

Prerequisites: Nexus Dashboard

- Prerequisites and guidelines for all enabled services, on page 1
- Communication Ports, on page 7
- Fabric Connectivity, on page 11
- Services Co-location Use Cases, on page 16
- Pre-Installation Checklist, on page 19

Prerequisites and guidelines for all enabled services

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.

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.
- Nexus Dashboard nodes must be in synchronization with the NTP server; however, there can be latency of up to 1 second between the Nexus Dashboard nodes. If the latency is greater than or equal to 1 second between the Nexus Dashboard nodes, this may result in unreliable operations on the Nexus Dashboard cluster.

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 1: 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 using SSH • DNS and NTP communication • Nexus Dashboard firmware upload • Intersight device connector • AAA traffic

The two networks have the following requirements:

- For all 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 using 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 Nexus Dashboard management interface (`bond1`) has internal iptables rules that rate-limit ICMP packets to an average of 6 packets per second with a burst limit of 5. This ICMP rate and burst limit also applies to the data network port (`bond0`). If you are using ICMP-based monitoring tools to track the health of the management network, you may observe intermittent packet drops if the polling frequency exceeds these limits. This is expected behavior designed to protect the management plane.

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 fabric controllers or switches depends on the specific service you plan to enable, see the "Network Requirements" sections in the service-specific chapters below.

Table 2: 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 using 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

Nexus Dashboard supports pure IPv4, pure IPv6, or dual stack IPv4/IPv6 configurations for the cluster nodes and services.

When defining an IP address configuration, the following guidelines apply:

- All nodes and networks in the cluster must have a uniform IP configuration, either pure IPv4, 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 IP addresses during the cluster bootstrap workflow.

    Management IP addresses 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 GUI 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 IP addresses are used to log in to the GUI 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 IP addresses

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 IP addresses 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 such as gratuitous ARP or neighbor discovery to advertise the persistent IP addresses within its Layer 3 network.

While that is still supported, this release also allows you to configure the persistent IP addresses feature even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IP addresses are advertised out of each node's data links using BGP, which we refer to as "Layer 3 mode". The IP addresses must also be part of a subnet that is not overlapping with any of the nodes' management or data subnets. If the persistent IP addresses are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IP addresses 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 IP addresses 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 address requirements listed in the service-specific sections that follow.

The total number of persistent IP addresses 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 3: 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 Used to check connectivity through polls
SSH	22	TCP	In/Out	CLI and CIMC of the cluster nodes Basic mechanism to access cluster nodes and CIMC
TACACS	49	TCP	Out	TACACS server If you are using TACACS authentication, cluster nodes need to connect to this port
DNS	53	TCP/UDP	Out	DNS server Used to resolve DNS names (for example, Intersight or CSSM)
HTTP	80	TCP	Out	Internet/proxy If Internet connection to connect to Intersight or CSSM uses proxy through port 80, this is required
NTP	123	UDP	Out	NTP server Used to get correct time set for cluster nodes

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
HTTPS	443	TCP	In/Out	UI, other clusters (for multi-cluster connectivity), fabrics, Internet/proxy Used to access Nexus Dashboard UI and provides means for the nodes to access another cluster's UI when you have multi-cluster connectivity configured
LDAP	389 636	TCP	Out	LDAP server If you are using LDAP authentication, cluster nodes need to connect to this port
Radius	1812	TCP	Out	Radius server If you are using RADIUS authentication, cluster nodes need to connect to this port
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics Used for Kubernetes cluster nodes for key management service
Infra-Service	30012 30021 30500-30600	TCP/UDP	In/Out	Other cluster nodes Used for Kubernetes cluster nodes internal communication

Table 4: 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 Used to check connectivity through polls

