

Prerequisites and Guidelines

- General prerequisites and guidelines, on page 1
- Prerequisites for the Nexus Dashboard data network and management network, on page 5
- Prerequisites for the Nexus Dashboard internal app and service networks, on page 7
- Prerequisites for LAN deployments, on page 7
- Prerequisites for SAN deployments, on page 21
- Nexus Dashboard persistent IP addresses, on page 30
- Round trip time requirements, on page 37
- Prerequisites for Nexus Dashboard storage (virtual form factors), on page 37
- Fabric connectivity, on page 39

General prerequisites and guidelines

This section describes requirements and guidelines for the Nexus Dashboard cluster regardless of the deployment type.

General deployment guidelines and restrictions

- Deploying virtual Nexus Dashboard VMs on remote storage is unsupported and may lead to unexpected behavior.
- Expanding the cluster size after initial deployment is not supported. If you need to increase the cluster size, deploy a new cluster with the required number of nodes and perform a backup and restore of the configuration.
- For virtual clusters built using three data nodes and three app nodes, you must configure the data nodes as primary and the app nodes as secondary.

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

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.

The following guidelines apply for DNS:

- For external DNS servers, both TCP and UDP traffic must be allowed. See Communication ports for LAN deployments, on page 11 and Communication ports for SAN deployments, on page 22 for more information.
- 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:

- We recommend that you do not use a Windows server as the NTP server.
- 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.
- These are the requirements for NTP delay, offset, and jitter:
  - Delay: < 100 ms
  - Offset: ±25 ms
  - Jitter: < 10 ms

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

IPv4 and IPv6 support

Nexus Dashboard supports IPv4-only, IPv6-only or dual-stack 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 IPv4-only, IPv6-only, or dual stack IPv4/IPv6.
- If you deploy the cluster in IPv4-only mode and want to switch to dual stack IPv4/IPv6 or IPv6-only, you must redeploy the cluster.
- For dual stack configurations:
  - The data, management, app, and service networks must be in dual stack mode.

    Mixed configurations, such as IPv4 data network and dual stack management network, are not supported.
  - For IPv6-based Nexus Dashboard deployments, the CIMCs of all physical servers must also have IPv6 addresses.
  - 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, provided both are configured.

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

    Clusters deployed through the vND deployment process on AWS public cloud do not support IPv6-only or dual stack 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.
  - IPv6 addresses are required for physical servers' CIMCs.

Necessary URLs for certain connections

There are certain URLs that Nexus Dashboard must reach that are necessary for these connections:

- Cisco Intersight: Connecting your Nexus Dashboard cluster to Cisco Intersight has these benefits:
  - Automatic meta data updates that certain features can use to provide updated data
  - TAC log collection and uploads
- Connecting to Smart Licensing
- Pulling energy management stats from electricity maps

These are the URLs that Nexus Dashboard must reach for these connections and why:

URL	Protocol/Port/Service	Description
amazontrust.com	TCP/80(HTTP) TCP/443(HTTPS)	Used to securely connect to Cisco Intersight
connectdna.cisco.com	TCP/443(HTTPS)	Used to securely connect to Cisco Intersight and Smart Licensing
swapi.cisco.com	TCP/443(HTTPS)	Used to securely connect to Cisco Smart Licensing
svc.ucs-connect.com	TCP/443(HTTPS)	Used to securely connect to Cisco Intersight
svc-static1.ucs-connect.com	TCP/443(HTTPS)	Used to securely connect to Cisco Intersight
svc.eu-central-1.intersight.com	TCP/443(HTTPS)	Used to securely connect to Cisco Intersight (EMEA Region)
svc-static1.eu-central-1.intersight.com	TCP/443(HTTPS)	Used to securely connect to Cisco Intersight (EMEA Region)

Prerequisites for the Nexus Dashboard data network and management network

Nexus Dashboard is deployed as a cluster, connecting each node to two networks. When first configuring Nexus Dashboard, for each cluster node, you will need to provide two IP addresses for the two Nexus Dashboard interfaces:

- One connected to the data network, which is used for back-end, cluster, and Infra connectivity for optimal performance
- The other connected to the management network, which is used for seamless GUI and front-end operations

Table 1: External network purpose

Data network	Management network
• Nexus Dashboard node clustering • Service to service communication • Nexus Dashboard nodes to Cisco APIC and NX-OS controller capability communication • 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 • Multi-cluster connectivity

The two networks have the following requirements:

- The management network and data network must be in different subnets.

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.

- Changing the data subnet requires re-deploying the cluster, so we recommend using a larger subnet (such as /27) than the bare minimum required by the nodes and features to account for any additional features that may require more IP addresses 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 and 443 as the 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.

- If you are using telemetry, by default, the data network must provide IP reachability to the in-band network of each fabric and of the Cisco APIC for an ACI fabric (if you are using the orchestration functionality), as well as to these integrations.

Note You can also define routes in the route table of the Nexus Dashboard and use the management network instead to reach to any of the following services.

  - For DNS integration, to 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 VMware vCenter integration, to the VMware vCenter.
  - For AppDynamics integration, to the AppDynamics controller.

For more information, see *Working with Integrations in Your Nexus Dashboard*.

Note If all the integrations are in same subnet as the management network, then they will use the management network.

Prerequisites for the Nexus Dashboard internal app and service networks

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

- App network--Used for applications internally within Nexus Dashboard. The app network must be a `/16` network for IPv4 or `/108` network for IPv6 and a default value is pre-populated during deployment.
- Service network--Used internally by the Nexus Dashboard. The service network must be a `/16` network for IPv4 or `/108` network for IPv6 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 app network and service network 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 the app or service network, 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.

Prerequisites for LAN deployments

Network prerequisites for LAN deployments

These network prerequisites apply for LAN deployments:

- All new Nexus Dashboard deployments must have the management network and data network in different subnets.

- Interfaces on both data and management networks can be either Layer 2 or Layer 3 adjacent. For data network Layer 3 adjacency, you must configure BGP during the bootstrap process. Management network interfaces do not support the BGP protocol. If different Nexus Dashboard nodes are deployed with management addresses in different subnets, those will simply be routed to one other.
- You must use persistent data IP addresses to bring up the cluster, so you must allocate a certain number of persistent IP addresses depending on your configuration.
  - If your cluster has 1 node, allocate 3 persistent IP addresses.
  - If your cluster has 3 or more nodes, allocate 5 persistent IP addresses.
  - If you configure dual stack IPv4 and IPv6, then add the same number of persistent IP addresses for IPv6 (in other words, 5 IPv4 and 5 IPv6 persistent IP addresses if you configure dual stack).

  For more information about persistent IP addresses, see Nexus Dashboard persistent IP addresses, on page 30. You must allocate the minimum required persistent IP addresses during the bootstrap process. You can allocate additional persistent IP addresses after the cluster is deployed using the External Service Pools configuration in the GUI.
- The pod profile policy is dynamically configured based on the number of nodes that you deploy.

Prerequisites for onboarding ACI fabrics in LAN deployments

These network prerequisites apply for onboarding ACI fabrics in LAN deployments:

- If you are planning to use orchestration 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 fabric's APIC cluster or both.

  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.

Additional prerequisites for using orchestration with ACI fabrics

If you plan to use orchestration with ACI fabrics, these prerequisites also apply:

- If you plan to use orchestration with ACI fabrics and remote leaf switches, these restrictions apply:
  - Remote leaf switches in one fabric cannot use another fabric's L3Out.
  - Stretching a bridge domain between one fabric (local leaf or remote leaf) and a remote leaf in another fabric is not supported.
- Orchestration is only supported on single-node Nexus Dashboard clusters (virtual-data profile or physical appliances) for non-production (lab) deployments. If you want to enable Orchestration on one of these form factors, it must be enabled using the built-in swagger API.
  1. From the Nexus Dashboard UI, click on the "?" icon and choose **Help Center**.
  2. In the **Help Center**, click on **API reference: Swagger (In-product)** .
  3. Within the API listing, click the **Infra** group from the left navigation.

