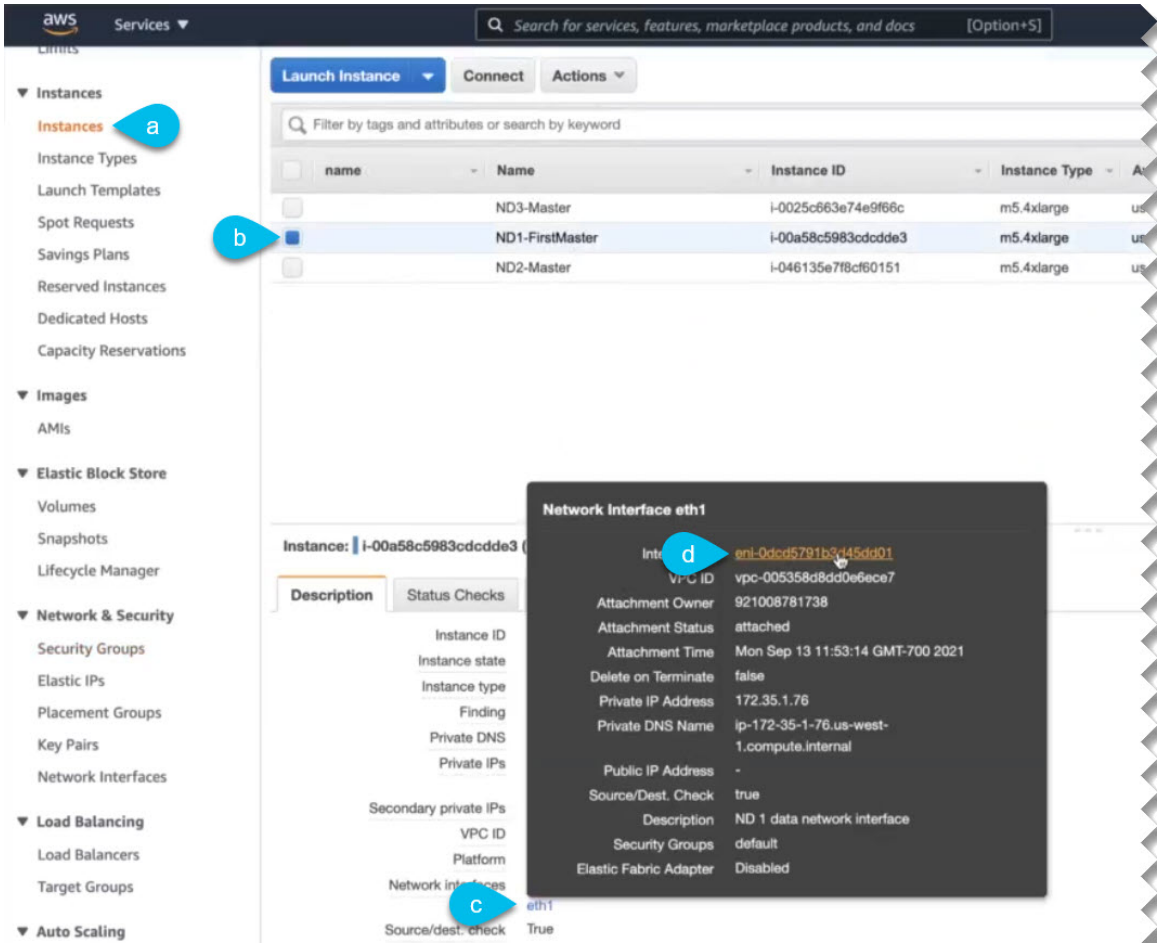
This is due to an issue with `etcd` on the node after a reboot of the pND (Physical Nexus Dashboard) cluster.

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

Step 20 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.

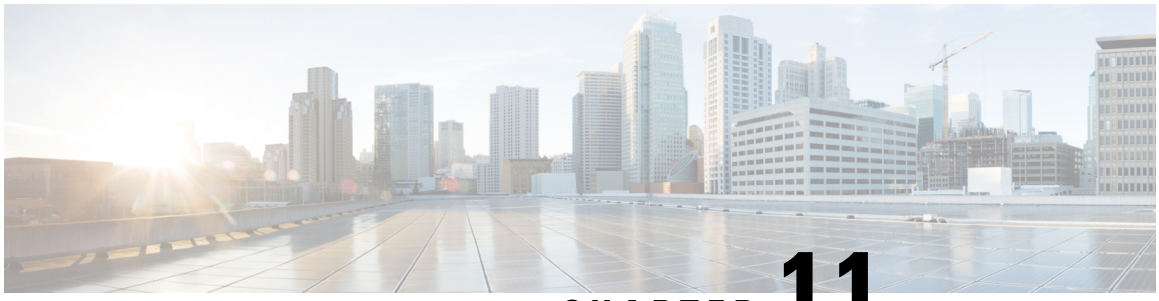
   Provide the following information for the new rule:

   - 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 11

Deploying in Microsoft Azure

- Prerequisites and Guidelines, on page 151
- Deploying Nexus Dashboard in Azure, on page 156

Prerequisites and Guidelines

Note Only the Nexus Dashboard Orchestrator service can be deployed on a cloud-hosted form factor.

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.

- Review and complete the general prerequisites described in the Deployment Overview, on page 5.
- 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.

- Ensure that the CPU family used for the Nexus Dashboard VMs supports AVX instruction set.
- 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 152 and Generating an SSH Key Pair in Windows, on page 153 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 153.

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 152.

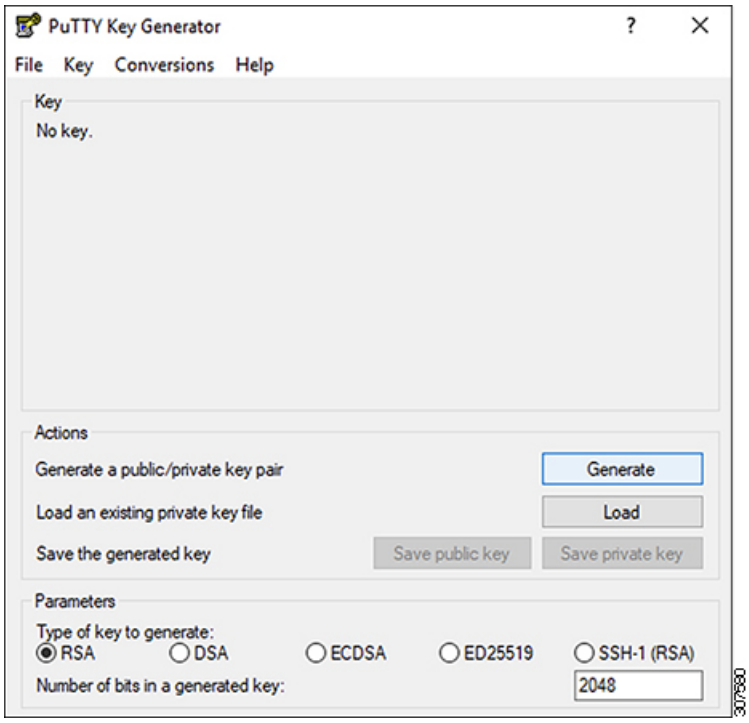
Procedure

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**.

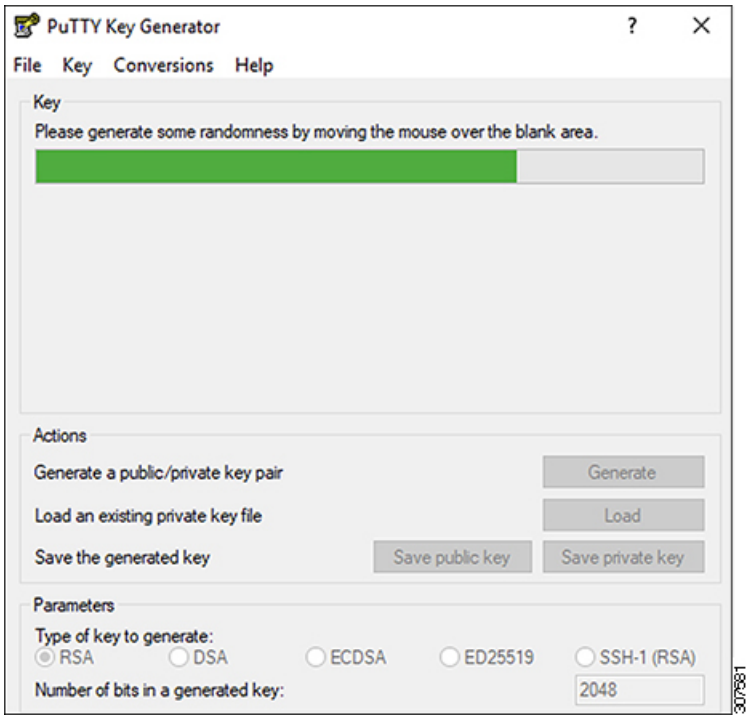
You will see a window for the PuTTY Key Generator on your screen.



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:

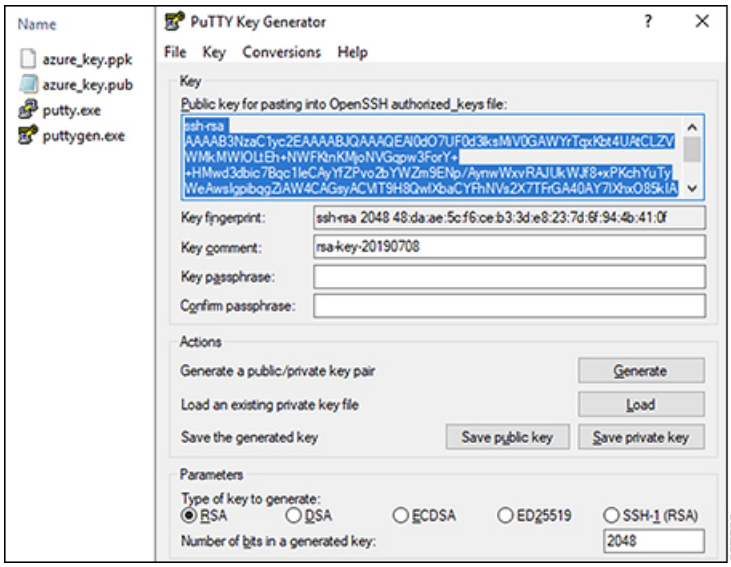
- Including the **ssh-rsa** text at the beginning of the public key.
- Excluding the following text string at the end:

```
== 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 155 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 151.

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 151.
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 151 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.

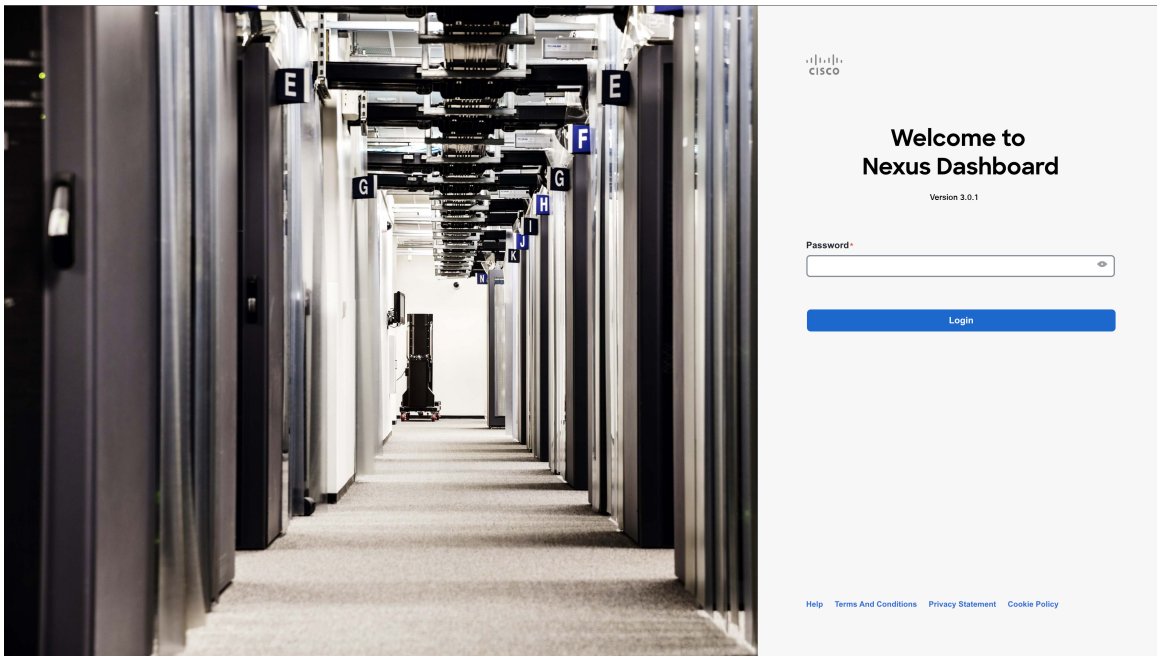
Also note which is the "first" node, which will be indicated by the node's VM name `vm-node1-<cluster-name>`. You will use this node's public IP address to complete cluster configuration.

Step 7 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.

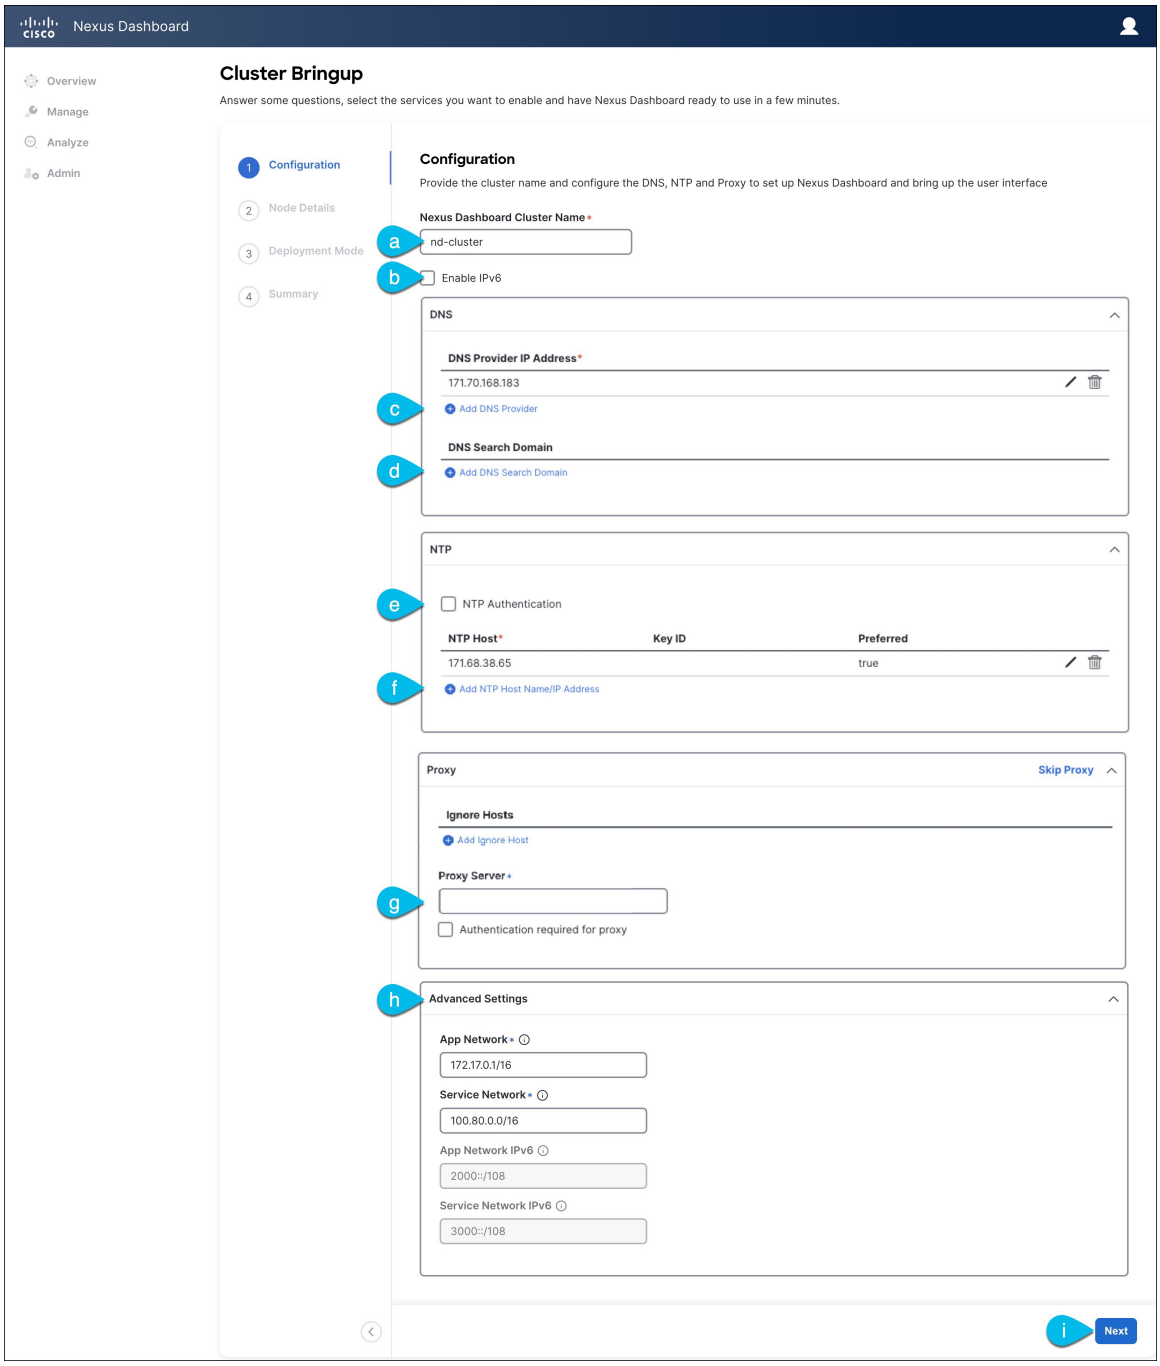
The rest of the configuration workflow takes place from the first node's GUI. You do not need to log in to or configure the other two nodes directly.

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



Step 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.

The cluster name must follow the RFC-1123 requirements.

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

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

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

d) (Optional) Click **+Add DNS Search Domain** to add a search domain.

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

e) (Optional) If you want to enable NTP server authentication, enable the **NTP Authentication** checkbox and click **Add NTP Key**.

In the additional fields, provide the following information:

- **NTP Key** – a cryptographic key that is used to authenticate the NTP traffic between the Nexus Dashboard and the NTP server(s). You will define the NTP servers in the following step, and multiple NTP servers can use the same NTP key.
- **Key ID** – each NTP key must be assigned a unique key ID, which is used to identify the appropriate key to use when verifying the NTP packet.
- **Auth Type** – this release supports `MD5`, `SHA`, and `AES128CMAC` authentication types.
- Choose whether this key is **Trusted**. Untrusted keys cannot be used for NTP authentication.

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

For the complete list of NTP authentication requirements and guidelines, see Prerequisites and Guidelines, on page 9.

f) Click **+Add NTP Host Name/IP Address** to add one or more NTP servers.

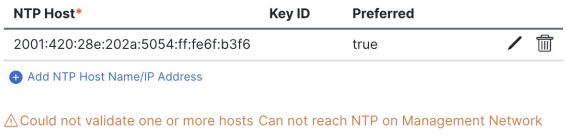
In the additional fields, provide the following information:

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

  If NTP authentication is disabled, this field is grayed out.
- Choose whether this NTP server is **Preferred**.

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

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



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

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

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

g) Provide a **Proxy Server**, then click **Validate** it.

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

You can also choose to provide one or more IP addresses communication with which should skip proxy by clicking **+Add Ignore Host**.

The proxy server must have the following URLs enabled:

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

If you want to skip proxy configuration, click **Skip Proxy**.

h) (Optional) If your proxy server required authentication, enable **Authentication required for Proxy**, provide the login credentials, then click **Validate**.

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

Under advanced settings, you can configure the following:

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

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

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

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

  Application and Services networks are described in the Prerequisites and Guidelines, on page 9 section earlier in this document.

j) 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) 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 10 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) Ensure that you select `Primary` for the node type.

Only 3-node clusters are supported for cloud deployments, so all nodes must be `Primary`.

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

Step 11 Repeat the previous step to add the 3rd node.

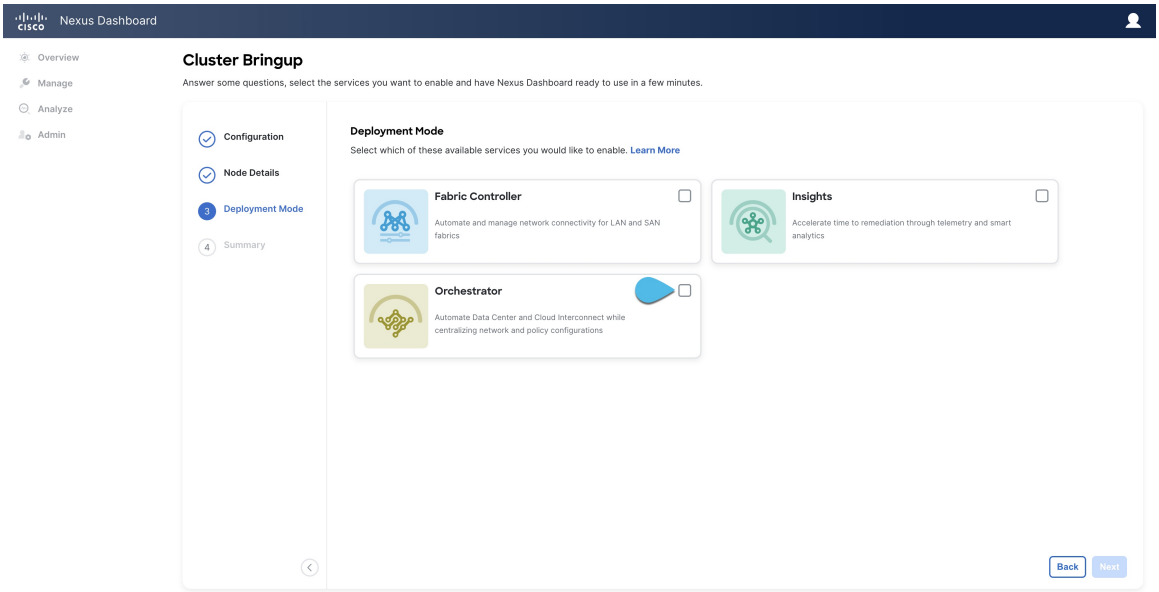
Step 12 In the **Node Details** page, click **Next** to continue.

Step 13 Choose the **Deployment Mode** for the cluster.

a) Choose the services you want to enable.

Prior to release 3.1(1), you had to download and install individual services after the initial cluster deployment was completed. Now you can choose to enable the services during the initial installation.

Note
Cloud deployments support the Orchestrator service, so you must not select any other deployment modes.



b) Click **Add Persistent Service IPs/Pools** to provide one or more persistent IPs required by Insights or Fabric Controller services.

For more information about persistent IPs, see the Prerequisites and Guidelines, on page 9 section.

c) Click **Next** to proceed.

Step 14 In the **Summary** screen, review and verify the configuration information, click **Save**, and click **Continue** to confirm the correct deployment mode and proceed with building the cluster.

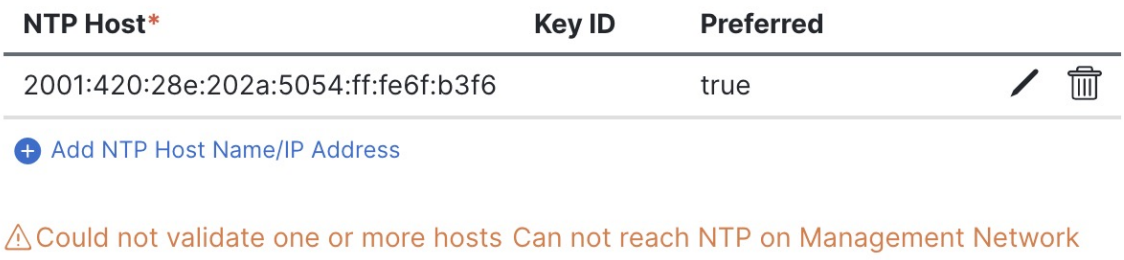
During the node bootstrap and cluster bring-up, the overall progress as well as each node's individual progress will be displayed in the UI. If you do not see the bootstrap progress advance, manually refresh the page in your browser to update the status.

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

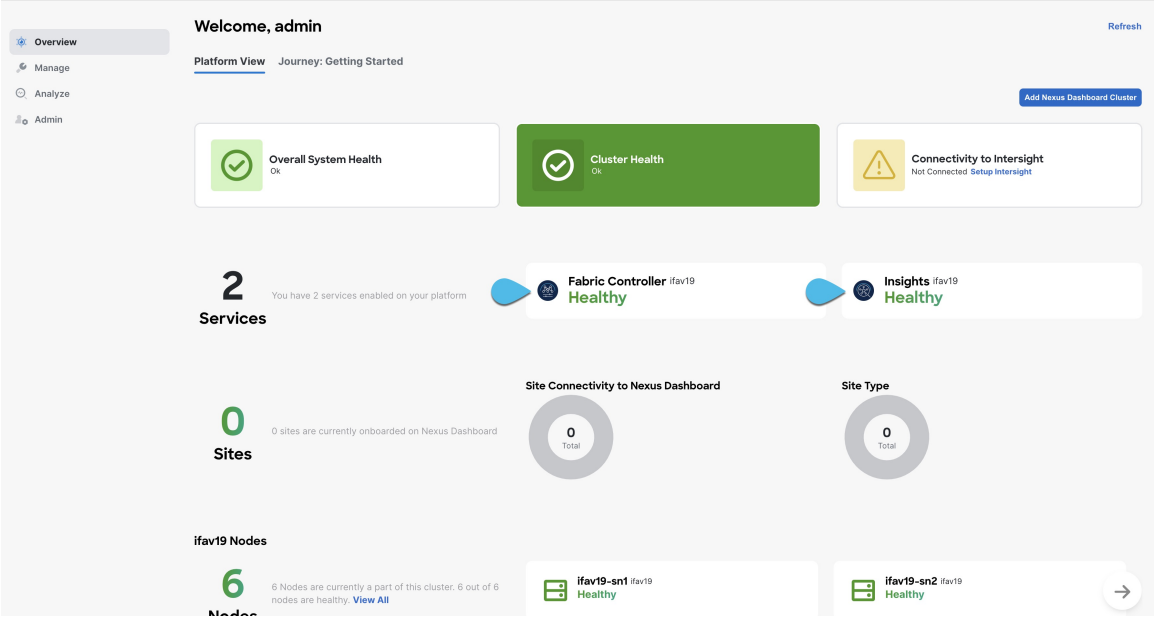
Step 15 Verify that the cluster is healthy.

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

After the cluster becomes available, you can access it by browsing to any one of your nodes' management IP addresses. The default password for the `admin` user is the same as the `rescue-user` password you chose for the first node. During this time, the UI will display a banner at the top stating "Service Installation is in progress, Nexus Dashboard configuration tasks are currently disabled":



After all the cluster is deployed and all services are started, you can check the **Overview** page to ensure the cluster is healthy:



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

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