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
SSH	22	TCP	Out	In-band of switches and APIC Used to collect info through SSH from switches and controllers
DNS	53	TCP/UDP	In/Out	Other cluster nodes and DNS server Used for Kubernetes services for intra-cluster connectivity and for DNS integration for Nexus Dashboard Insights
NFSv3	111	TCP/UDP	In/Out	Remote NFS server Used for backup and storing of the Nexus Dashboard Insights flows
HTTPS	443	TCP	Out	In-band of switches and APIC/NDFC Used for REST API
NFSv3	608	UDP	In/Out	Remote NFS server Used for backup and storing of the Nexus Dashboard Insights flows
SSH	1022	TCP/UDP	In/Out	Other cluster nodes Used for Kubernetes infrastructure inter-cluster nodes
NFSv3	2049	TCP	In/Out	Remote NFS server Used for backup and storing of the Nexus Dashboard Insights flows
VXLAN	4789	UDP	In/Out	Other cluster nodes Used for VTEP connectivity for Kubernetes cluster nodes data connectivity
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics

Service	Port	Protocol	Direction `In`—towards the cluster `Out`—from the cluster towards the fabric or outside world	Connection
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 Used for Kubernetes infrastructure inter-cluster nodes
Infra-Service	30016 30017	TCP/UDP	In/Out	Other cluster nodes Used for Kubernetes infrastructure inter-cluster nodes
Infra-Service	30019	UDP	In/Out	Other cluster nodes Used for Kubernetes infrastructure inter-cluster nodes
Infra-Service	30500-30600	TCP/UDP	In/Out	Other cluster nodes Used for Kubernetes infrastructure inter-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 fabrics. 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 fabric'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 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