4. Locate the **System Settings** sub-menu and click the arrow to expand it, if necessary, then search for `/settings/general/actions/enableOrchestration`.
5. Expand the API and click **Try it Out**.

Orchestration services will now be enabled on your cluster.

Additional prerequisites for using telemetry with ACI fabrics

If you plan to use telemetry with ACI fabrics, these prerequisites also apply:

- Telemetry collection is supported with OOB networks of APIC and switches as long as they're running ACI version 6.1(2f) or later.
- You have configured NTP settings on Cisco APIC.

  For more information, see Configure NTP in ACI Fabric Solution.
- If you plan to use the following flow telemetry or traffic analytics functions, 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, Precision Time Protocol (PTP) must be enabled on Cisco APIC so that telemetry 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 using 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 flow telemetry, hence a GNSS/GPS source is typically not required.

  For a single-pod ACI fabric, you can connect your PTP GM using 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 using leaf switches or using 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 using IPN devices.

  See section "Precision Time Protocol" in the *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 telemetry 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.

Note There is an issue with streaming telemetry from ACI Spines to Nexus Dashboard when data network and ACI in-band are in the same subnet.

Prerequisites for onboarding NX-OS, IOS XR, and IOS XE devices in LAN deployments

These network prerequisites apply for onboarding NX-OS, IOS XR, and IOS XE devices in LAN deployments:

- If you are planning to use orchestration to manage NX-OS fabrics, the data network must have in-band reachability for NX-OS fabrics.

Additional prerequisites for using telemetry with NX-OS fabrics or standalone NX-OS switches

If you plan to use telemetry with NX-OS fabrics or standalone NX-OS switches, these prerequisites also apply:

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

Note If you are using the Flow Telemetry feature, the data network must have IP reachability to the fabric's 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 telemetry.

  In both managed and monitor fabric mode, you must ensure PTP is correctly configured on all nodes in the fabric. You can enable PTP in the fabric 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 fabric.

  The quality of the time synchronization using 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 flow telemetry, hence a GNSS/GPS source is typically not required.

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

Communication ports for LAN deployments

Nexus Dashboard uses TLS or mTLS with encryption to protect data privacy and integrity while in transit.

Open all ports described in the tables in this section.

This table lists the management network communication ports for LAN deployments, where in the **Direction** column:

- **In** means toward the cluster
- **Out** means from the cluster toward the fabric or outside world

Note Ensure that ICMP traffic is permitted between the Nexus Dashboard nodes' management interfaces and the Cisco Smart Software Manager (CSSM) On-Prem installation.

Beginning with Nexus Dashboard release 4.2.1, you must permit ICMP traffic between the Nexus Dashboard nodes' management interfaces and the CSSM On-Prem installation. This requirement differs from previous releases and is necessary for successful connectivity verification during Smart Licensing configuration.

Table 2: Management network communication ports for LAN deployments

Service	Port	Protocol	Direction (`In` /`Out`)	Connection
CSSM On-prem	ICMP	ICMP	Out	Ensure that ICMP traffic is permitted between the Nexus Dashboard nodes' management interfaces and the Cisco Smart Software Manager (CSSM) on-premises installation. **Note** Beginning with Cisco Nexus Dashboard release 4.2.1, you must permit ICMP traffic between the Nexus Dashboard nodes' management interfaces and the Cisco Smart Software Manager (CSSM) on-premises installation. This requirement differs from previous releases and is necessary to ensure successful connectivity.

Service	Port	Protocol	Direction (In/Out)	Connection
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, CIMC, default gateway, switch discovery. **Note** Adding or discovering LAN devices uses ICMP echo packets as part of the discovery process. So if you have a firewall between the Nexus Dashboard cluster and your switches, it must allow ICMP messages through or the discovery process will fail. ICMP traffic on the management interface is rate-limited to an average of 6 packets/sec and a burst of 5. Monitoring systems should be configured with this limit in mind to avoid false-positive alerts regarding packet loss.
DHCP	67	UDP	In	If the local DHCP server is configured for bootstrap or POAP purposes. **Note** When using Nexus Dashboard as a local DHCP server for POAP purposes, all Nexus Dashboard primary node IP addresses must be configured as DHCP relays. Whether the Nexus Dashboard nodes' management IP addresses are bound to the DHCP server is determined by the LAN Device Management Connectivity in the server settings.
DHCP	68	UDP	Out	
DNS	53	TCP/UDP	Out	DNS server
HTTP	80	TCP	Out	Internet/proxy
HTTP (PnP)	9666	TCP	In	Cisco Plug and Play (PnP) for Catalyst devices is accomplished using Nexus Dashboard 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, as with POAP, runs on a persistent IP address that is associated with either the management or data subnet. The persistent IP subnet is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.

Service	Port	Protocol	Direction (In /Out)	Connection
HTTP (POAP)	80	TCP	In	Only used for device zero-touch provisioning using POAP, where devices can send (limited jailed write-only access to Nexus Dashboard) basic inventory information to Nexus Dashboard to start secure POAP communication. Nexus Dashboard Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
HTTPS	443	TCP	In/Out	UI, other clusters (for multi-cluster connectivity), fabrics, Internet/proxy
HTTPS/HTTP (NX-API)	443/80	TCP	Out	NX-API HTTPS/HTTP client connects to device NX-API server on port 443/80, which is also configurable. NX-API is an optional feature, used by limited set of Nexus Dashboard functions.
HTTPS (PnP)	9667	TCP	In	Cisco Plug and Play (PnP) for Catalyst devices is accomplished using Nexus Dashboard 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, as with POAP, runs on a persistent IP address that is associated with either the management or data subnet. The persistent IP subnet is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.

Service	Port	Protocol	Direction (In/Out)	Connection
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished using the Nexus Dashboard HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP address assigned to that pod. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
Infra-Service	30012 30021 30500-30600	TCP/UDP	In/Out	Other cluster nodes
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
LDAP	389 636	TCP	Out	LDAP server
NTP	123	UDP	Out	NTP server
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.
Radius	1812	TCP	Out	Radius server
SCP	22	TCP	In/Out	SCP is used by various features to transfer files between devices and Nexus Dashboard, such as for archiving backup files to remote server.. The Nexus Dashboard 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 Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.

Service	Port	Protocol	Direction (In /Out)	Connection
SCP/Show Techcollection	22	TCP	Out	Transport tech-support file from persistent IP address of Nexus Dashboard POAP-SCP pod to a separate ND cluster running telemetry. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings
SMTP	25	TCP	Out	You can configure the SMTP port on the **Admin** > **Server Settings** > **General** page. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from Nexus Dashboard to devices.
SNMP Trap	2162	UDP	In	SNMP traps from devices to Nexus Dashboard are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings,
SSH	22	TCP	In/Out	CLI and CIMC of the cluster nodes
TAC Assist	8884	TCP	In/Out	Other cluster nodes Used for TAC Assist, which is a service to collect **show tech** from switches and upload the information to Intersight. This port is used to exchange **show tech** data across cluster nodes.
TACACS	49	TCP	Out	TACACS server

Service	Port	Protocol	Direction (In/Out)	Connection
TFTP (POAP)	69	TCP	In	Only used for device zero-touch provisioning using POAP, where devices can send (limited jailed write-only access to Nexus Dashboard) basic inventory information to Nexus Dashboard to start secure POAP communication. Nexus Dashboard bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.

