



Cisco Nexus Dashboard Infrastructure Management, Release 3.1.x

Table of Contents

System Settings	1
Persistent IP Addresses	5
Persistent IP Guidelines and Limitations	5
Enabling BGP On All Nodes	6
Configuring Persistent IPs	7
Multi-Cluster Connectivity	9
Guidelines and Limitations	9
Connecting Multiple Clusters	10
Disconnecting Clusters	12
Deploying Additional Physical Nodes	13
Prerequisites and Guidelines for Physical Nodes	13
Deploying Physical Nodes	15
Deploying Additional Virtual Nodes in VMware ESX	16
Prerequisites and Guidelines for ESX Nodes	16
Deploying ESX Node Using vCenter	18
Deploying ESX Node Directly in ESXi	23
Deploying Additional Virtual Nodes in Linux KVM	26
Prerequisites and Guidelines for KVM Nodes	26
Deploying KVM Nodes	27
Managing Worker Nodes	31
Adding Worker Nodes	31
Deleting a Worker node	32
Managing Standby Nodes	33
Adding Standby Nodes	33
Replacing Single Primary Node with Standby Node	34
Replacing Two Primary Nodes with Standby Nodes	35
Deleting Standby Nodes	39
Trademarks	40

System Settings

The **System Settings** GUI screen allows you to configure a number of options specific to the Nexus Dashboard cluster and its nodes. It will also display information about any issues that may be present in your Nexus Dashboard cluster.

The screenshot displays the **System Settings** page for the **ifav74-Cluster-SN123** cluster. The page is divided into several sections, each with an **Edit** button:

- General** (selected):
 - Cluster Details**: Shows Name (ifav74-Cluster-SN123), App Subnet (172.17.0.1/16), Service Subnet (100.80.0.0/16), App Subnet IPv6 (2000::/108), and Service Subnet IPv6 (3000::/108).
 - Proxy Configuration**: Shows Type (HTTP, HTTPS) and Server (http://proxy-wsa.esl.cisco.com).
 - Ignore Hosts**: Shows 2001:420:28e:2023::/64.
 - Routes**: Shows Management Network Routes (172.23.48.0/21) and Data Network Routes (100.15.15.0/24, 2001:6116:116::/64, 2001:7117:117::/64).
 - Network Scale**: Shows Number of Sites, Number of Switches, and Flows per second.
 - External Service Pools**: Shows Management Service IP Usage (0) and Data Service IP Usage (33/23 Available, 10 In Use). It includes a table for Management Service IP's and a table for Data Service IP's.
- Multi-Cluster Connectivity** (1):
 - Error** (2): 1 Error on this page. Collapse to hide. kafka service: Statefulset(kafka) not in desired state.
 - NTP** (7): Shows Key (06984f8339438b228273b82bfd8a3b6b9a6e4310) and NTP Host Name/IP Address (172.31.172.23, 2001:420:28E:2023::1:7).
 - DNS** (8): Shows Domain Name (ifav74-cluster-sn123.case.local) and Providers IP Addresses (10.195.200.12).
 - Syslog** (9): Shows Remote Destinations (172.23.50.101).
 - Network-Attached Storage** (10): Shows Name (ifav74inbv4, ifav74inbv6, ifav74mgmtv4, ifav74mgmtv6).

Figure 1. System Settings

1. The **Multi-cluster Connectivity** tab allows you to connect multiple clusters together for a single pane of glass view and administration of the clusters and their sites, services, and configurations.

For more information, see [Multi-Cluster Connectivity](#).

2. The errors and warning tile will display the number of existing issues in your cluster. You can click **Expand** to see the full list of specific issues.
3. To configure a proxy for the Nexus Dashboard, click the **Edit** icon in the **Proxy Configuration** tile.

In certain deployment scenarios, such as with a combination of on-premises and cloud sites and the Nexus Dashboard cluster deployed inside a corporate network, you may have to access the Internet and the cloud sites through a proxy.



This release supports adding a single proxy server.

Note that Nexus Dashboard uses 2 main route tables—one for the Management network and one for the Data network—and by default, it will use the routing table of the originating IP address. In other words, Nexus Dashboard will attempt to reach the proxy from the routing table of the POD/Service that is trying to use the proxy.

For example, if you configure a proxy and establish Intersight connectivity from your Nexus Dashboard and then attempt to configure AppD integration from the Insights service running in the cluster, you may get an error stating that AppD host is not reachable. This happens because the proxy is only accessible from the management interface, so in such cases you also need to add a management network route for the proxy IP address, as described in "Management Network or Data Network routes" below.

To add a proxy server:

- a. Click **+Add Server** in the proxy configuration window.
- b. From the **Type** dropdown, select the type of traffic that you want to be proxied.
- c. In the **Server** field, provide the full address for the proxy server including the port if required.

For example <http://proxy.company.com:80>.

- d. If the server requires login credentials, provide the **Username** and **Password**.
- e. (Optional) Click **Add Ignore Host** to provide any hosts that will ignore the proxy.