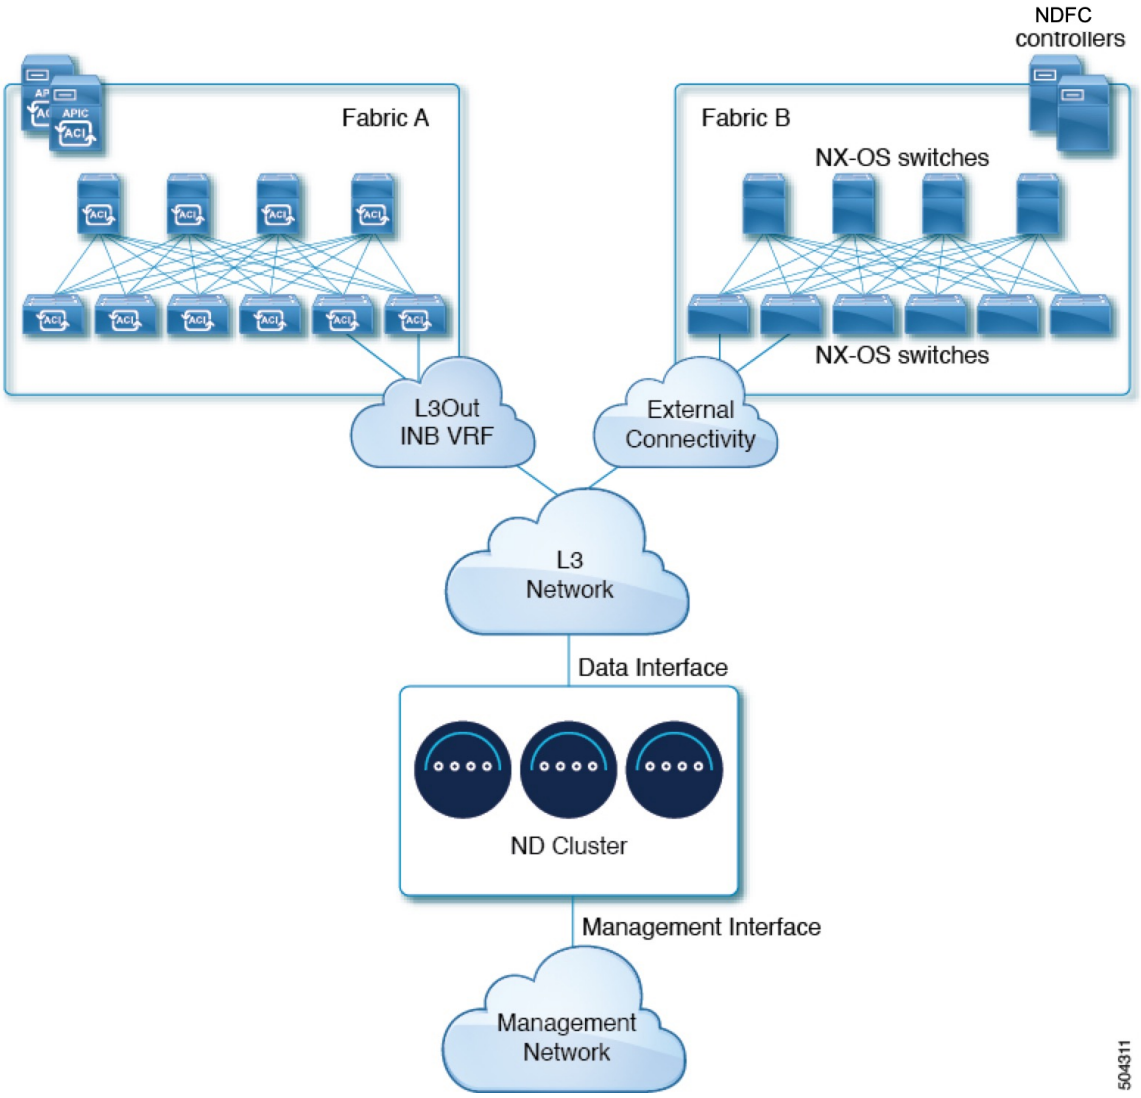