This table lists the data network communication ports for LAN deployments, where in the **Direction** column:

- **In** means toward the cluster
- **Out** means from the cluster toward the fabric or outside world

Table 3: Data network communication ports for LAN deployments

Service	Port	Protocol	Direction (In/Out)	Connection
DHCP	67	UDP	In	If the Nexus Dashboard local DHCP server is configured for bootstrap or POAP purposes. **Note** When using Nexus Dashboard as a local DHCP server for POAP purposes, all Nexus Dashboard primary node IP addresses must be configured as DHCP relays. Whether the Nexus Dashboard nodes' data IP addresses are bound to the DHCP server is determined by the LAN Device Management Connectivity in the server settings.
DHCP	68	UDP	Out	
DNS	53	TCP/UDP	In/Out	Other cluster nodes and DNS server
Flow Telemetry	5640-5671	UDP	In	In-band of switches Used to receive flow telemetry from fabrics

Service	Port	Protocol	Direction (In /Out)	Connection
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 using software telemetry to a persistent IP address associated with a Nexus Dashboard GRPC receiver service pod.
HTTP (PnP)	9666	TCP	In	Cisco Plug and Play (PnP) for Catalyst devices is accomplished using Nexus Dashboard 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, as with POAP, runs on a persistent IP address that is associated with either the management or data subnet. The persistent IP subnet is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
HTTP (POAP)	80	TCP	In	Only used for device zero-touch provisioning using POAP, where devices can send (limited jailed write-only access to Nexus Dashboard) basic inventory information to Nexus Dashboard to start secure POAP communication. Nexus Dashboard Bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
HTTPS	443	TCP	Out	In-band of switches and APIC and NX-OS fabrics
HTTPS/HTTP (NX-API)	443/80	TCP	Out	NX-API HTTPS/HTTP client connects to device NX-API server on port 443/80, which is also configurable. NX-API is an optional feature, used by limited set of Nexus Dashboard functions.

Service	Port	Protocol	Direction (In/Out)	Connection
HTTPS (PnP)	9667	TCP	In	Cisco Plug and Play (PnP) for Catalyst devices is accomplished using Nexus Dashboard 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, as with POAP, runs on a persistent IP address that is associated with either the management or data subnet. The persistent IP subnet is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
HTTPS (POAP)	443	TCP	In	Secure POAP is accomplished using the Nexus Dashboard HTTPS Server on port 443. The HTTPS server is bound to the SCP-POAP service and uses the same persistent IP address assigned to that pod. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	Nexus Dashboard 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
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, default gateway

Service	Port	Protocol	Direction (In /Out)	Connection
Infra-Service	3379 3380 4430 8989 9090 9969 9979 9989 15223 30002-30006 30009-30010 30012 30014-30021 30022 30025 30027 30028	TCP	In/Out	Other cluster nodes
Infra-Service	30016 30017	TCP/UDP	In/Out	Other cluster nodes
Infra-Service	30019	UDP	In/Out	Other cluster nodes
Infra-Service	30500-30600	TCP/UDP	In/Out	Other cluster nodes
Kafka	30001	TCP	In/Out	In-band IP of switches and APIC/Controller
KMS	9989	TCP	In/Out	Other cluster nodes and ACI fabrics
NFSv3	111	TCP/UDP	In/Out	Remote NFS server
NFSv3	608	UDP	In/Out	Remote NFS server
NFSv3	2049	TCP	In/Out	Remote NFS server
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.

Service	Port	Protocol	Direction (In/Out)	Connection
SCP	22	TCP	In/Out	SCP is used by various features to transfer files between devices and Nexus Dashboard, such as for archiving backup files to remote server.. The Nexus Dashboard 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 Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
SCP	22	TCP	Out	Transport tech-support file from persistent IP address of Nexus Dashboard POAP-SCP pod to a separate ND cluster running telemetry. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings
SMTP	25	TCP	Out	You can configure the SMTP port on the **Admin** > **Server Settings** > **General** page. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from Nexus Dashboard to devices.
SNMP Trap	2162	UDP	In	SNMP traps from devices to Nexus Dashboard are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings,
SSH	22	TCP	Out	UI, in-band of switches, and APIC
SSH	1022	TCP/UDP	In/Out	Other cluster nodes

Service	Port	Protocol	Direction (In /Out)	Connection
SW Telemetry	5695 30000 57500 30570	TCP	In/Out	Other cluster nodes Used to collect various telemetry information from fabrics Port 57500 is needed between switches and Nexus Dashboard for telemetry and NX-OS based switches
Tech support	2022	TCP	In	This is used to upload tech support from fabric to cluster.
TFTP (POAP)	69	TCP	In	Only used for device zero-touch provisioning using POAP, where devices can send (limited jailed write-only access to Nexus Dashboard) basic inventory information to Nexus Dashboard to start secure POAP communication. Nexus Dashboard bootstrap or POAP can be configured for TFTP or HTTP/HTTPS. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
VXLAN	4789	UDP	In/Out	Other cluster nodes

Prerequisites for SAN deployments

Network prerequisites for SAN deployments

These network prerequisites apply for SAN deployments:

- In SAN deployments, management and data networks can use the same subnet.
- Persistent IP addresses are supported on the data network only with Layer 2 adjacent and Layer 3 adjacent with BGP configured. However, if you configure the management and data networks to use the same subnet, then you cannot use Layer 3 adjacency because you must configure BGP during the bootstrap process for the data network Layer 3 adjacency, but the management network doesn't support Layer 3 adjacency with BGP configured.

  In this situation, you can:

  - Use Layer 2 adjacency instead if you want to configure the management and data networks to use the same subnet, or

- Use different subnets for the management and data networks, where you can configure Layer 2 adjacency on the management network and Layer 3 adjacency with BGP configured on the data network

- You must allocate some number of persistent IP addresses depending on your configuration. For more information about persistent IP addresses, see Nexus Dashboard persistent IP addresses, on page 30.

Communication ports for SAN deployments

Nexus Dashboard uses TLS or mTLS with encryption to protect data privacy and integrity while in transit.

This table lists the management network communication ports for SAN deployments.

Table 4: Management network communication ports for SAN deployments

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
DNS	53	TCP/UDP	Out	DNS server
GRPC (Telemetry)	33000	TCP	In	SAN Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to Nexus Dashboard persistent IP address.
HTTP	80	TCP	Out	Internet/proxy
HTTPS	443	TCP	In/Out	UI, other clusters (for multi-cluster connectivity), fabrics, Internet/proxy
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	Nexus Dashboard 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
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, CIMC, default gateway
Infra-Service	30012 30021 30500-30600	TCP/UDP	In/Out	Other cluster nodes

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
LDAP	389 636	TCP	Out	LDAP server
NTP	123	UDP	Out	NTP server
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.
Radius	1812	TCP	Out	Radius server
SCP	22	TCP	In/Out	SCP is used by various features to transfer files between devices and Nexus Dashboard, such as for archiving backup files to remote server.. The Nexus Dashboard 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 Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
SCP	22	TCP	Out	Transport tech-support file from persistent IP address of Nexus Dashboard POAP-SCP pod to a separate ND cluster running telemetry. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings
SMTP	25	TCP	Out	You can configure the SMTP port on the **Admin** > **Server Settings** > **General** page. This is an optional feature.

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
SNMP	161	TCP/UDP	Out	SNMP traffic from Nexus Dashboard to devices.
SNMP Trap	2162	UDP	In	SNMP traps from devices to Nexus Dashboard are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings,
SSH	22	TCP	In/Out	CLI and CIMC of the cluster nodes
Syslog	514	UDP	In	When Nexus Dashboard is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
TACACS	49	TCP	Out	TACACS server

This table lists the data network communication ports for SAN deployments.