```
$ acs health
k8s install is in-progress

$ acs health
k8s services not in desired state - [...]

$ acs health
k8s: Etcd cluster is not ready
```

- When the cluster is up and running, the following output will be displayed:

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

Note
In some situations, you might power cycle a node (power it off and then back on) and find it stuck in this stage:

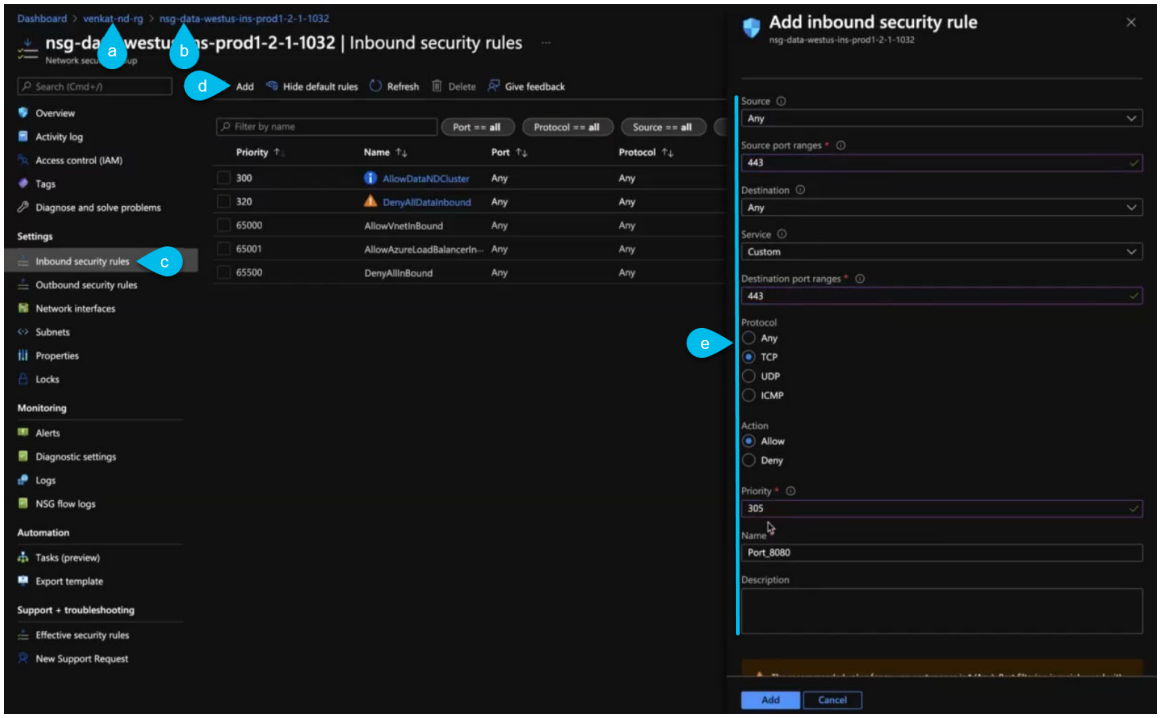
`deploy base system services`

This is due to an issue with `etcd` on the node after a reboot of the pND (Physical Nexus Dashboard) cluster.

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

Step 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`.

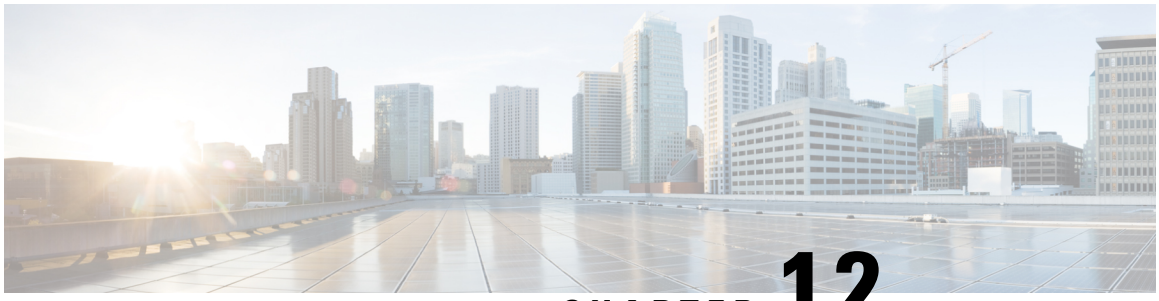
  For example, `305`.
- Provide a **Name** for the rule.

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 12

Onboarding Fabrics

- Onboarding ACI Fabrics, on page 165
- Onboarding NDFC Fabrics, on page 166
- Onboarding NX-OS Switches, on page 168

Onboarding ACI Fabrics

This section describes how to onboard one or more ACI fabrics to your Nexus Dashboard.

Before you begin

- You can on-board only 1 type of sites (ACI, NDFC, or Standalone NX-OS) within the same cluster.

  Onboarding a mix of ACI and NDFC, ACI and NX-OS, or NDFC and NX-OS within the same cluster is not supported.
- Fabric connectivity must be already configured as described in Fabric Connectivity, on page 17.
- EPG/L3Out for Nexus Dashboard data network IP connectivity must be already configured as described in Fabric Connectivity, on page 17.
- IP connectivity from Nexus Dashboard to Cisco APIC in-band IP over the data network must be already configured.
- IP connectivity from Nexus Dashboard to the leaf nodes' and spine nodes' in-band IPs over the data network must be already configured.

Procedure

Step 1 Navigate to **Manage** > **Sites**.

Step 2 Click **Add Site**.

This starts the site onboarding workflow.

Step 3 In the **Add Site** screen, choose **Controller Based Site**.

If you don't have the Insights service installed, this selection will not be visible and site onboarding defaults to this option.

Step 4 Provide site information.

- **Host Name/IP Address** — provide the IP address used to communicate with the Cisco APIC.

  Note
  When providing the address, do not include the protocol (`http://` or `https://`) as part of the URL string or site onboarding will fail.

- **User Name** and **Password** — login credentials for a user with admin privileges on the site you are adding.
- (Optional) **Login Domain** — if you leave this field empty, the site's local login is used.
- (Optional) **Validate Peer Certificate** — allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).

  Note
  You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the onboarding workflow and follow the instructions described in the "Administrative Tasks" article in the Nexus Dashboard *documentation library*; then after you have imported the certificates, add the site as described here. If you enable the Verify Peer Certificate option but don't import the valid certificate, site onboarding will fail.

- (Optional) Enable the **Use Proxy** option if connectivity to this site's controller requires the proxy.

  The proxy must be already configured in the Nexus Dashboard's **Admin Console**.

Step 5 Provide additional site **Details**.

- **Name** — a descriptive name for the site.
- **Location** — site's geographical location. This option is available only for on-premises sites.

Step 6 In the **Summary** page, verify the information and click **Save** to finish adding the site.

Onboarding NDFC Fabrics

This section describes how to onboard one or more NDFC fabrics to your Nexus Dashboard.

Note After the cluster is deployed, ensure that you configure the NDFC deployment persona by navigating to **Fabric Controller** > **System Settings** > **Feature Management** and selecting one of the supported modes.

When you create fabrics in your NDFC service, they are automatically added as sites to the Nexus Dashboard. The following steps are only required if you are on-boarding sites from a different Nexus Dashboard cluster for Fabric Controller and Insights co-location use case where each service is deployed in separate clusters.

Before you begin

- You can on-board only 1 type of sites (ACI, NDFC, or Standalone NX-OS) within the same cluster.

Onboarding a mix of ACI and NDFC, ACI and NX-OS, or NDFC and NX-OS within the same cluster is not supported.

- Fabric connectivity must be already configured as described in Fabric Connectivity, on page 17.
- Layer 3 connectivity to the fabric and switches must be already configured.
- If your cluster is deployed in AWS or Azure, you must configure inbound rules on the data interface.

Procedure

Step 1 Navigate to **Manage** > **Sites**.

Step 2 Click **Add Site**.

This starts the site onboarding workflow.

Step 3 In the **Add Site** screen, choose **Controller Based Site**.

Step 4 Provide site information.

- **Host Name/IP Address** — provide the IP address used to communicate with the Cisco NDFC.

  Note
  For NDFC sites, this must be the in-band IP address of NDFC.

  When providing the address, do not include the protocol (`http://` or `https://`) as part of the URL string or site onboarding will fail.

- **User Name** and **Password** — login credentials for a user with admin privileges on the site you are adding.
- (Optional) **Login Domain** — if you leave this field empty, the site's local login is used.
- (Optional) **Validate Peer Certificate** — allows Nexus Dashboard to verify that the certificates of hosts to which it connects (such as site controllers) are valid and are signed by a trusted Certificate Authority (CA).

  Note
  You must have the certificate for this site already imported into your Nexus Dashboard before you can add a site using this option. If you have not yet added the certificates, cancel the onboarding workflow and follow the instructions described in the "Administrator" article in the Nexus Dashboard *documentation library*; then after you have imported the certificates, add the site as described here. If you enable the Verify Peer Certificate option but don't import the valid certificate, site onboarding will fail.

Step 5 Provide additional site **Details**.

- **Name** — a descriptive name for the site.
- **Location** — site's geographical location. This option is available only for on-premises sites.

Step 6 In the **Summary** page, verify the information and click **Save** to finish adding the site.

Onboarding NX-OS Switches

This section describes how to onboard one or more standalone NX-OS switches to your Nexus Dashboard.

Note When onboarding standalone NX-OS switches without a controller (such as APIC or NDFC), the following restrictions apply:

- You can on-board only 1 type of sites (ACI, NDFC, or Standalone NX-OS) within the same cluster.

  Onboarding a mix of ACI and NDFC, ACI and NX-OS, or NDFC and NX-OS within the same cluster is not supported.
- Only Nexus Dashboard Insights service supports standalone NX-OS switches.
- Only physical Nexus Dashboard clusters support onboarding NX-OS switches.
- You must not install the NDFC service in the same Nexus Dashboard cluster where you will onboard standalone NX-OS switches.
- Before onboarding standalone NX-OS switches, you must enable "NX-OS switch discovery" in your cluster as mentioned in Step 3 below.

  Enabling NX-OS switch discovery must be done by an Admin user.
- You must also configure 10 persistent IPs (if using IPv4) and 8 IPs (if using IPv6) in the data network.

  Persistent IPs can be configured in the **Nexus Dashboard** > **Admin Console** > **System Settings** > **External Service Pools** > **Data Service IPs** page.
- You must enable NX-OS switches auto-discovery Cisco Discovery Protocol (CDP) in all NX-OS Switches.
- You must configure Nexus Dashboard reachability to NX-OS switches' management network because NX-OS switches auto-discovery uses the switches' management interfaces.

  You must configure Nexus Dashboard data network reachability to NX-OS switches' in-band network.

Procedure

Step 1 Navigate to **Manage** > **Sites**.

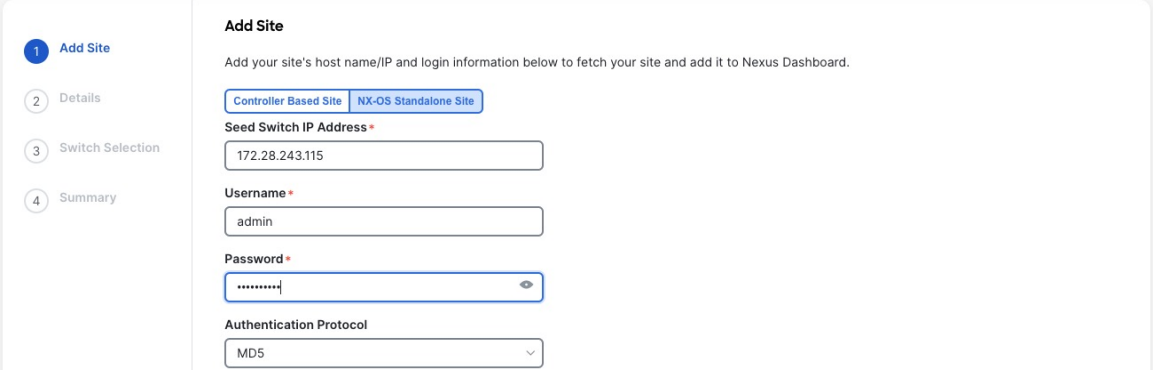
Step 2 Click **Add Site**.

This starts the site onboarding workflow.

Step 3 In the **Add Site** screen, choose **NX-OS Standalone Site**.

Note
If this is the first time you are onboarding NX-OS switches without a controller, click **Enable NXOS Discovery**.

Step 4 Provide site information.