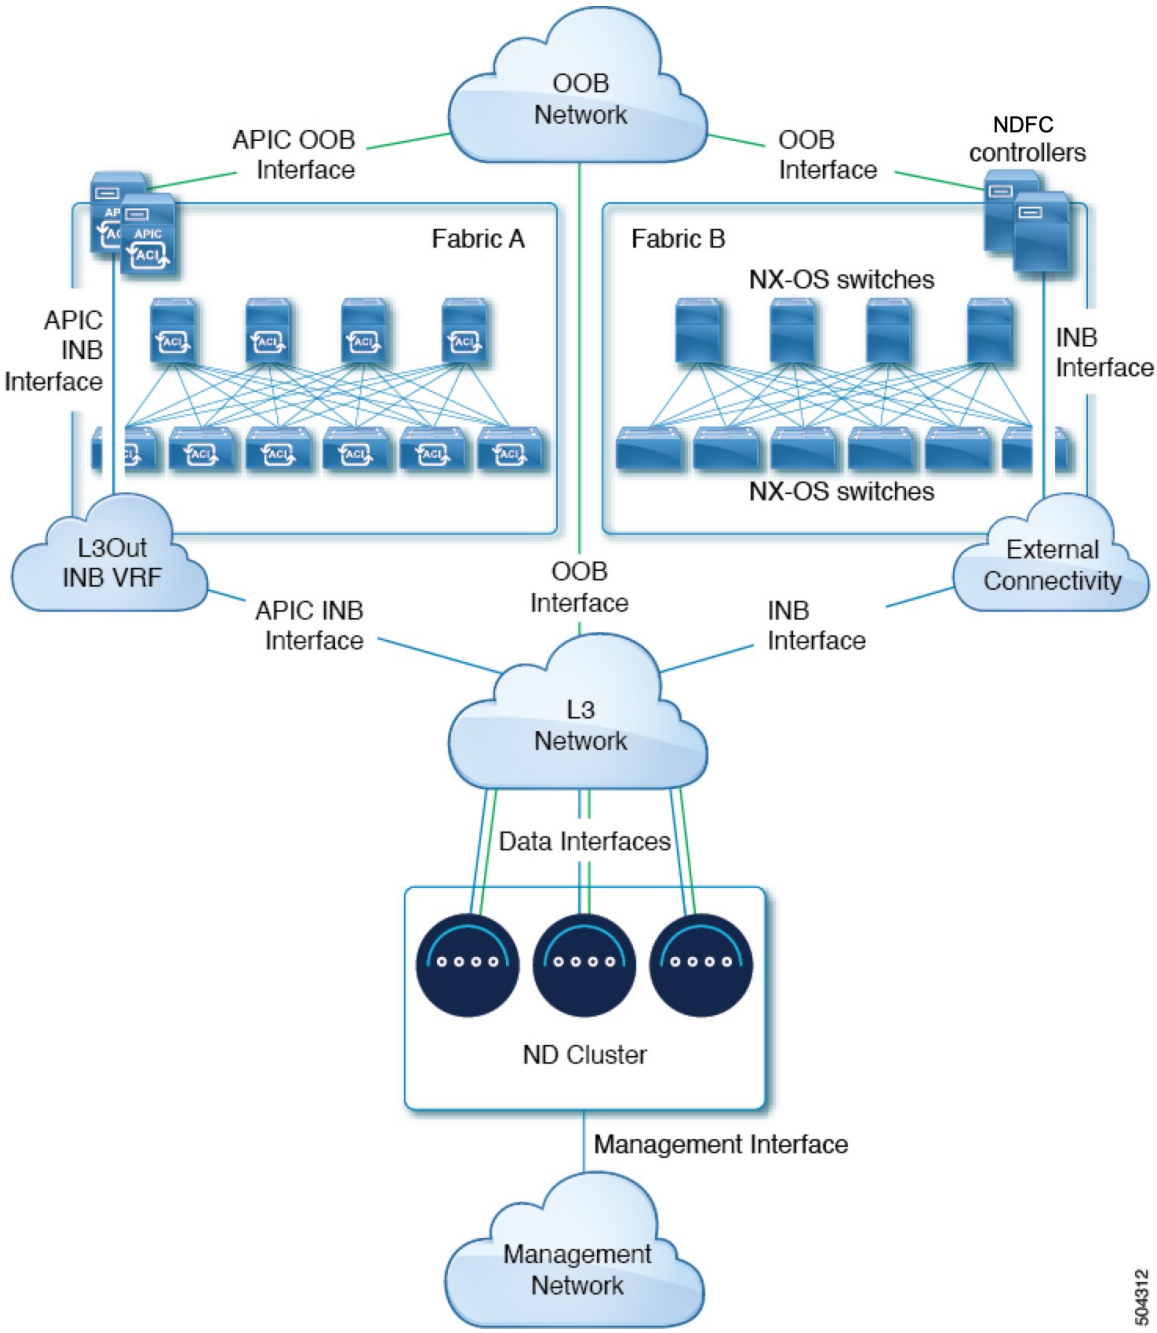