Table 5: Data network communication ports for SAN deployments

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
DNS	53	TCP/UDP	In/Out	Other cluster nodes and DNS server
GRPC (Telemetry)	33000	TCP	In	SAN Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to Nexus Dashboard persistent IP address.
HTTPS	443	TCP	Out	In-band of switches and APIC and NX-OS fabrics
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	Nexus Dashboard 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
ICMP	ICMP	ICMP	In/Out	Other cluster nodes, default gateway

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward 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
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
KMS	9880	TCP	In/Out	Other cluster nodes and ACI fabrics
NFSv3	111	TCP/UDP	In/Out	Remote NFS server
NFSv3	608	UDP	In/Out	Remote NFS server
NFSv3	2049	TCP	In/Out	Remote NFS server
NX-API	8443	TCP	In/Out	Used by Cisco MDS 9000 Series switches with NX-OS release 9.x and later for performance monitoring.

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
SCP	22	TCP	In/Out	SCP is used by various features to transfer files between devices and Nexus Dashboard, such as for archiving backup files to remote server.. The Nexus Dashboard 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 Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
SCP	22	TCP	Out	Transport tech-support file from persistent IP address of Nexus Dashboard POAP-SCP pod to a separate ND cluster running telemetry. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings
SMTP	25	TCP	Out	You can configure the SMTP port on the **Admin** > **Server Settings** > **General** page. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from Nexus Dashboard to devices.

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection
SNMP Trap	2162	UDP	In	SNMP traps from devices to Nexus Dashboard are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings,
SSH	22	TCP	Out	In-band of switches and APIC
SSH	1022	TCP/UDP	In/Out	Other cluster nodes
Syslog	514	UDP	In	When Nexus Dashboard is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
VXLAN	4789	UDP	In/Out	Other cluster nodes

This table lists the ports that the Nexus Dashboard SAN deployments on single-node clusters require.

Table 6: Nexus Dashboard ports for SAN deployments on single-node clusters

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
GRPC (Telemetry)	33000	TCP	In	SAN Telemetry Server which receives SAN data (such as storage, hosts, flows, and so on) over GRPC transport tied to Nexus Dashboard persistent IP address.
HTTPS (vCenter, Kubernetes, OpenStack, Discovery)	443	TCP	Out	Nexus Dashboard 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.
SCP	22	TCP	In/Out	SCP is used by various features to transfer files between devices and Nexus Dashboard, such as for archiving backup files to remote server.. The Nexus Dashboard 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 Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.
SCP	22	TCP	Out	Transport tech-support file from persistent IP address of Nexus Dashboard POAP-SCP pod to a separate ND cluster running telemetry. The SCP-POAP service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings

Service	Port	Protocol	Direction `In`—toward the cluster `Out`—from the cluster toward the fabric or outside world	Connection (Applies to both LAN and SAN deployments, unless stated otherwise)
SMTP	25	TCP	Out	You can configure the SMTP port on the **Admin** > **Server Settings** > **General** page. This is an optional feature.
SNMP	161	TCP/UDP	Out	SNMP traffic from Nexus Dashboard to devices.
SNMP Trap	2162	UDP	In	SNMP traps from devices to Nexus Dashboard are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings,
SSH	22	TCP	Out	SSH is a basic mechanism for accessing devices.
Syslog	514	UDP	In	When Nexus Dashboard is configured as a Syslog server, Syslogs from the devices are sent out toward the persistent IP address associated with the SNMP-Trap/Syslog service pod. The SNMP-Trap-Syslog service in Nexus Dashboard has a persistent IP address that is associated with either the management or data subnet. This is controlled by the **LAN Device Management Connectivity** setting in the Nexus Dashboard server settings.

Nexus Dashboard persistent IP addresses

Persistent IP addresses, also known as external service IP addresses, are IP addresses that are used for various controller and telemetry functions within the Nexus Dashboard cluster. The word “persistent” is used because while the service may move between different Nexus Dashboard nodes in the event of a node or pod failure, the IP address for the service being referred by the switches in the fabric is preserved. This ensures that no

configuration updates to the switches are needed in case of a Nexus Dashboard-related failure event. Persistent IP addresses can be programmed in both the management and data subnets depending on the features deployed.

You can view the configured persistent IP addresses on your Nexus Dashboard by navigating to:

Admin > **System Settings** > **General**

Locate the **External Pools** tile and click **View all** at the bottom left area of the **External Pools** tile to view the configured persistent IP addresses on your Nexus Dashboard.

Persistent IP address updates in release 4.x

This section provides information on the changes in Nexus Dashboard release 4.x for the persistent IP addresses. It also explains how to make certain updates in the number of persistent IP addresses before proceeding to upgrade to Nexus Dashboard release 4.x.

- Reduction of number of IP addresses needed

  Starting with release 4.1.1, some services that needed exclusive IP addresses in previous Nexus Dashboard releases were merged with others. For example, on a per Nexus Dashboard node basis, software telemetry, flow telemetry and IPFM telemetry collectors on the data network have been merged in such a way that one IP address per collector service is sufficient to serve all three functions.

- LAN Device Connectivity

  You can set the type of LAN device connectivity (Data or Management) under **Admin** > **System Settings** > **Fabric management** > **Advanced settings** > **Admin** > **LAN Device Management Connectivity**.

  Prior to release 4.1.1, the default setting for LAN Device connectivity was Management. Beginning with release 4.1.1, this default has been changed to Data. However, when upgrading from Nexus Dashboard release 3.2.x to 4.x, any user-configured setting for LAN device connectivity is preserved.

- Layer-3 Persistent IP Support for Telemetry

  Starting with Nexus Dashboard release 4.1.1, telemetry-collector-X persistent IPs are supported in layer-3 adjacent Nexus Dashboard clusters. See below for Layer 3 BGP deployment details.

The number of persistent IP addresses and how they are mapped to services has changed in Nexus Dashboard release 4.x. The following services consume persistent IP addresses on Nexus Dashboard:

LAN deployments:

- Telemetry collector-x: There is a requirement for 1 persistent IP address for a 1-node cluster, or 3 persistent IP addresses for a 3-node or larger cluster, on the data network.
- SNMP trap and syslog receiver: 1 persistent IP address on the data network if the LAN device connectivity type is set to Data or 1 persistent IP address on the management network if the LAN device connectivity type is set to Management.
- Switch bootstrap service (POAP/PnP): 1 persistent IP address on the data network if the LAN device connectivity type is set to Data or 1 persistent IP address on the management network if the LAN device connectivity type is set to Management.
- (Optional) Endpoint Locator (EPL): 1 persistent IP address on the data network for each fabric where EPL is enabled. The EPL feature can be enabled for up to 4 fabrics in a given Nexus Dashboard cluster.
- (Optional) IPFM (IP Fabric for Media) telemetry collector-x: If LAN device connectivity is set to Data, then no additional persistent IPs are required. However if LAN device connectivity type is set to

Management, there is a requirement for 1 persistent IP address for a 1-node cluster, or 3 persistent IP addresses for a 3-node or larger cluster, on the management network.

SAN deployments:

- SNMP trap and syslog receiver: 1 persistent IP address on the data network.
- Switch bootstrap service: 1 persistent IP address on the data network.
- (Optional) SAN Insights receiver-x: There is a requirement for 1 persistent IP address for a 1-node cluster, or 3 persistent IP addresses for a 3-node or larger cluster, on the data network.

For an IPv4-only Nexus Dashboard cluster deployment, each service listed above will consume 1 persistent IPv4 address. For an IPv6-only Nexus Dashboard cluster deployment, each service will consume 1 persistent IPv6 address. For a dual-stack Nexus Dashboard cluster deployment, each service needing a persistent IP will consume an IPv4 address and an IPv6 address.

Fresh installation or upgrade

The total number of persistent IP addresses that you will need won't necessarily change based on whether you are performing a fresh installation or an upgrade. In addition, for a fresh installation of Nexus Dashboard (greenfield deployment), you must configure persistent IP addresses only on the data network. You can change the LAN device connectivity type from Data to Management after the cluster installation is complete.