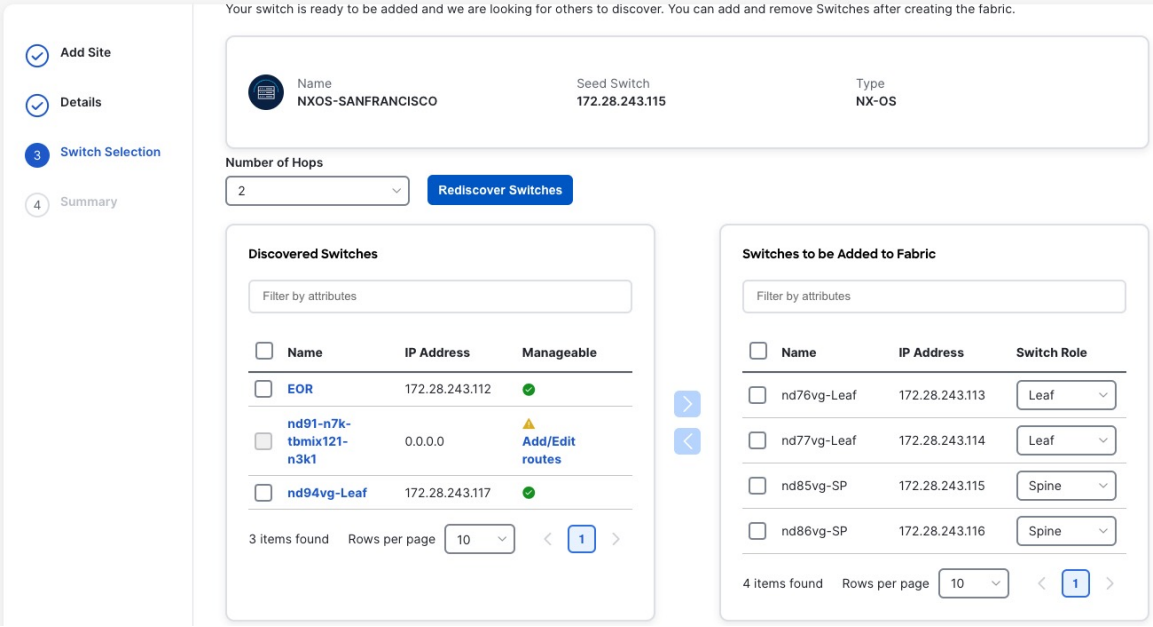
- **Seed Switch IP Address** — provide the IP address of the seed switch used to discover other switches in the site.
- **Username** and **Password** — login credentials on the seed switch.

Step 5 Provide additional site **Details**.

- **Name** — a descriptive name for the site.
- **Location** — site's geographical location. This option is available only for on-premises sites.

Step 6 In the **Switch Selection** page, select one or more switches to add to the site.

By default, the switch discovery process will show switches that are 2 hops away from the seed switch. You can change the default setting using the **Number of Hops** dropdown and clicking **Rediscover Switches**.



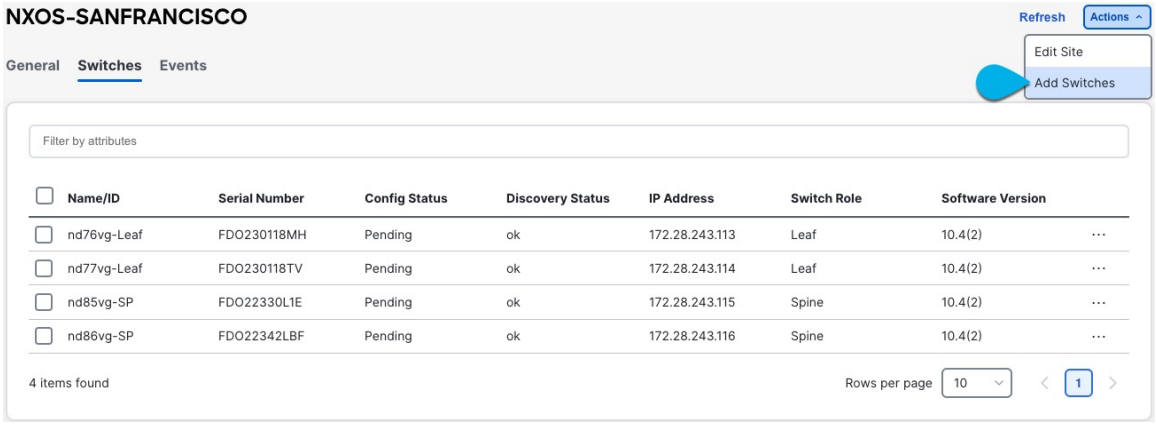
After the switches are discovered, simply select all the switches you want to add to the site from the list on the left and click on the right arrow to move them to right-hand list.

The switches are added with the default `Leaf` role, but you can change it to other roles as required, then click **Next** to continue.

Step 7 In the **Summary** page, verify the information and click **Save** to finish adding the site.

Step 8 (Optional) Add switches to an existing standalone NX-OS site.

After you first add the site, you can **Add Switches** by selecting the site in the GUI:

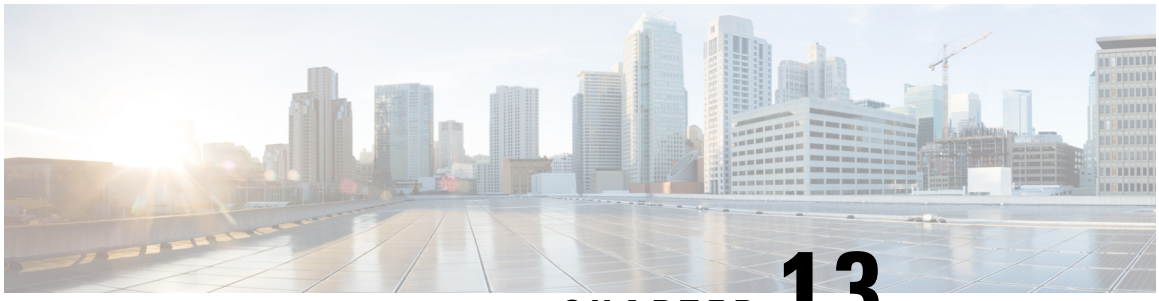




PART III

Upgrading or Migrating to This Release

- Upgrading Existing ND Cluster to This Release, on page 173
- Migrating From DCNM to NDFC, on page 187



CHAPTER 13

Upgrading Existing ND Cluster to This Release

- Prerequisites and Guidelines, on page 173
- Supported Upgrade Paths, on page 175
- Upgrading Nexus Dashboard, on page 178
- Troubleshooting Upgrades, on page 183
- Resolving Nexus Dashboard Orchestrator Configuration Issues, on page 183

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.
- Note that after you upgrade to this release, you can no longer change the number of services enabled in your cluster.

  Beginning with release 3.1.1, each cluster has a "deployment mode" which defines the combination of enabled services and cannot be changed after the cluster is deployed or upgraded. In other words, you will not be able to add or remove services after upgrading to this release without redeploying the cluster. If you were planning to add or remove services in your cluster, we recommend doing so before upgrading to this release.

Note In some cases, a deployment mode supported in this release may have not been supported in a prior release (for example, cohosting Insights and Orchestrator is not supported in virtual clusters in release 3.0.1). In this case, if you have a single service (such as Insights) deployed in your current cluster but you plan to add another service (such as Orchestrator) after the upgrade:

1. Disable existing Insights service in your current cluster.
2. Install the additional Orchestrator service in your current cluster.
3. Enable the Orchestrator service in your current cluster.

   At this point you would have both Insights and Orchestrator in your current cluster with Insights disabled and Orchestrator enabled. Note that you must not enable both services at the same time if it is not a supported configuration in your current release.
4. Disable the Orchestrator service and proceed with the upgrade.

- If you are running Nexus Dashboard Insights service in a 4-node or 5-node physical cluster, you can simply upgrade the cluster and the service to this release as you typically would and continue using the 4-node or 5-node cluster.

  This release of Nexus Dashboard with Nexus Dashboard Insights supports only 3-node and 6-node profiles for greenfield deployments. However, if you are upgrading an existing 4-node or 5-node cluster from an earlier release without changing your current scale, you can continue using it with this release.
- 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" article in the Nexus Dashboard *documentation library*.
- If you are upgrading a virtual Nexus Dashboard cluster deployed in Linux KVM, you must enable the `Copy host CPU configuration` option in the **Virtual Machine Manager** UI.

  This release supports CentOS 7.9 or Red Hat Enterprise Linux 8.6 with the following Kernel and KVM versions:

  - For CentOS 7.9, Kernel version `3.10.0-957.el7.x86_64` and KVM version `libvirt-4.5.0-23.el7_7.1.x86_64`
  - For RHEL 8.6, Kernel version `4.18.0-372.9.1.el8.x86_64` and KVM version `libvert 8.0.0`
- 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.

  This release supports VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3, 8.0.



Note If you need to upgrade the ESX server, you must do that before upgrading your Nexus Dashboard. ESX upgrades are outside the scope of this document, but in short:

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 this document.

- 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`.

- If you are upgrading from release 3.0(1), you must disable all services running in the cluster before upgrading to this release.

  If you are upgrading from release 3.1(1), you do not need to disable the services.

  Note that because of the unified installation introduced in release 3.1(1), all of your existing services will be automatically upgraded to the versions compatible with this Nexus Dashboard release while preserving their configuration. The services will also be automatically re-enabled after the upgrade is completed.

- 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.
- Ensure that no configuration changes are made to the cluster, such as adding `secondary` 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.

Supported Upgrade Paths

Prior to Nexus Dashboard release 3.1(1), the platform software and individual services could be upgraded independent of each other but within a specific set of supported version combinations. Because of that, prior releases' documentation provided specific upgrade paths across multiple releases to ensure compatibility between all components at each upgrade stage.

This requirement has now been removed due to a much tighter coupling between the Nexus Dashboard and individual services with only a single version of each service compatible with each version of the platform.

As a result, as long as you are currently on the minimum required version of the Nexus Dashboard software (listed in Prerequisites and Guidelines, on page 173), you can upgrade both the platform and all currently enabled services directly to this release.

The following table provides a few example scenarios for specific deployment combinations:

Table 22:

Current Nexus Dashboard Release	Compatible Services (depending on form factor and cluster size, you may have one or more of these services currently enabled)	Upgrade Workflow
3.0(1)	Fabric Controller: 12.1(3) Orchestrator: 4.2(x) Insights: 6.3(1)	Upgrade directly to release 3.1.1 as described in the following section. All services will be automatically upgraded along with the platform.
2.3(2d)	Fabric Controller: 12.1(2) Orchestrator: 4.1(x) Insights: 6.2(x)	**1.** If you are running Orchestrator release 4.1(1) or earlier, upgrade the Orchestrator service to release 4.1(2) first. If you do not have the Orchestrator service enabled, skip this step. **Note** When upgrading from Nexus Dashboard release 2.3.2d with Orchestrator service enabled, the Orchestrator service must be version 4.1(2) or later to upgrade directly to Nexus Dashboard 3.1(1). **2.** Upgrade from release 2.3(2d) to release 3.1(1). All services will be automatically upgraded along with the platform.

Current Nexus Dashboard Release	Compatible Services (depending on form factor and cluster size, you may have one or more of these services currently enabled)	Upgrade Workflow
2.3(2b)	Fabric Controller: 12.1(2) Orchestrator: 4.1(x) Insights: 6.2(x)	1. Upgrade the Nexus Dashboard platform **only** to release 2.3(2d) as described in *Nexus Dashboard Deployment Guide, Release 2.3.x* 2. If you are running Orchestrator release 4.1(1) or earlier, upgrade the Orchestrator service to release 4.1(2) first. If you do not have the Orchestrator service enabled, skip this step. **Note** When upgrading from Nexus Dashboard release 2.3.2d with Orchestrator service enabled, the Orchestrator service must be version 4.1(2) or later to upgrade directly to Nexus Dashboard 3.1(1). 3. Upgrade from release 2.3(2d) to release 3.1(1). All services will be automatically upgraded along with the platform.

