



Prerequisites

This chapter provides release-specific prerequisites information for your deployment of *Cisco Nexus Dashboard Fabric Controller*.

- [Prerequisites, on page 1](#)

Prerequisites

Before you install the Cisco Nexus Dashboard Fabric Controller on Cisco Nexus Dashboard, you must need to meet the following prerequisites:

Nexus Dashboard

You must have Cisco Nexus Dashboard cluster deployed and its fabric connectivity configured, as described in [before proceeding with any additional requirements and the Nexus Dashboard Fabric Controller service installation described here](#).

Nexus Dashboard Fabric Controller Release	Minimum Nexus Dashboard Release
Release 12.1.1p	Cisco Nexus Dashboard, Release or later Note Cisco Nexus Dashboard cluster in Linux KVM does not support Nexus Dashboard Fabric Controller Release 12.1.1p.

Nexus Dashboard Networks

When first configuring Nexus Dashboard, on every node, you must provide two IP addresses for the two Nexus Dashboard interfaces—one connected to the Data Network and the other to the Management Network. The data network is typically used for the nodes' clustering and north-south connectivity to the physical network. The management network typically connects to the Cisco Nexus Dashboard Web UI, CLI, or API.

For enabling the Nexus Dashboard Fabric Controller, the Management and Data Interfaces on a Nexus Dashboard node must be in different subnets. Different nodes that belong to the same Nexus Dashboard cluster can either be Layer-2 adjacent or Layer-3 adjacent. Refer to [Layer 3 Reachability Between Cluster Nodes](#) for more information.

Connectivity between the Nexus Dashboard nodes is required on both networks with the round trip time (RTT) not exceeding 50ms. Other applications running on the same Nexus Dashboard cluster may have lower RTT requirements and you must always use the lowest RTT requirement when deploying multiple applications in the same Nexus Dashboard cluster. Refer to for more information.

Management Interface	Data Interface	Persistent IPs
Layer 2 adjacent	Layer 2 adjacent	<p>One of the following for LAN:</p> <ul style="list-style-type: none"> • If using default LAN Device Management Connectivity (set to Management): <ul style="list-style-type: none"> • 2 IPs in management network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in management network if IP Fabric for Media is enabled • If LAN Device Management Connectivity is set to Data: <ul style="list-style-type: none"> • 2 IPs in data network for SNMP/Syslog and SCP services • Plus one IP per fabric for EPL (if enabled) in data network • Plus one IP for Telemetry receiver in data network if IP Fabric for Media is enabled
Layer 3 adjacent	Layer 3 adjacent	<p>For LAN:</p> <ul style="list-style-type: none"> • LAN Device Management Connectivity on NDFC must be set to Data • 2 IPs for SNMP/Syslog and SCP/POAP services • Plus one IP per fabric for EPL <p>These IPs must be part of a subnet that is different from Nexus Dashboard management and Nexus Dashboard data subnets associated with any of Nexus Dashboard nodes. These IPs must belong to the Layer-3 External Persistent Service Pool.</p>

Virtual Nexus Dashboard (vND) Prerequisites

For virtual Nexus Dashboard deployments, each vND node has 2 interfaces or vNICs. The Data vNIC maps to bond0 (also known as bond0br) interface and Management vNIC maps to bond1 (also known as bond1br) interface. The requirement is to enable/accept promiscuous mode on the port groups that are associated with the Nexus Dashboard Management and/or Data vNICs where IP stickiness is required. The Persistent IP addresses are given to the pods (for example, SNMP Trap or Syslog receiver, Endpoint Locator instance per

Fabric, and so on). Every POD in Kubernetes can have multiple virtual interfaces. Specifically for IP stickiness, an extra virtual interface is associated with the POD that is allocated an appropriate free IP from the external service IP pool. The vNIC has its own unique MAC address that is different from the MAC addresses associated with the vND virtual vNICs. Moreover, all North-to-South communication to and from these pods go out of the same bond interface. By default, the VMware ESXi systems check if the traffic flows out of a particular VM vNIC that matches the Source-MAC that is associated with that vNIC. If NDFC pods with an external service IP, the traffic flows are sourced with the Persistent IP addresses of the given pods that map to the individual POD MAC associated with the virtual POD interface. Therefore, enable the required settings on the VMware side to allow this traffic to flow seamlessly in and out of the vND node.

When vND nodes are deployed with the new Layer-3 HA feature, you need not enable Promiscuous mode on the vND vNIC interfaces. Promiscuous mode is required only for vND deployments when the vNDs are layer-2 adjacent from each other.

For more information, refer to .

Nexus Dashboard Cluster Sizing

Nexus Dashboard supports cohosting of services. Depending on the type and number of services you choose to run, you may be required to deploy extra worker nodes in your cluster. For cluster sizing information and recommended number of nodes based on specific use cases, see the [Cisco Nexus Dashboard Capacity Planning](#) tool.

If you plan to host other applications in addition to the Nexus Dashboard Fabric Controller, ensure that you deploy and configure additional Nexus Dashboard nodes based on the cluster sizing tool recommendation, as described in the , which is also available directly from the Nexus Dashboard Web UI.

Network Time Protocol (NTP)

Nexus Dashboard Fabric Controller uses NTP for clock synchronization, so you must have an NTP server configured in your environment.

Clocks on all nodes must be synchronized within the same second. Any delta between two nodes that exceeds more than 1 second could affect database consistency mechanism between the nodes.