As mentioned earlier, there is a change in the default setting in Nexus Dashboard 4.x compared to previous Nexus Dashboard releases. In Nexus Dashboard 4.x, the default LAN Device Management Connectivity is set to Data, while in earlier releases it was set to Management. As part of the effort towards delivering the unified Nexus Dashboard, the goal is to make the recommended best practice deployment as easy as possible. The reachability from Nexus Dashboard to and from switches is recommended to be over the Nexus Dashboard data interface. The Nexus Dashboard management interface should primarily be used for UI/API access and reachability for AAA, DNS, Proxy, NTP, Intersight, and so on. Finally, note that the connectivity setting set by the user is preserved when doing an inline upgrade from Nexus Dashboard release 3.2.x to Nexus Dashboard release 4.x.

Additional considerations to keep in mind

Along with the factors listed above, there are several additional considerations that you should keep in mind regarding persistent IP addresses:

- Nexus Dashboard deployment mode:
  - Layer 2: Here the Nexus Dashboard nodes within the cluster are layer-2 adjacent. This means that all Nexus Dashboard nodes share the same management and data subnet respectively. Persistent IP addresses need to be on the same network as the data network or management network.
  - Layer 3 BGP: In this mode, the Nexus Dashboard nodes within the cluster are layer-3 adjacent. In other words, unique management and data subnets are associated with each Nexus Dashboard node in the cluster. There needs to be IP reachability between the nodes to form the cluster. Persistent IP addresses cannot be from a subnet that belongs to any of the Nexus Dashboard nodes' Data or Management interface subnets. In this case, LAN Device Management Connectivity must be set to Data and cannot be changed.

Updates to mapping

The mapping for persistent IP addresses has been updated to show correct service names.

Older	Newer
cisco-nir-collectorpersistent1-service	Telemetry collector-1
cisco-nir-collectorpersistent2-service	Telemetry collector-2
cisco-nir-collectorpersistent3-service	Telemetry collector-3
cisco-ndfc-dcnm-poap-data-http-ssh	Switch Bootstrap server
cisco-ndfc-dcnm-poap-mgmt-http-ssh	Switch Bootstrap server
cisco-ndfc-dcnm-syslog-trap-data	SNMP trap and syslog receiver
cisco-ndfc-dcnm-syslog-trap-mgmt	SNMP trap and syslog receiver
cisco-ndfc-pmn-telemetry-mgmt-worker-0	IPFM telemetry collector-1
cisco-ndfc-pmn-telemetry-mgmt-worker-1	IPFM telemetry collector-2
cisco-ndfc-pmn-telemetry-mgmt-worker-2	IPFM telemetry collector-3
cisco-ndfc-dcnm-san-insight-receiver-1	SAN Insights receiver-1
cisco-ndfc-dcnm-san-insight-receiver-2	SAN Insights receiver-2
cisco-ndfc-dcnm-san-insight-receiver-3	SAN Insights receiver-3

Determining the total number of persistent IP addresses that you will need

All the factors listed above come into play when you are trying to determine the total number of persistent IP addresses that you will need and which network they come from. Make sure you review you final Nexus Dashboard deployment configuration to verify that you have enough persistent IP addresses in the proper subnet range for your deployment, and that you have additional persistent IP addresses if necessary, depending on the type of LAN device connectivity that you set and for services that you might enable, such as Endpoint Locator (EPL).

Following is an example scenario that demonstrates how persistent IP addresses are used:

Fresh installation:

First, at cluster bringup, you will need a certain number of persistent IP addresses on the data network, based on the cluster size, as mentioned earlier:

- **1-node cluster, with physical or virtual nodes**: Minimum of **3** persistent IP addresses needed on the data network
- **3-node or larger cluster, with physical or virtual nodes**: Minimum of **5** persistent IP addresses needed on the data network

Note These values are valid for either IPv4 or IPv6, but double the number for dual-stack IPv4 and IPv6. For example, for a 3-node or larger cluster, a minimum of 10 persistent IP addresses would be needed on the data network for dual-stack IPv4 and IPv6 (5 IPv4 and 5 IPv6 persistent IP addresses).

After bootstrapping, you may need to add additional persistent IP addresses as needed, depending on these scenarios:

- If you leave the LAN device connectivity type set to **Data**, you won't need any additional persistent IP addresses unless you enable the Endpoint Locator (EPL) feature, which requires 1 additional persistent IP address on the data network per fabric where EPL is enabled.
- If you change the LAN device connectivity type from **Data** to **Management**:
  - You will need 2 additional persistent IP addresses on the management network for Syslog/SNMP trap and switch bootstrap functionality.
  - (Optional) If you want to enable Endpoint Locator (EPL), you will need 1 persistent IP address on the data network per fabric where EPL is enabled.
  - (Optional) If you want IP Fabric for Media (IPFM) fabrics, you will need 1 persistent IP address for a 1-node cluster, or 3 persistent IP addresses for a 3-node or larger cluster, on the management network.

Table 7: Persistent IP requirements: Fresh installation of 4.x

Number of ND nodes	LAN Device Connectivity	Mandatory persistent IP addresses	Optional persistent IP addresses	Other common persistent IP addresses
1	Data [1]	3 in data network	N/A	1 in data network per fabric where EPL is enabled
	Management	2 in management network 1 in data network	1 in management network for IPFM fabrics	
3 or more	Data [1]	5 in data network	N/A	
	Management	2 in management network 3 in data network	3 in management network for IPFM fabrics	

[1] Indicates default option set during ND bootstrap process

Upgrade:

Now assume that you want to upgrade from Nexus Dashboard 3.2.x to 4.x. The number of persistent IP addresses that you will need in Nexus Dashboard 4.x will vary, depending on the services that you were running and how they were configured on Nexus Dashboard 3.2.x, and the size of your cluster in Nexus Dashboard 4.x. Note that the LAN Device Management Connectivity that you set in the Nexus Dashboard 3.2.x release is preserved as-is when performing an inline upgrade to the Nexus Dashboard 4.x release.

- If you had only **NDFC** running in your Nexus Dashboard 3.2.x system, and
  - If you had **Data** set as the type of LAN device connectivity in Nexus Dashboard 3.2.x, and
    - You have a 1-node cluster that you want to upgrade to Nexus Dashboard 4.x, then you'll need 3 persistent IP addresses on the data network.
    - You have a 3-node or larger cluster that you want to upgrade to Nexus Dashboard 4.x, then you'll need 5 persistent IP addresses on the data network.

    - If you had **Management** set as the type of LAN device connectivity in Nexus Dashboard 3.2.x, and
        - You have a 1-node cluster that you want to upgrade to in Nexus Dashboard 4.x, then you would already have either 2 or 3 persistent IP addresses (additional IP is required if IPFM/PTP feature was enabled) in the management network. In addition, you will need 1 persistent IP address in the data network, otherwise the upgrade to 4.x will fail during the pre-upgrade validation step.
        - You have a 3-node or larger cluster that you want to upgrade to in Nexus Dashboard 4.x, then you would already have either 2 or 5 persistent IP addresses (3 additional IPs are required if IPFM/PTP feature was enabled) in the management network. You will need to configure 3 persistent IP addresses in the data network; only then will the upgrade to 4.x proceed.