Current Nexus Dashboard Release	Compatible Services (depending on form factor and cluster size, you may have one or more of these services currently enabled)	Upgrade Workflow
Any release prior to 2.3(2b)	Fabric Controller: 12.1(1) or earlier Orchestrator: 4.0(2) or earlier Insights: 6.1(3) or earlier	1. Upgrade the Nexus Dashboard platform to release 2.3(2d) as described in *Nexus Dashboard Deployment Guide, Release 2.3.x* 2. Upgrade the individual services to any version compatible with Nexus Dashboard release 2.3(2d) as you previously would using the upgrade paths or minimum version requirements described in the service-specific documentation: • Fabric Controller • Orchestrator • Insights 3. Upgrade from release 2.3(2d) to release 3.1(1). All services will be automatically upgraded along with the platform.

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 3.0(1). 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 173

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 Nexus Dashboard image for your target release.

Note
The upgrade process is the same for all Nexus Dashboard form factors and uses the Nexus Dashboard ISO image (`nd-dk9.<version>.iso`). In other words, even if you used the virtual form factors (such as the ESX `.ova`) or a cloud provider's marketplace for initial cluster deployment, you must still use the `.iso` image for upgrades.

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

We recommend hosting the image on a 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, which can significantly speed up the process.

Step 2 Log in to your current Nexus Dashboard's **Admin Console** as an `Administrator` user.

Step 3 Disable any existing services installed in the cluster.

Note
You must disable all services before you upgrade the cluster. You must not delete any services after disabling them. The disabled services will be automatically re-activated once the upgrade process is complete.

a) From the main navigation menu, select **Services** (release 2.3.2) or **Operate** > **Services** (release 3.0.1 or later).

b) In the service's tile, click the actions (**...**) menu and choose **Disable**.

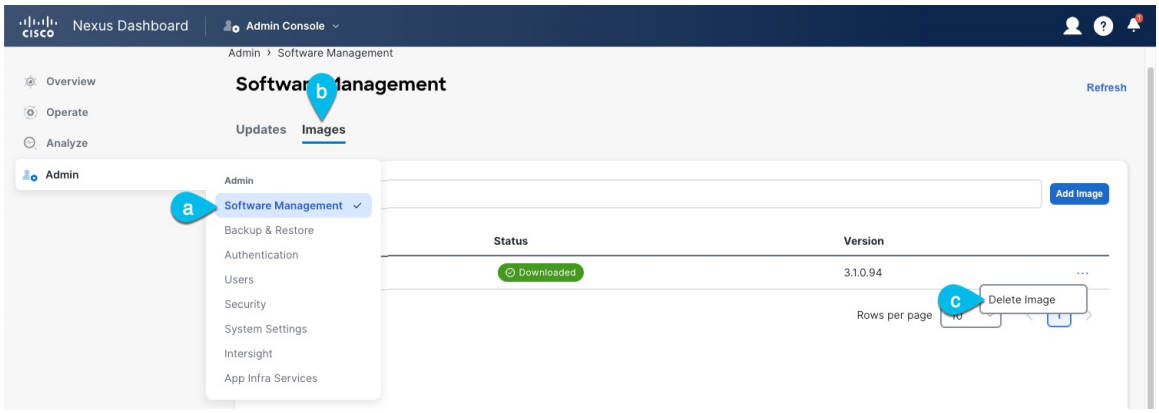
c) Repeat this step for all services deployed in the cluster.

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.

Note
In release 2.3.2, this page is located under **Operations** > **Firmware Management** instead.



a) Navigate to **Admin** > **Software 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) In the **Admin** > **Software Management** page's **Images** tab, click **Add Image**.
b) In the **Add Software Image** window, select whether the image is **Local** on your machine or **Remote** on a web server.
c) Click **Choose file** or provide the **URL** to the image you downloaded in the first step.
d) Click **Upload** to add the image.
e) Wait for the image status to change to `Downloaded`.

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 6 Set up the upgrade.

a) Navigate to **Admin** > **Software Management**.

Note
In release 2.3.2, this page is located under **Operations** > **Firmware Management** instead.

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.

d) In the **Setup** > **Version selection** screen, select the firmware version you uploaded and click **Next**.
e) In the **Setup** > **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.

f) 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 **Software Management** screen and click **Continue**.

This step may take up to 20 minutes, during which the upgrade 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.

Step 7 Activate the new image.

If you never navigated away from the upgrade screen, simply click **Activate** to activate the new image, otherwise:

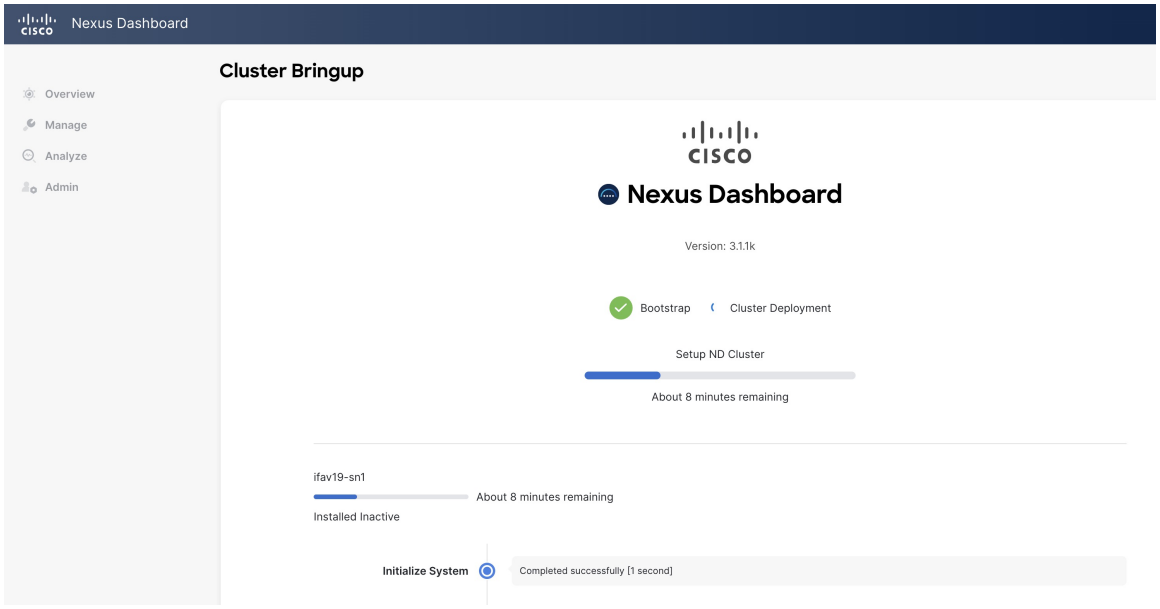
a) Navigate back to the **Admin** > **Software Management** screen.

In release 2.3.2, this page is located under **Operations** > **Firmware Management** instead.

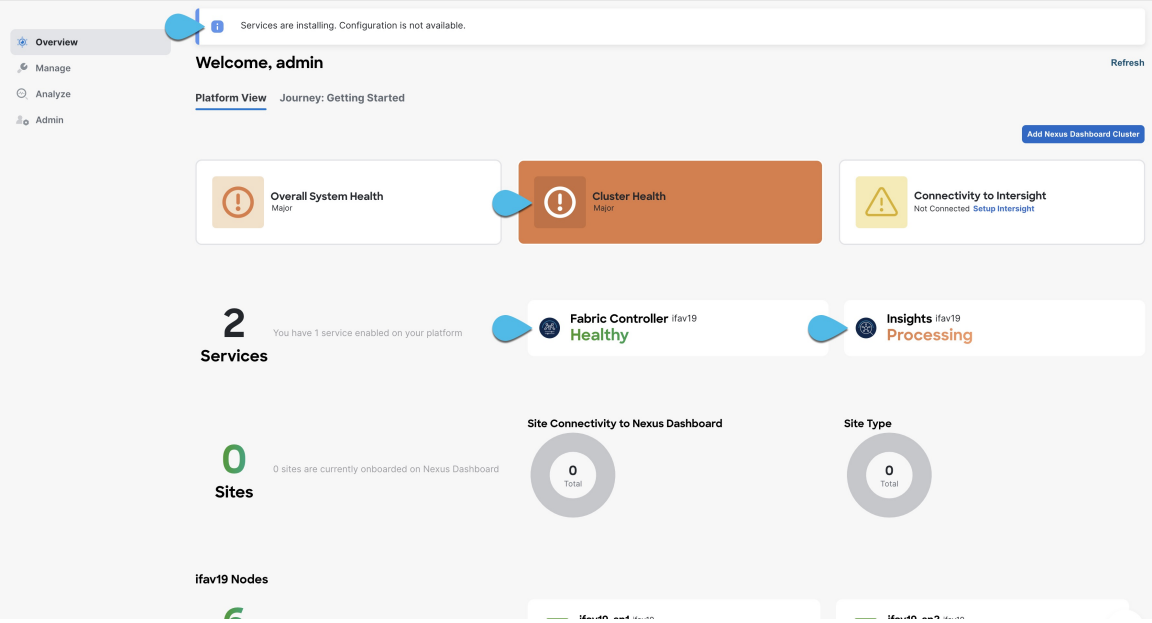
b) In the **Last Update Status** tile, click **Continue**.

c) In the **Firmware Update** > **Install** screen, click **Activate**.

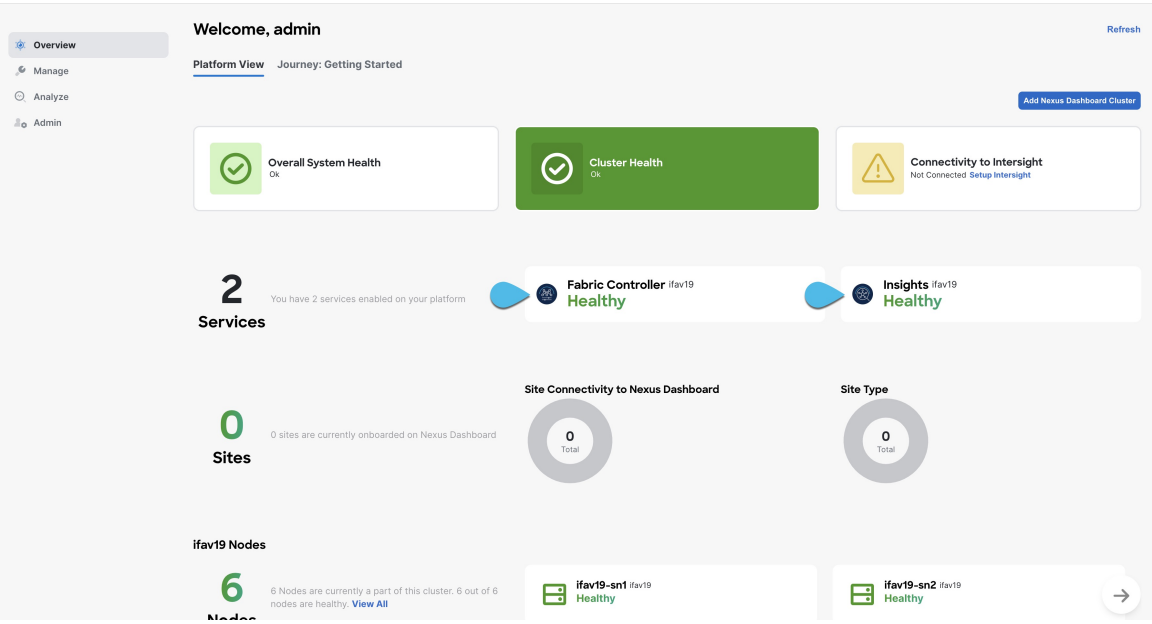
After you click **Activate**, the cluster will bring down the background services, which may take several minutes, and then restart. Note that all nodes will restart simultaneously during the activation stage and it may take up to 20 additional minutes for all the cluster services to start and the GUI to become available after the nodes restart:



You can check the **Overview** page for the progress and service status:



Once the upgrade is finalized, any existing services will show as `Healthy` in the **Overview page**:



Step 8 (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" article in the Nexus Dashboard *documentation library*.

Troubleshooting Upgrades

After all the nodes restart during new image activation stage described in the previous section, you may log in to the GUI to check the status of the upgrade workflows. Initially, you can see the bootstrap process similar to the initial cluster deployment and once the nodes come up, you can see additional information about service activation in the GUI's **Overview** page.

In case the upgrade fails for any reason, the GUI will display the error and additional workaround steps. However, if you are unable to resolve the issues through the GUI, you can re-try the upgrade manually by logging in to the nodes as the `rescue-user` and running the commands described in this section.

Procedure

Step 1 Log in to all of your Nexus Dashboard cluster nodes as the `rescue-user`.

You will need to perform recovery commands on all nodes simultaneously, so log in to each node before continuing with the next step.

Step 2 Ensure that you are logged in to all nodes as the `rescue-user`.

Step 3 Run the required commands depending on the specific scenarios.

If the upgrade failed where one or more nodes did not reboot and are still running a older release, you can attempt a manual upgrade; however, you must contact Cisco TAC for assistance with the manual upgrade because this requires root access that only Cisco TAC can generate.

Step 4 Wait for the installer to finish on all nodes.

Step 5 Restart all nodes simultaneously using the `acs reboot` command.

After the nodes restart, you can log into the UI to see the bootstrap progress similar to regular UI-based upgrades.

Step 6 After the node upgrade tasks are completed, verify that the nodes are healthy and you can log into the UI.

Once the bootstrap process completes, you can view the Nexus Dashboard UI as you typically would.

You can check the **Overview** page for overall system health and the **Manage** > **Software Management** page to see the current `Running` version.

In addition, verify the services' status in the **Analyze** > **Service Status** page.

Resolving Nexus Dashboard Orchestrator Configuration Issues

The following topics provide additional procedures to help you resolve Nexus Dashboard Orchestrator configuration issues after an upgrade.

Resolve Configuration Inconsistencies In Templates

Check for and resolve any configuration inconsistencies in your templates. The inconsistencies will be displayed as a warning message in a bulleted list.

Procedure

Step 1 Navigate to your schemas and check for any templates marked with a warning. Click on the ellipsis (…) at the restored file location to view the error message.

Step 2 Open a template that contains a warning and resolve any configuration inconsistencies.

Step 3 Save and deploy the template.

Step 4 Repeat these steps for all templates that have a warning.

Resolve Configuration Drifts

In some cases you may run into a situation where the configuration actually deployed in the site's controller is different from the configuration defined in the Nexus Dashboard Orchestrator. These configuration discrepancies are referred to as **Configuration Drifts** and are indicated by an `Out of Sync` warning next to the site name in the template view page

After upgrading your Nexus Dashboard Orchestrator and restoring the previous configuration backup, we recommend that you check for and resolve any configuration drifts that were not automatically resolved by the upgrade process as described in this section.



Note Deploying any templates before resolving configuration drifts would push the configuration defined in the Orchestrator and overwrite the values defined in the fabrics' controllers.

Procedure