Figure 2: Connecting via Layer 3 Network, Nexus Dashboard Orchestrator



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 fabric'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

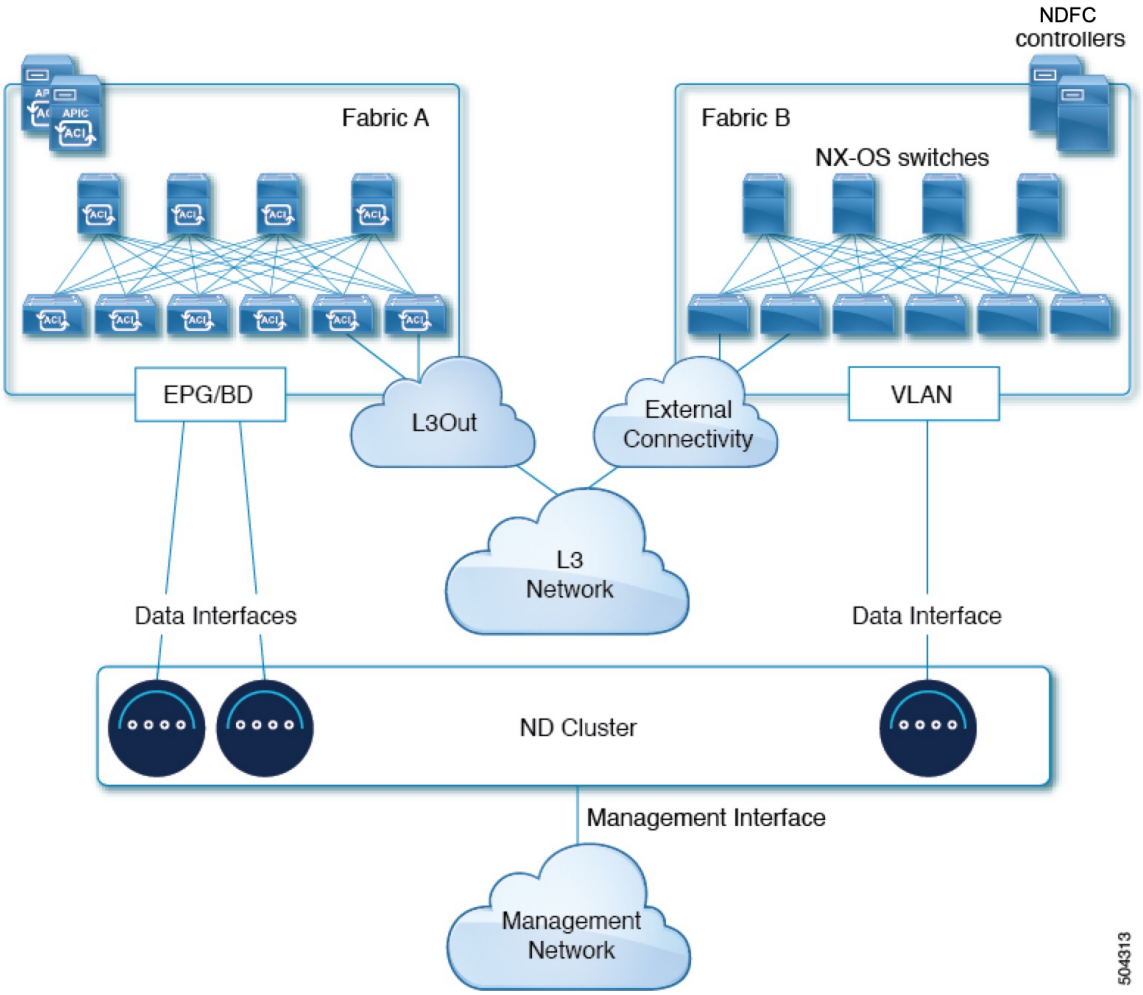
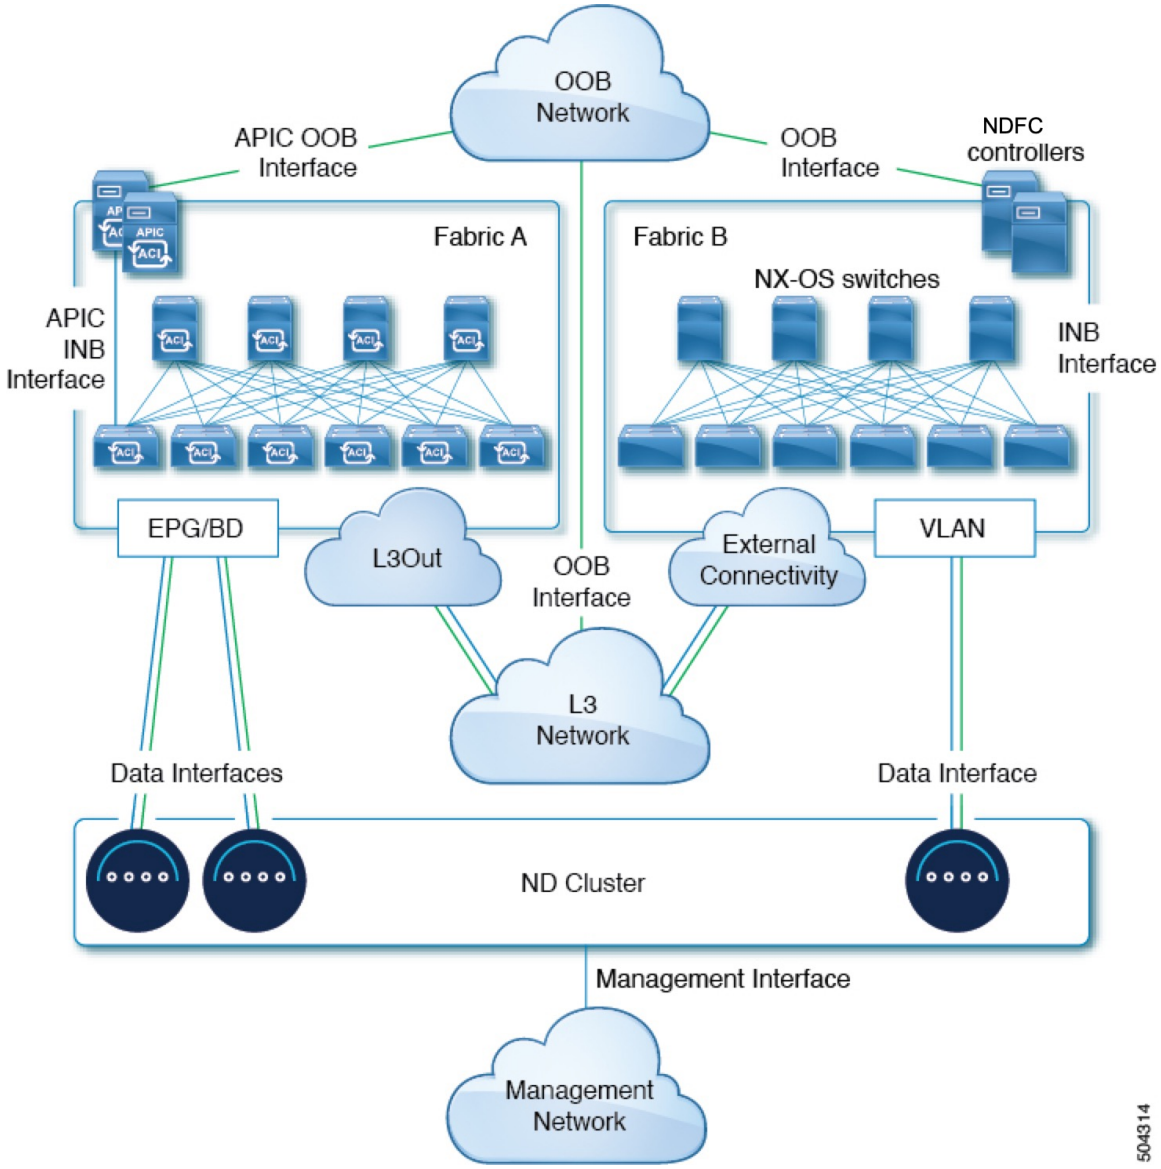