- If you had only **NDI** running in your Nexus Dashboard 3.2.x system, and
    - You have a 1-node cluster that you want to upgrade to Nexus Dashboard 4.x, then you would already have configured 4 persistent IP addresses in the data network. After the upgrade to 4.x, only 3 persistent IP addresses will be in use. The remaining can be reclaimed.
    - You have a 3-node or larger cluster that you want to upgrade to Nexus Dashboard 4.x, then you would already have configured 8 persistent IP addresses in the data network and 2 additional data IPs for support of standalone NX-OS deployments. After the upgrade to 4.x, only 5 of these data IP addresses will be in use. The remaining can be reclaimed.
- If you had only **NDO** running in your Nexus Dashboard 3.2.x system, then you did not have any persistent IP addresses in your Nexus Dashboard 3.2.x system. When you upgrade to Nexus Dashboard 4.x, if you have a 3-node cluster that you want to upgrade to Nexus Dashboard 4.x, then you'll need 5 persistent IP addresses on the data network before the upgrade can proceed.
- If you had an **NDO** and **NDI** deployment mode in your Nexus Dashboard 3.2.x system, and you have a 3-node or larger cluster that you want to upgrade to Nexus Dashboard 4.x, then you would already have configured 8 persistent IP addresses in the data network. After the upgrade to 4.x, only 5 of these data persistent IP addresses will be in use. The remaining persistent IP addresses can be reclaimed.
- If you had only **NDFC** and **NDI** deployment mode in your Nexus Dashboard 3.2.x system, and you have a 3-node or larger physical ND cluster that you want to upgrade to Nexus Dashboard 4.x, then there were two options based on the LAN Device Management Connectivity setting:
    - If you had **Management** set as the type of LAN device connectivity in Nexus Dashboard 3.2.x, you would already have configured 8 persistent IP addresses in the data network for NDI and 2 persistent IP addresses in the management network for NDFC. After the upgrade to 4.x, the 2 persistent IP addresses in the management subnet will be used along with only 3 of the data persistent IP addresses. The remaining persistent IP addresses can be reclaimed.
    - If you had **Data** set as the type of LAN device connectivity in Nexus Dashboard 3.2.x, you would already have configured 8 persistent IP addresses in the data network for NDI and 2 additional data persistent IP addresses for NDFC. After the upgrade to 4.x, only 5 of these data persistent IP addresses will be in use. The remaining persistent IP addresses can be reclaimed.

Table 8: Persistent IP requirements: Upgrade from 3.2.x to 4.x

ND 3.2.x deployment mode	Number of ND nodes	LAN Device Connectivity	ND 3.2.x persistent IP address requirement	ND 4.x persistent IP address requirement
NDFC	1	Data	2 in data network, plus 1 in data network if IPFM/PTP is enabled	3 in data network
		Management	2 in management network, plus 1 in management network if IPFM/PTP is enabled	2 in management network, plus 1 in management network for IPFM fabrics 1 in data network
NDFC	3 or more	Data	2 in data network, plus 3 in data network if IPFM/PTP is enabled	5 in data network
		Management	2 in management network, plus 3 in management network if IPFM/PTP is enabled	2 in management network, plus 3 in management network for IPFM fabrics 3 in data network
NDFC + NDI	3 physical	Data	10 in data network	5 in data network
		Management	2 in management network 8 in data network	2 in management network, plus 3 in management network for IPFM fabrics 3 in data network
NDI	1	N/A	3 in data network	3 in data network
NDI	3 or more	N/A	10 in data network	5 in data network
NDO	3	N/A	None	5 in data network
NDO + NDI	3 or more	N/A	8 in data network	5 in data network

Note EPL persistent IP address requirements remain the same in release 4.x as they were in release 3.2.x.

BGP configuration and persistent IP addresses

Some prior releases of Nexus Dashboard allowed you to configure one or more persistent IP addresses 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 address functionality, you must:

- Ensure that the peer routers exchange the advertised persistent IP addresses between the nodes' Layer 3 networks.
- For data network Layer 3 adjacency, you must configure BGP during the bootstrap process. BGP does not support management network Layer 3 adjacency.
- Ensure that the persistent IP addresses you allocate do not overlap with any of the nodes' management or data subnets.

Round trip time requirements

Connectivity between the nodes is required on both networks with additional round trip time (RTT) requirements, as listed in this table.

Table 9: Cluster round trip time requirements

Connectivity	Maximum RTT
Between nodes within the same Nexus Dashboard cluster (intra-cluster connectivity)	50 ms
Between nodes in one cluster and nodes in a different cluster if the clusters are connected using multi-cluster connectivity (inter-cluster connectivity) For more information about multi-cluster connectivity, see *Connecting Clusters*.	500 ms
Between external DNS servers and the Nexus Dashboard cluster	5 seconds
Between cluster nodes and switches	150 ms

Prerequisites for Nexus Dashboard storage (virtual form factors)

This section of the document provides guidance for designing storage infrastructure for Cisco Nexus Dashboard deployments (OVA, KVM, Nutanix, and so on).

The Nexus Dashboard offers features that process high-frequency telemetry and flow data from network devices. As a result, storage plays a critical role in ensuring:

- Reliable data ingestion
- Accurate analytics and insights
- Consistent system performance at scale

This section outlines the recommended storage characteristics, performance expectations, and architectural considerations for production environments.

Key workload characteristics

These are the key workload characteristics:

- Write-intensive ingestion with sustained throughput requirements
- Mixed I/O patterns:
  - Sequential writes (data ingestion)
  - Random reads/writes (queries and analytics)
- Latency-sensitive operations, particularly during ingestion

Baseline storage requirements

These are the baseline storage requirements:

Capacity

- Storage capacity: Two storage disks are required:
  - One disk of size 50G used as boot disk
  - Another disk to store data
- The size of data depends on platform flavor. For example: Virtual-data requires 3 TB of disk.

Performance

Designing storage for Nexus Dashboard requires careful attention to performance, not just capacity. You should ensure:

- Sync latency is less than 10ms.
- The actual throughput and IOPS requirement depends on scale and can vary from:
  - IOPS: 1.5K io/s to 80K io/s sustained
  - Throughput: 10 MB/s to 1.5 GB/s sustained

You can also verify disk I/O latency using the `acs diag diskio` command from Nexus Dashboard.

Recommendations for designing storage

To meet these requirements, these are the recommendations for designing storage for Nexus Dashboard.

Storage media selection and placement

For production deployments, we strongly recommend using NVMe or SSDs. From a placement point of view, direct-attached storage (DAS) typically provides the most predictable results. If you are using RAID controllers, configure in JBOD mode to minimize overhead.

Using remote or shared storage

You can use remote storage solutions (such as SAN, NAS, or software-defined storage platforms) if they consistently meet performance requirements. Note that storage must deliver required IOPS and throughput continuously, not just during short bursts. In many cases, locally attached storage provides the most consistent and predictable performance for Nexus Dashboard workloads.

- In case of shared storage, performance variability may occur due to other workloads and resource contention can introduce latency spikes.
- Therefore, we recommend that you have dedicated disks and not shared along with other nodes within the Nexus Dashboard cluster or with other workloads.

Caching

- From the Nexus Dashboard point of view, every write is flushed via hypervisor; therefore, caching plays a critical role to achieve performance.
- Without a cache, every single 4KB write has to be written all the way to the physical medium. You should always enable cache policy to write back instead of Write Through.
- Note that hardware vendors typically will set the default to Write Through to guarantee data integrity. For example, if the power cord is pulled or the building loses power at that exact second, the data is already on the disk. It is completely safe.
- Modern systems use battery (BBU) and they use cache only if the battery or SuperCap is healthy. Make sure that your disk controller is provisioned to fallback automatically to Write Through if the battery fails or is charging.

Fabric connectivity

These 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 more information on configuring the fabric to enable the in-band telemetry functions, see the following documents:

- *Getting Your Cisco ACI Fabrics Ready for Cisco Nexus Dashboard Insights*
- *Deploying Nexus Fabrics with Telemetry on Cisco Nexus Dashboard*

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

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

Connecting using an external Layer 3 network

We recommend connecting the Nexus Dashboard cluster to the fabrics using 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 using orchestration 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 fabric's APIC or both.