Step 1 In your Nexus Dashboard Orchestrator, navigate to **Manage** > **Tenant Templates**.

Step 2 Choose the **Applications** tab.

Step 3 Select the first schema and check its templates for configuration drifts.

You will repeat the following steps for every schema and template in your deployment

You can check for configuration drifts in one of the following two ways:

- Check the template deployment status icon for each site to which the template is assigned.
- Select the template and click **Deploy to sites** to bring up the configuration comparison screen to check which objects contain configuration drifts.

Step 4 If the template contains a configuration drift, resolve the conflicts.

For more information about configuration drifts, check the "Configuration Drifts" chapter in the *Cisco Nexus Dashboard Orchestrator Configuration Guide for ACI Fabrics*.

a) Close the template deployment dialog to return to the Schema view.

   Deploying any templates at this point would push the values in the Orchestrator database and overwrite any existing settings in the fabrics.

b) From the template's **Actions** menu, select **Reconcile Configuration Drifts**.

   The **Drift Reconciliation** wizard opens.

c) In the **Drift Reconciliation** screen, compare the template-level configurations for each site and choose the one you want.

   Template-level properties are common across all sites associated to the template. You can compare the template level properties defined on Nexus Dashboard Orchestrator with the configuration rendered in each site and decide what should become the new configuration in the Nexus Dashboard Orchestrator template. Selecting the site configuration will modify those properties in the existing Nexus Dashboard Orchestrator template, whereas selecting the Nexus Dashboard Orchestrator configuration will keep the existing Nexus Dashboard Orchestrator template settings as is

d) Click **Go to Site Specific Properties** to switch to site-level configuration.

   You can choose a site to compare that specific site's configuration. Unlike template-level configurations, you can choose either the Nexus Dashboard Orchestrator-defined or actual existing configurations for each site individually to be retained as the template's site-local properties for that site.

   Even though in most scenarios you will make the same choice for both template-level and site-level configuration, the drift reconciliation wizard allows you to choose the configuration defined in the site's controller at the "Template Properties" level and the configuration defined in Nexus Dashboard Orchestrator at the "Site Local Properties" level or vice versa.

e) Click **Preview Changes** to verify your choices.

   The preview will display full template configuration adjusted based on the choices picked in the **Drift Reconciliation** wizard

f) Save the schema.

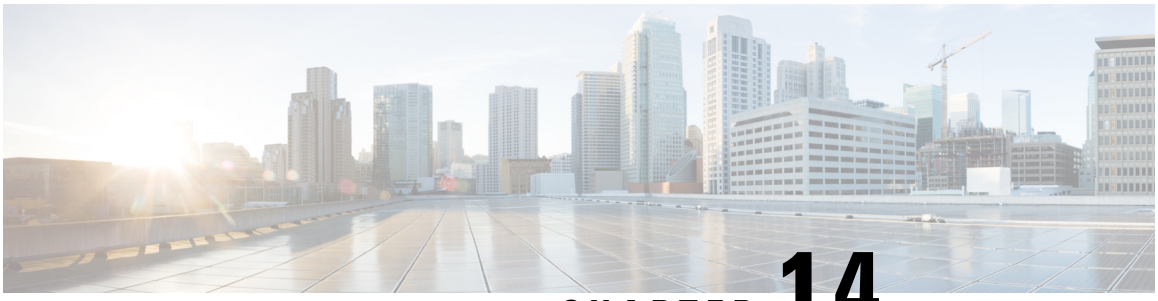
g) Click **Deploy to sites** to deploy the configuration and finish reconciling the drift for that template

Step 5 Repeat the above steps for every schema and template in your Nexus Dashboard Orchestrator.

Step 6 Check audit logs to verify that all templates have been re-deployed.

You can view the audit logs in the **Operations** tab.

Audit Logs page and confirm that all templates show as `Redeployed` to ensure that full re-deployment successfully completed.



CHAPTER 14

Migrating From DCNM to NDFC

- Prerequisites and Guidelines, on page 187
- Migrate Existing DCNM Configuration to NDFC, on page 189

Prerequisites and Guidelines

Note If you are already running Nexus Dashboard with Fabric Controller service, skip this section and upgrade as described inUpgrading Existing ND Cluster to This Release, on page 173 instead.

Upgrading from DCNM 11.5(4) consists of the following workflow:

1. Ensure you complete the prerequisites and guidelines described in this section.
2. Back up your existing configuration using a migration tool specific to the target NDFC release.
3. Deploy a brand new Nexus Dashboard cluster with Fabric Controller (NDFC) service.

Note

- Unlike in previous releases where you had to install the service and enable it after the cluster was already deployed, in this release you enable the service during initial cluster deployment due to the introduction of the unified installation.
- Deploy the Nexus Dashboard cluster only with the Fabric Controller (NDFC) service; do not deploy with other services (for example, with Fabric Controller + Insights). Doing so will cause the migration to fail when you restore the configuration backup in the next step with the message `Backup file is not compatible with this deployment mode.`

4. Restore the configuration backup you created in step 1.

Note Before you proceed with the upgrade, validate each fabric's credentials.

For LAN fabrics, navigate to the **Web UI** > **Administration** > **Credentials Management** > **LAN Credentials** page, select each fabric, and choose **Validate** to validate credentials.

For SAN fabrics, navigate to the **Web UI** > **Administration** > **Credentials Management** > **SAN Credentials** page, select each fabric, and choose **Validate** to validate credentials.

Persona Compatibility

By using the appropriate Upgrade Tool, you can restore data that is backed up from DCNM Release 11.5(4) on a newly deployed Nexus Dashboard Fabric Controller for the personas as mentioned in the following table:

Backup from DCNM 11.5(4)	Persona Enabled in NDFC After Upgrade
DCNM 11.5(4) LAN Fabric Deployment on OVA/ISO/SE	Fabric Controller + Fabric Builder
DCNM 11.5(4) PMN Deployment on OVA/ISO/SE	Fabric Controller + IP Fabric for Media (IPFM)
DCNM 11.5(4) SAN Deployment on OVA/ISO/SE	SAN Controller
DCNM 11.5(4) SAN Deployment on Linux	SAN Controller
DCNM 11.5(4) SAN Deployment on Windows	SAN Controller

Feature Compatibility Post Upgrade

The following table lists caveats associated with features that are restored from DCNM 11.5(4) backup after upgrading.

Note SAN Insights and VMM Visualizer features are not enabled after restore; you can choose to enable them in the **Settings** > **Feature Management** page of the Nexus Dashboard Fabric Controller UI.

Feature in DCNM 11.5(4)	Upgrade Support
Nexus Dashboard Insights configured Refer to Cisco Nexus Dashboard User Guide for more information.	Supported
Container Orchestrator (K8s) Visualizer	Supported
VMM Visibility with vCenter	Supported
Nexus Dashboard Orchestrator configured	Not Supported
Preview features configured	Not supported
LAN switches in SAN installations	Not supported
Switches discovered over IPv6	Not supported

Feature in DCNM 11.5(4)	Upgrade Support
DCNM Tracker	Not supported
Fabric Backups	Not supported
Report Definitions and Reports	Not supported
Switch images and Image Management policies	Not supported
SAN CLI templates	Not carried over from 11.5(4)
Switch images/Image Management data	Not carried over from 11.5(4)
Slow drain data	Not carried over from 11.5(4)
Infoblox configuration	Not carried over from 11.5(4)
Endpoint Locator configuration	You must reconfigure Endpoint Locator (EPL) post upgrade. However, historical data is retained up to a maximum size of 500 MB.
Alarm Policy configuration	Not carried over from 11.5(4)
Performance Management data	CPU/Memory/Interface statistics up to 90 days is restored post upgrade.
Temperature data	Temperature data is not saved in the backup and as a result is not restored after the migration. You must re-enable temperature data collection after the migration.

Migrate Existing DCNM Configuration to NDFC

This section describes how to back up your existing DCNM 11.5(4) configuration, deploy a new Nexus Dashboard cluster, and restore the configuration to finish the migration.

Procedure

Step 1 Download the upgrade tool.

a) Navigate to the NDFC download page..

https://software.cisco.com/download/home/281722751/type/282088134/

b) In the **Latest Releases** list, choose the target release.

c) Download the upgrade tool appropriate for your deployment type.

DCNM 11.5(4) deployment type	Upgrade Tool File Name
ISO/OVA	`DCNM_To_NDFC_12_2_1_Upgrade_Tool_OVA_ISO.zip`

DCNM 11.5(4) deployment type	Upgrade Tool File Name
Linux or Windows	DCNM_To_NDFC_12_2_1_Upgrade_Tool_LIN_WIN.zip

d) Copy the upgrade tool image to your existing DCNM 11.5(4) server using the **sysadmin** account.

Step 2 Extract the archive and validate the signature for Linux/Windows deployments.

Note
If you are using the ISO/OVA archive, skip to the next step.

a) Ensure that you have Python 3 installed.