Figure 4: Connecting Directly to Leaf Switches, Nexus Dashboard Orchestrator



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.

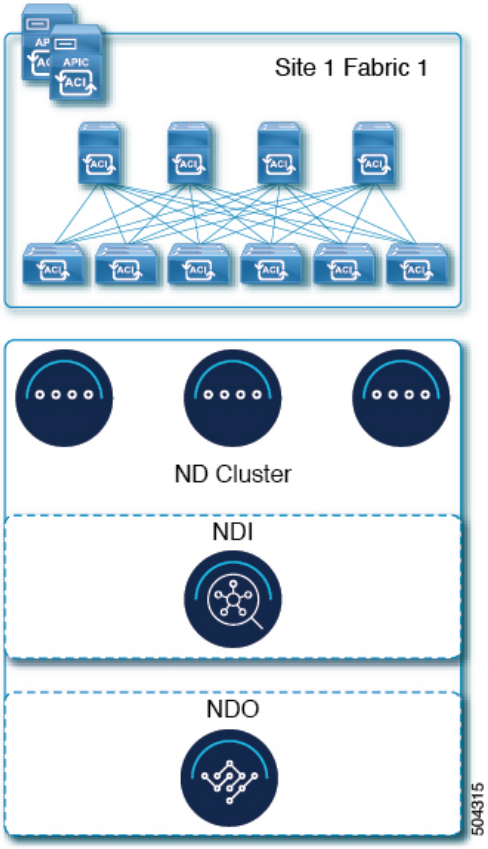
Note This release does not support co-hosting services in Nexus Dashboard clusters that are deployed in Linux KVM, AWS, Azure, or RHEL. All services co-hosting scenarios below apply for physical or VMware ESX cluster form factors only. For additional cluster sizing and deployment planning reference information,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