  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 using telemetry, 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 the Cisco APIC Layer 3 Networking Configuration Guide.

- 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 using a Layer 3 network, where the first figure shows a mix of ACI and NX-OS fabrics, and the second figure shows only ACI fabrics.

Figure 1: Connecting using Layer 3 network, With a mix of ACI and NX-OS fabrics

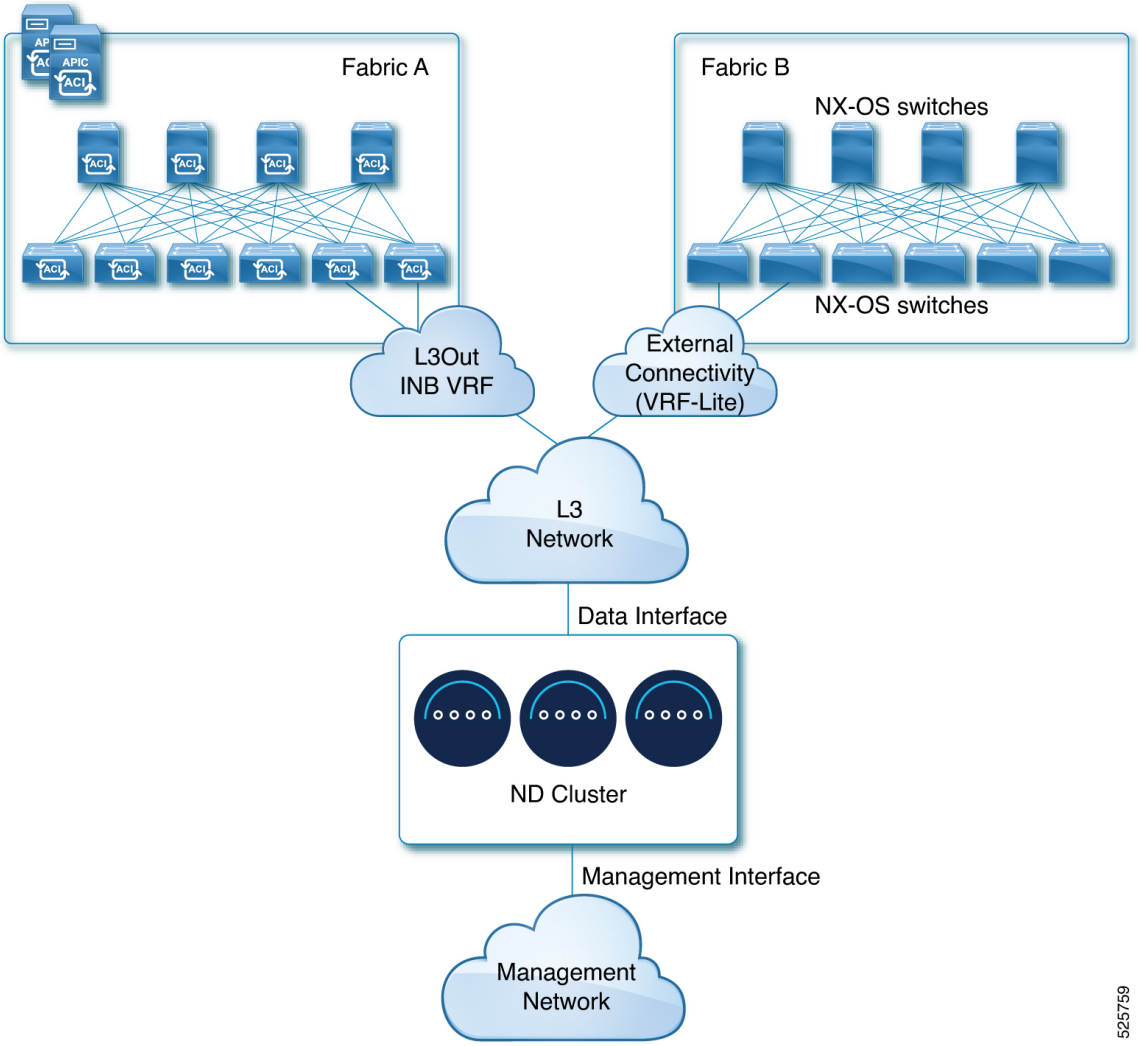
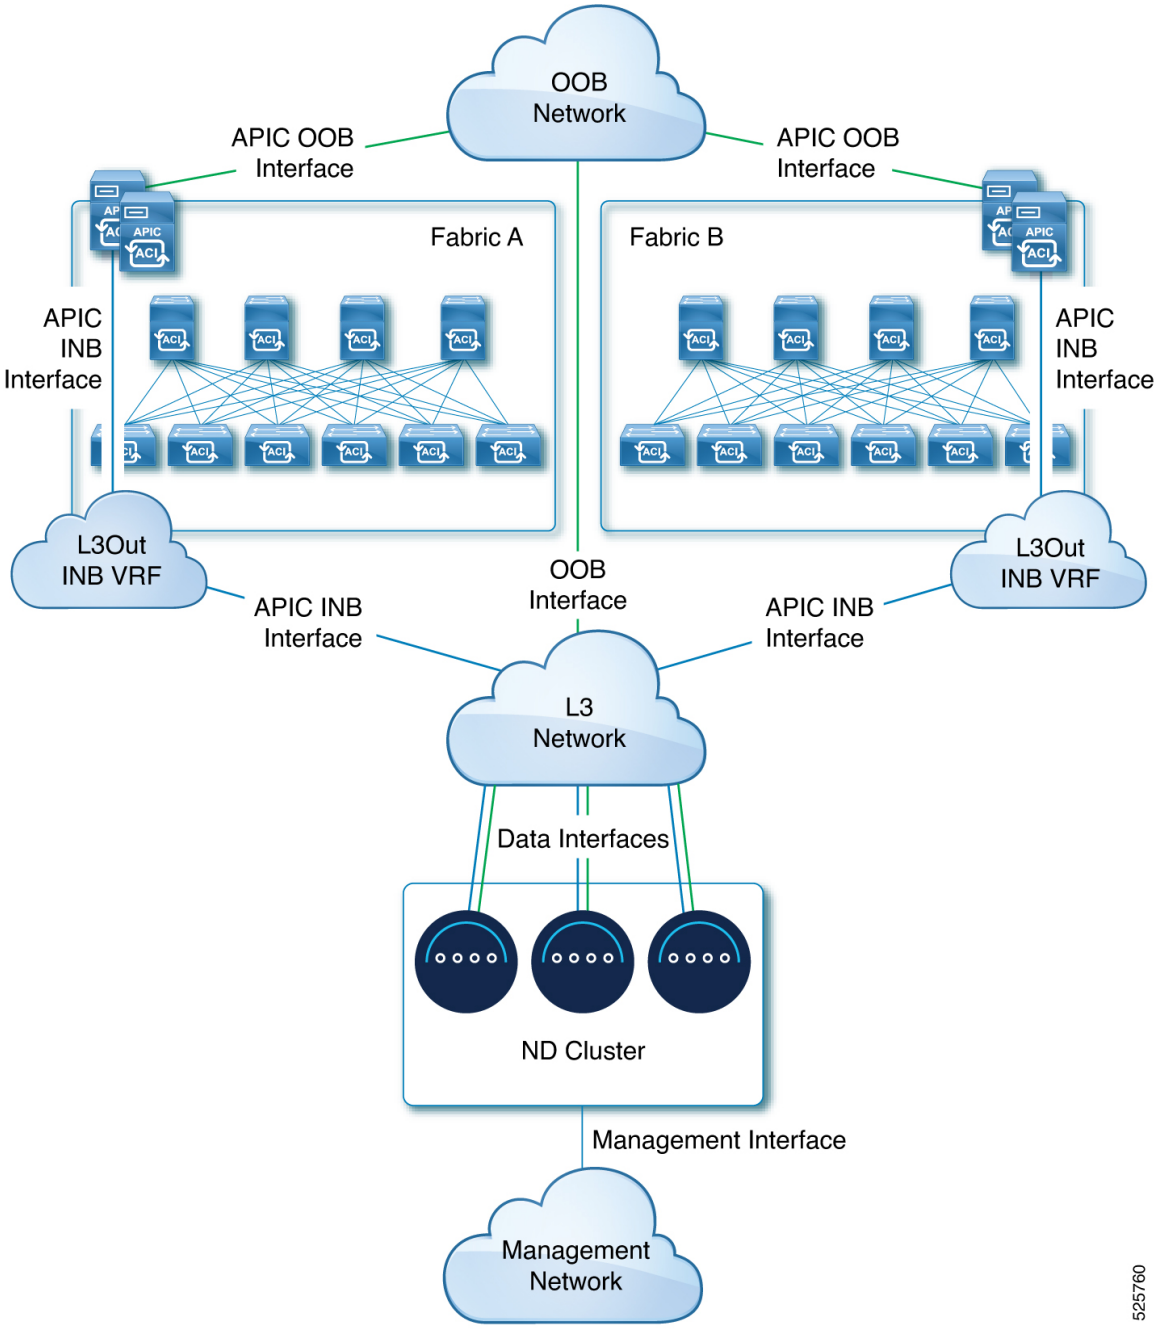