```
$ python3 --version
Python 3.9.6
```

b) Extract the downloaded archive.

```
# unzip DCNM_To_NDFC_12_2_1_Upgrade_Tool_LIN_WIN.zip
Archive:  DCNM_To_NDFC_12_2_1_Upgrade_Tool_LIN_WIN.zip
 extracting: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
 extracting: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip.signature
  inflating: ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
  inflating: cisco_x509_verify_release.py3
```

c) Validate signature.

Inside the ZIP archive, you will find the upgrade tool as well as the signature file. Use the following commands to validate the upgrade tool:

```
# ls -l
total 4624
-rw-rw-r-- 1 root root    1422 Aug 11  2023 ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
-rwxr-xr-x 1 root root   16788 Feb 26 15:57 cisco_x509_verify_release.py3
-rw-r--r-- 1 root root 2344694 Feb 27 07:51 DCNM_To_NDFC_12_2_1_Upgrade_Tool_OVA_ISO.zip
-rwxr-xr-x 1 root root 2359065 Feb  2 09:19 DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
-rw-rw-r-- 1 root root     256 Feb 26 16:54 DCNM_To_NDFC_Upgrade_Tool_OVA_ISO.signature

# ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM -i
DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip -s DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip.signature -v dgst
 -sha512

Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Successfully verified the signature of DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip using
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
```

d) Once the validation script signature is verified, extract the script itself.

```
# unzip DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
Archive:  DCNM_To_NDFC_Upgrade_Tool_LIN_WIN.zip
   creating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/log4j2.properties
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/DCNMBackup.sh
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/DCNMBackup.bat
   creating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/jarchivelib-0.7.1-jar-with-dependencies.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/bcprov-jdk15on-1.68.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/not-going-to-be-commons-ssl-0.3.20.jar
```

```
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/jnm.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/slf4j-simple-1.7.21.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/log4j.properties
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/dcnmbackup.jar
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/sequences.info.oracle
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/tables.info.postgres
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/sequences.info.postgres
  inflating: DCNM_To_NDFC_Upgrade_Tool_LIN_WIN/jar/tables.info.oracle
```

Step 3 Extract the archive and validate the signature for ISO/OVA deployments.

Note
If you are using the Linux/Windows archive, skip to the next step.

a) Extract the downloaded archive.

```
# unzip DCNM_To_NDFC_12_2_1_Upgrade_Tool_OVA_ISO.zip
Archive:  DCNM_To_NDFC_12_2_1_Upgrade_Tool_OVA_ISO.zip
inflating: DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
extracting: DCNM_To_NDFC_Upgrade_Tool_OVA_ISO.signature
inflating: ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
inflating: cisco_x509_verify_release.py3
```

b) Validate signature.

Inside the ZIP archive, you will find the upgrade tool as well as the signature file. Use the following commands to validate the upgrade tool:

```
$ ./cisco_x509_verify_release.py3 -e ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM -i
DCNM_To_NDFC_Upgrade_Tool_OVA_ISO -s DCNM_To_NDFC_Upgrade_Tool_OVA_ISO.signature -v dgst -sha512

Retrieving CA certificate from https://www.cisco.com/security/pki/certs/crcam2.cer ...
Successfully retrieved and verified crcam2.cer.
Retrieving SubCA certificate from https://www.cisco.com/security/pki/certs/innerspace.cer ...
Successfully retrieved and verified innerspace.cer.
Successfully verified root, subca and end-entity certificate chain.
Successfully fetched a public key from ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM.
Successfully verified the signature of DCNM_To_NDFC_Upgrade_Tool_OVA_ISO using
ACI_4070389ff0d61fc7fbb8cdfdec0f38f30482c22e.PEM
```

Step 4 Back up existing configuration.

The backup tool collects last 90 days Performance Management data.

a) Log in to your DCNM Release 11.5(4) appliance console.
b) Create a screen session.

The following command creates a session which allows you to execute additional commands:

```
dcnm# screen
```

Note that the commands continue to run even when the window is not visible or if you get disconnected.

c) Gain super user (`root`) access.