Single Fabric, Nexus Dashboard Insights and Orchestrator

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

Note For Nexus Dashboard Orchestrator, we recommend a centralized, single-fabric deployment.

Figure 5: Single Fabric, Nexus Dashboard Insights and Orchestrator

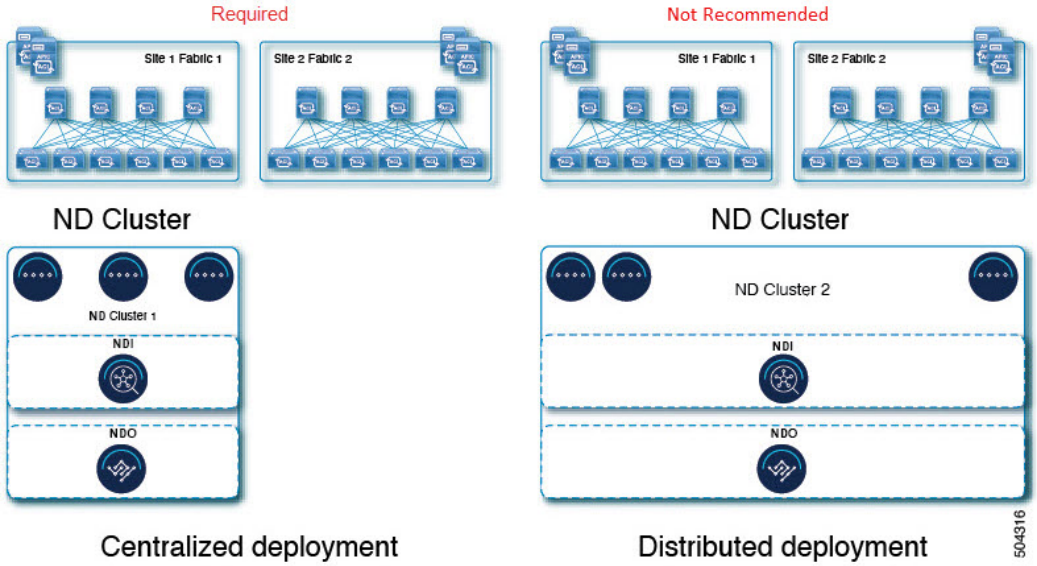


Multiple Fabrics, Single Cluster for Nexus Dashboard Insights and Orchestrator

In a multiple fabrics 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 fabrics, however since the Insights and Orchestrator services do not gain redundancy benefits

from a distributed cluster and could instead be exposed to interconnection failures when nodes are in different fabrics, we require the deployment option on the left:

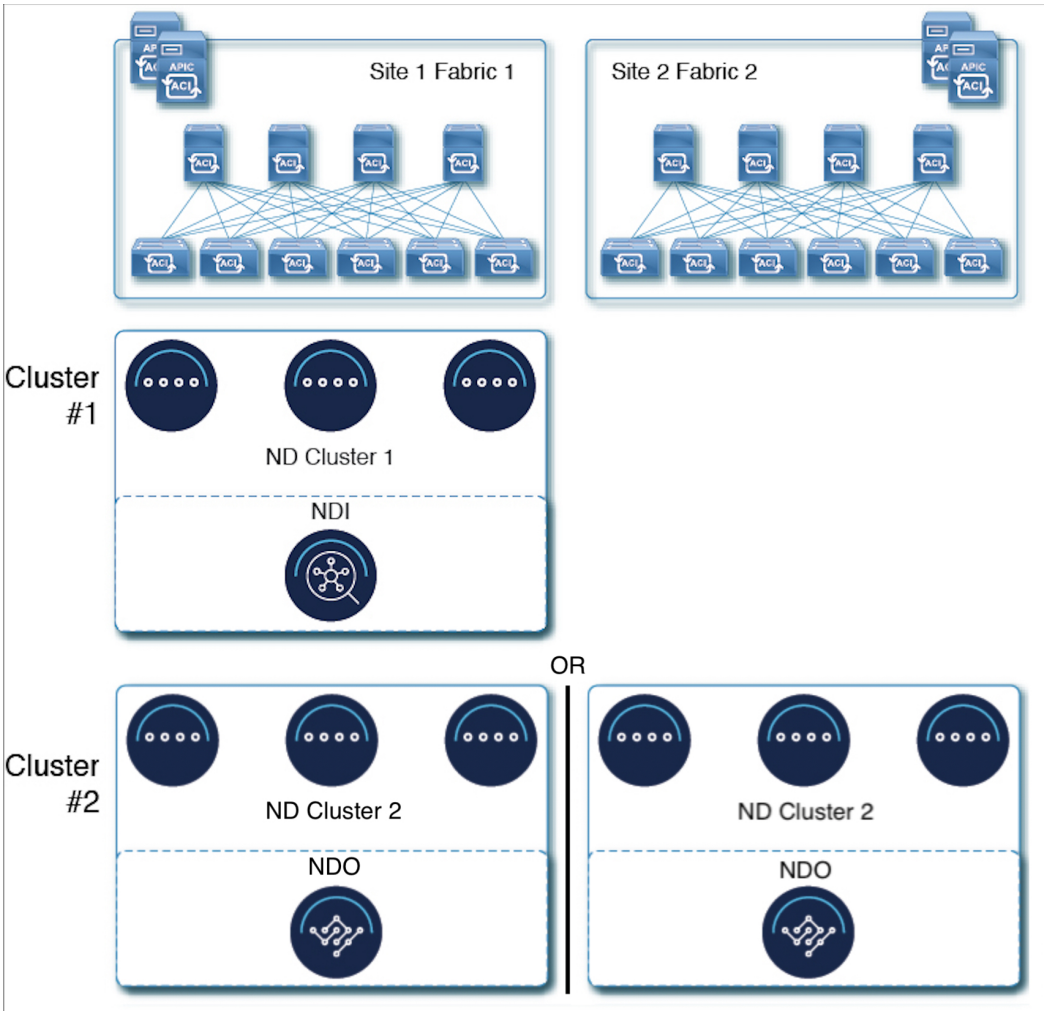
Figure 6: Multiple Fabrics, Single Cluster for Nexus Dashboard Insights and Orchestrator



Multiple Fabrics, Multiple Clusters for Nexus Dashboard Insights and Orchestrator

In this case, we recommend deploying two Nexus Dashboard clusters, with one of them dedicated to the Nexus Dashboard Orchestrator service using the virtual or cloud form factor. All the nodes for the Nexus Dashboard Orchestrator service should be in either Site 1 or Site 2, but not both.

Figure 7: Multiple Fabrics, 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 5: Cluster Details

Parameters	Example	Your Entry
Cluster **Name**	`nd-cluster`	
NTP Server	`170.78.48.55`	
DNS Provider	`170.71.68.83`	
DNS Search Domain	`cisco.com`	

Parameters	Example	Your Entry
App Network	`172.17.0.0/16`	
Service Network	`100.80.0.0/16`	

Note

- If you use multi-cluster connectivity to establish connectivity between multiple Nexus Dashboard clusters, the names of the clusters that you plan to connect together must be unique.
- 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 6: 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`	

Parameters	Example	Your Entry
(Optional) **Data Network VLAN** of the first node	`101`	
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`	

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