Figure 2: Connecting using Layer 3 network, with ACI fabrics only



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 using orchestration 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 fabric's APIC or both.

  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 using telemetry, you can establish connectivity from the data interface to the in-band or out-of-band (OOB) interface of each fabric. However, you must add the route if you establish connectivity from the data interface to the out-of-band interface.

  For ACI fabrics, the data interface IP subnet connects to an EPG/ or bridge domain 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 using an L3Out.

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 using 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 configurations on the APIC side, following are recommended configurations:
  - We recommend configuring the bridge domain, 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 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.

The following graphics show these types of connections:

- Connecting directly to ACI fabric
- Connecting directly to NX-OS fabric
- Connecting directly to ACI and NX-OS fabrics

Figure 3: Connecting Directly to ACI Fabric

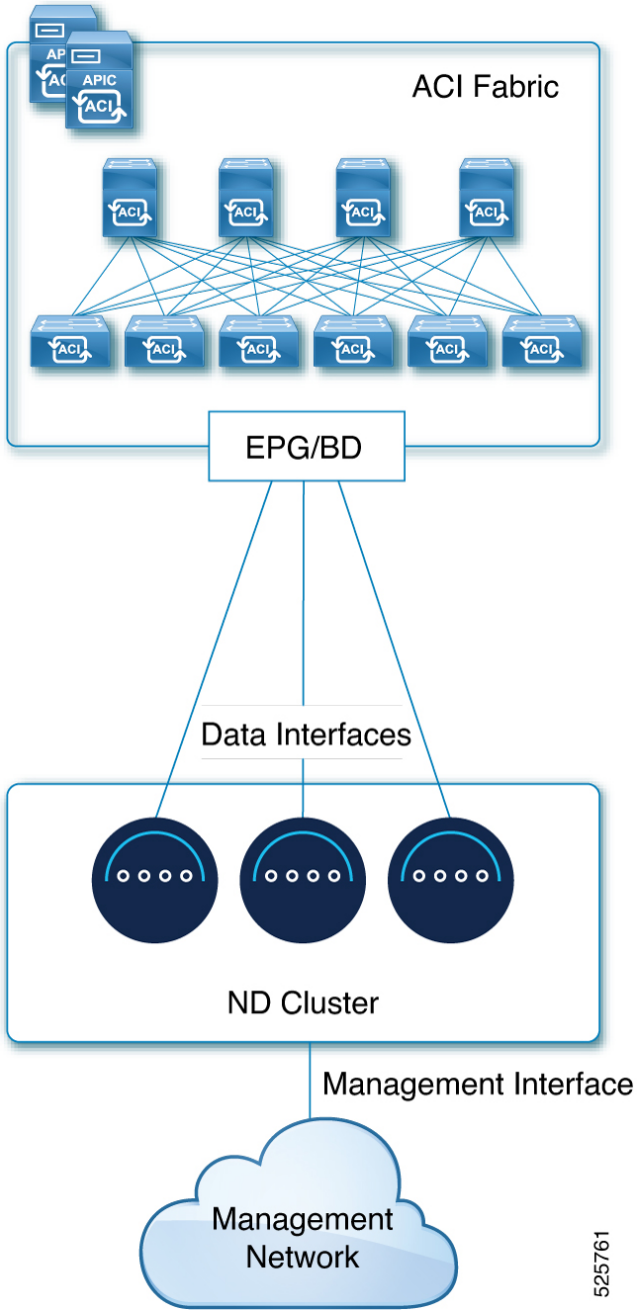


Figure 4: Connecting Directly to NX-OS Fabric

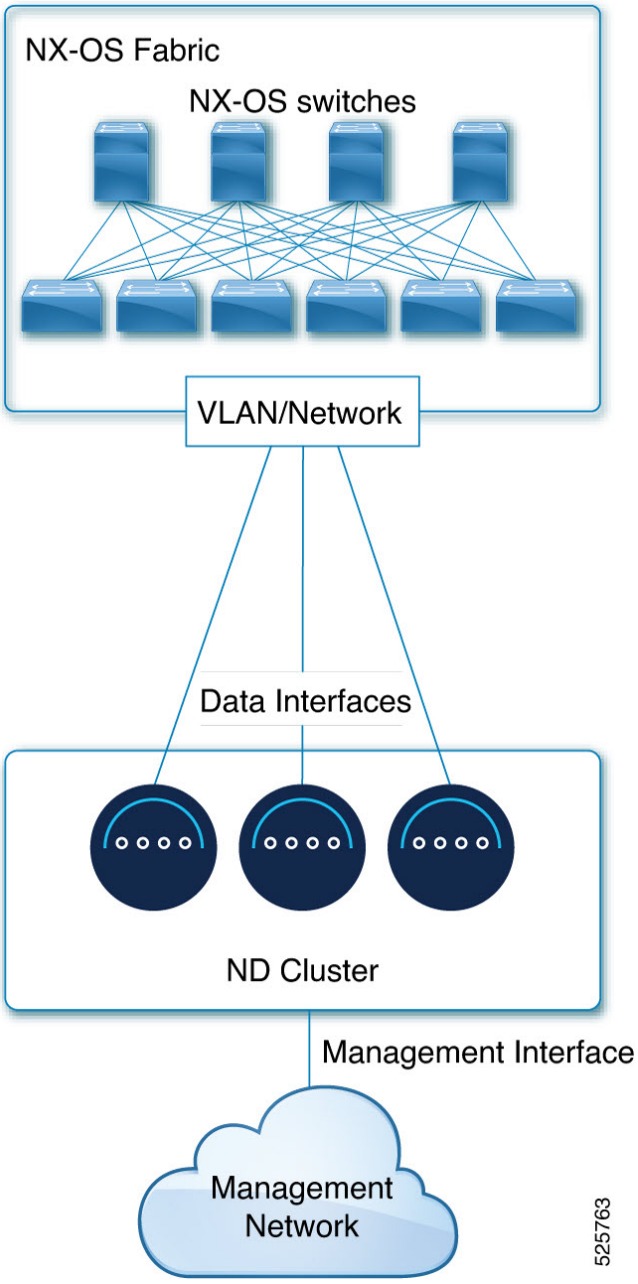


Figure 5: Connecting Directly to ACI and NX-OS Fabrics