```
dcnm# su
Enter password: <root-password>
[root@dcnm]#
```

d) For OVA and ISO, enable execution permissions for the upgrade tool.

```
[root@dcnm]# chmod +x ./DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
```

e) Run the upgrade tool you downloaded in the previous step.

- For Windows:

```
G:\DCNM_To_NDFC_Upgrade_Tool_LIN_WIN>DCNMBackup.bat
DCNMBackup.bat
Enter DCNM root directory [C:\Program Files\Cisco Systems\dcm]:

Initializing, please wait...

********************************************************************************

Welcome to DCNM-to-NDFC Upgrade Tool for Linux/Windows.

This tool will analyze this system and determine whether you can move to NDFC 12.2.1 or not.

If upgrade to NDFC 12.2.1 is possible, this tool will create files to be used for performing
 the upgrade.

Thank you!

********************************************************************************

This tool will backup config data. Exporting Operational data like Performance(PM) might take
 some time.

Do you want to export operational data also? [y/N]: y
********************************************************************************


Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring
the backup file generated by this tool.


Please enter the encryption key:
Enter it again for verification:
….
2024-02-26 17:57:32,247 [main] INFO  DCNMBackup - Creating final tar.gz file....
2024-02-26 17:57:32,649 [main] INFO  DCNMBackup - Final tar.gz elapsed time: 402 in ms
2024-02-26 17:57:32,650 [main] INFO  DCNMBackup - Backup done.
2024-02-26 17:57:32,657 [main] INFO  DCNMBackup - Log file: backup.log
2024-02-26 17:57:32,658 [main] INFO  DCNMBackup - Backup file:
backup11_win57_20240226-172247.tar.gz
```

- For Linux:

```
# ./DCNMBackup.sh
Enter DCNM root directory [/usr/local/cisco/dcm]:

Initializing, please wait...

********************************************************************************

Welcome to DCNM-to-NDFC Upgrade Tool for Linux/Windows.

This tool will analyze this system and determine whether you can move to NDFC 12.2.1 or not.

If upgrade to NDFC 12.2.1 is possible, this tool will create files to be used for performing
 the upgrade.

Thank you!

********************************************************************************

This tool will backup config data. Exporting Operational data like Performance(PM) might take
 some time.
```

```
Do you want to export operational data also? [y/N]: y
********************************************************************************


Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring
the backup file generated by this tool.


Please enter the encryption key:
Enter it again for verification:
2024-02-27 07:53:46,562 [main] INFO  DCNMBackup - Inside init() method
2024-02-27 07:53:46,564 [main] INFO  DCNMBackup - Loading properties....
2024-02-27 07:53:46,649 [main] INFO  DCNMBackup - Inside checkLANSwitches...
2024-02-27 07:53:46,732 [main] INFO  fms.db - set database url
as:jdbc:postgresql://localhost:5432/dcmdb
2024-02-27 07:53:46,887 [main] INFO  DCNMBackup - LAN Switch count: 0
2024-02-27 07:53:46,889 [main] INFO  DCNMBackup - Inside exportDBTables...
2024-02-27 07:53:46,892 [main] INFO  DCNMBackup - Exporting ---------> statistics
2024-02-27 07:53:46,903 [main] INFO  DCNMBackup - Exporting ---------> sequence
2024-02-27 07:53:46,964 [main] INFO  DCNMBackup - Exporting ---------> clustersequence
2024-02-27 07:53:46,965 [main] INFO  DCNMBackup - Exporting ---------> logicsvr_fabric
……..
2024-02-27 07:53:49,147 [main] INFO  DCNMBackup - Creating final tar.gz file....
2024-02-27 07:53:49,183 [main] INFO  DCNMBackup - Final tar.gz elapsed time: 35 in ms
2024-02-27 07:53:49,183 [main] INFO  DCNMBackup - Backup done.
2024-02-27 07:53:49,183 [main] INFO  DCNMBackup - Log file: backup.log
2024-02-27 07:53:49,183 [main] INFO  DCNMBackup - Backup file:
backup11_onefiveseven.cisco.com_20240227-72149.tar.gz
```

- For OVA:

```
# ./DCNM_To_NDFC_Upgrade_Tool_OVA_ISO
********************************************************************************

Welcome to DCNM-to-NDFC Upgrade Tool for OVA/ISO.

This tool will analyze this system and determine whether you can move to
NDFC 12.2.1 or not.

If upgrade to NDFC 12.2.1 is possible, this tool will create files
to be used for performing the upgrade.

NOTE:
Only backup files created by this tool can be used for upgrading,
older backup files created with 'appmgr backup' CAN NOT be used
for upgrading to NDFC 12.2.1

Thank you!

********************************************************************************

Continue? [y/n]: y

Collect operational data (e.g. PM, EPL)? [y/n]: y

Does this DCNM 11.5(4) have DCNM Tracker feature enabled on any switch on any fabric? [y/n]:
 n

Sensitive information will be encrypted using an encryption key.
This encryption key will have to be provided when restoring
the backup file generated by this tool.
```

```
Please enter the encryption key:
Enter it again for verification:

Adding backup header
Collecting DB table data
Collecting DB sequence data
Collecting stored credentials
Collecting Custom Templates
Collecting CC files
Collecting L4-7-service data
Collecting CVisualizer data
Collecting EPL data
Collecting PM data - WARNING: this will take a while!


Collecting AFW app info
Decrypting stored credentials
Adjusting DB tables
Creating backup file
Done.
Backup file: backup11_host108_20240227-153940.tar.gz
```

Step 5 Deploy a brand new Nexus Dashboard cluster as described in one of the earlier chapters in this document.

Ensure that you complete all guidelines and prerequisites for the Nexus Dashboard platform, the Fabric Controller service, and the specific form factor listed in the deployment chapters above.

Note

- Deploy the Nexus Dashboard cluster only with the Fabric Controller (NDFC) service; do not deploy with other services (for example, with Fabric Controller + Insights). Doing so will cause the migration to fail when you restore the configuration backup in the next step with the message `Backup file is not compatible with this deployment mode`.
- You must provide the required number of Persistent IP addresses in the Nexus Dashboard Fabric Controller UI before proceeding with restoring your DCNM configuration..
- If your existing configuration used smart licensing with direct connectivity to Cisco Smart Software Management (CSSM), you must ensure that your new Nexus Dashboard has the routes required to reach the CSSM website.

  Ensure that subnets for IP addresses on `smartreceiver.cisco.com` are added to the route table in the Nexus Dashboard's **Admin** > **System Settings** > **Routes** page for the Nexus Dashboard management network.

  You can ping `smartreceiver.cisco.com` to find the most recent subnet, for example:

  ```
  $ ping smartreceiver.cisco.com
  PING smartreceiver.cisco.com (146.112.59.81): 56 data bytes
  64 bytes from 146.112.59.81: icmp_seq=0 ttl=52 time=48.661 ms
  64 bytes from 146.112.59.81: icmp_seq=1 ttl=52 time=44.730 ms
  64 bytes from 146.112.59.81: icmp_seq=2 ttl=52 time=48.188 ms
  ```

  In addition, because NDFC is considered a new product instance, you must re-establish trust. If you took the backup with an expired Trust Token, you must manually run the Smart Licensing Configuration wizard and enter a valid token after the upgrade.

Step 6 Restore the configuration backup in the new cluster.

a) Log in to your Nexus Dashboard using an admin account.
b) From the dropdown menu at the top, choose **Fabric Controller**.
c) From the left navigation menu, choose **Admin** > **Backup and Restore**.

d) In the main pane, click **Restore**.
e) In the **Restore Now** window, provide the details.
   - Choose **Config Only** or **Full** based on the backup that you created in the previous step.
   - Choose the **Source** where the backup file is located, then upload the file or provide the remote server location and path.
   - Enter the **Encryption Key** you provided during configuration backup.
   - Ensure that the **Ignore External Service IP Configuration** option is unchecked.

f) Click **Next**, verify the information, and proceed to **Restore** the configuration.

   The UI is locked while the restore is in progress; the time required to restore depends on the data in the backup file.

   Once the restore process is completed, click the `x` icon in the **Restore Now** pop-up window to close it.

   After successful restoration, click **Reload the page** or refresh the browser page to complete restore and begin using you Nexus Dashboard Fabric Controller.

Step 7 Complete the post-upgrade tasks.

a) If you are using the SAN Controller persona:

   After restoring the data from backup, all the server-smart licenses are **OutofCompliance**.

   You can migrate to Smart Licensing using Policy from the **Operations** > **License Management** > **Smart** page in the UI and establish trust with CCSM using SLP.

b) If you are using the Fabric Controller persona:

   The following features are not carried over when you upgrade from DCNM 11.5(4):
   - Endpoint Locator must be reconfigured
   - IPAM Integration must be reconfigured
   - Alarm Policies must be reconfigured
   - Custom topologies must be recreated and saved
   - PM collection must be re-enabled on fabrics
   - Temperature data collection must be re-enabled to start collecting data
   - Switch images must be uploaded

Deployment Type in Release 11.5(4)	In 11.5(4), trap IP address is collected from	LAN Device Management Connectivity	Trap IP address after upgrade	Result
LAN Fabric Media Controller	`eth1` (or `vip1` for HA systems)	Management	Belongs to Management subnet	Honored There is no configuration difference. No further action required.

Deployment Type in Release 11.5(4)	In 11.5(4), trap IP address is collected from	LAN Device Management Connectivity	Trap IP address after upgrade	Result
LAN Fabric Media Controller	`eth0` (or `vip0` for HA systems)	Management	Does not belong to Management subnet	Ignored, another IP from the Management pool will be used as trap IP. Configuration difference is created. On the **Web UI > LAN > Fabrics > Fabrics**, double click on the Fabric to view **Fabric Overview**. From **Fabrics Actions** drop-down list, select **Recalculate Config**. Click **Deploy Config**.
LAN Fabric Media Controller	`eth0` (or `vip0` for HA systems)	Data	Belongs to Data subnet	Honored There is no configuration difference. No further action required.
LAN Fabric Media Controller	`eth0` (or `vip0` for HA systems)	Data	Does not belong to Data subnet	Ignored, another IP from the Data pool will be used as trap IP. Configuration difference is created. On the **Web UI > LAN > Fabrics > Fabrics**, double click on the Fabric to view **Fabric Overview**. From **Fabrics Actions** drop-down list, select **Recalculate Config**. Click **Deploy Config**.

Deployment Type in Release 11.5(4)	In 11.5(4), trap IP address is collected from	LAN Device Management Connectivity	Trap IP address after upgrade	Result
SAN Management	OVA/ISO – • `trap.registaddress` (if set) • eth0 (if `trap.registaddress` is not set) Windows/Linux – • `trap.registaddress` (if set) • Interface based on event-manager algorithm (if `trap.registaddress` is not set	Not applicable	Belongs to Data subnet	Honored There is no configuration difference. No further action required.
		Not applicable	Does not belong to Data subnet	Ignored, another IP from the Data pool will be used as trap IP.