You can add one or more hosts with which the cluster will communicate directly bypassing the proxy.

4. To add one or more Management Network or Data Network routes, click the **Edit** icon in the **Routes** tile.

Here you can define static routes for the management or data interfaces. For example, adding **10.195.216.0/21** as a Data Network route will cause all traffic destined to that subnet to transit out of the data network interface.

- o To add a management network route, click **Add Management Network Routes** and provide the destination subnet.
- o To add a data network route, click **Add Data Network Routes** and provide the destination subnet.

5. To configure **Network Scale**, click the **Edit** icon in the **Network Scale** tile.



In this release, these settings affect the Nexus Dashboard Insights service only.

Other services do not consider these settings and support the scale limits described in their respective Verified Scalability Guides.

Modifying the network scale requires a restart of the Insights service.

- a. In the **Number of Sites** field, provide the target number of sites for your deployment that this Nexus Dashboard cluster will manage.
 - b. In the **Number of Switches** field, provide the target number of switch nodes for your deployment.
 - c. In the **Flows per second** field, provide the target number of flows for your Nexus Dashboard Insights service.
6. To add one or more External Service Pools, click the **Edit** icon in the **External Service Pools** tile.

This allows you to provide persistent IP addresses for services that require to retain the same IP addresses even in case they are relocated to a different Nexus Dashboard node.

For detailed information and configuration steps, see [Persistent IP Addresses](#).

7. To configure NTP settings, click the **Edit** icon in the **NTP** tile.

By default, the NTP server that you configured when deploying the Nexus Dashboard cluster is listed here.

You can provide additional NTP servers by clicking **+Add NTP Server**.

You can remove existing NTP server by clicking the **Delete** icon next to it. Keep in mind that at least one NTP server must be configured in your cluster.

8. To configure DNS settings, click the **Edit** icon in the **DNS** tile.

By default, the DNS server and search domain that you configured when deploying the Nexus Dashboard cluster are listed here.

You can provide additional DNS servers and search domains by clicking **+Add a Provider** or **+Add a Search Domain** respectively.

You can remove existing DNS server by clicking the **Delete** icon next to it.

9. To provide one or more **syslog** servers to stream event logs to, click the **Edit** icon in the **Syslog** tile.

In the **Syslog** dialog that opens, click **+Add Remote Destinations** to add a new server. Then provide the IP address, protocol, and port number for the server and choose whether you want to enable streaming to this syslog server at this time.

For more information, see [\[Event Analytics\]](#).

10. To configure Network-Attached Storage (NAS), click the **Edit** icon in the **Network-Attached Storage** tile.

Beginning with release 3.0(1), you can add a NAS server at the Nexus Dashboard level which can be utilized by the services running in your cluster.



This release supports only NFSv3 remote storage.

To add a NAS:

- a. Click **+Add Network-Attached Storage** in the NAS configuration window.
- b. Choose whether Nexus Dashboard has **Read Only** or **Read Write** access to this server.
- c. Provide the **Name** for the NAS server.
- d. (Optional) Provide the **Description**.
- e. Provide the **IP address** used to connect to the server.
- f. Specify the **Port** used to establish the connection if it is different from the default port **2049**.
- g. Provide the **Export Path** to a directory on the NAS server where information will be stored or read.
- h. Specify the **Alert Threshold**.
- i. Specify the storage **Limit**.

This limits the amount of storage that can be requested on the server by Nexus Dashboard. You can provide the values in Mibibytes or Gibibytes, for example **300Mi** or **10Gi**.

- j. From the **Allowed Apps** dropdown, select which Nexus Dashboard services can access this storage.

Persistent IP Addresses

You can provide persistent IP addresses for services that require to retain the same IP addresses even in case they are relocated to a different Nexus Dashboard node.

Nexus Dashboard Insights requires some services (such as SNMP trap, syslog, and others) to stream data from the switches in your fabrics to the service. An IP address is configured on the switches for this purpose. Typically, if the IP address changes when the service is relocated, the service will reconfigure the new IP address on the switches.

In order to avoid this IP reconfiguration impact on the fabric switches, the service can request that the services IP addresses are preserved, in which case you will need to define a set of IP addresses which can be assigned to the service for this purpose.

If a service requires persistent IP addresses, you will not be able to enable that service in the Nexus Dashboard until enough IP addresses are defined as described below.



This feature is supported for Nexus Dashboard Insights with NDFC fabrics only. In addition, if you are using Layer 2 functionality only (IPs configured as part of the management and data subnets) and your Nexus Dashboard is deployed in VMware ESX, you must enable promiscuous mode for both management and data network interface portgroups, as described in <https://kb.vmware.com/s/article/1004099>.

Prior to Release 2.2(1), this feature was supported only for clusters where all nodes were part of the same Layer 3 network and the persistent IPs were defined as part of the node's management or data networks. Here the application uses Layer 2 mechanisms like Gratuitous ARP or Neighbor Discovery to advertise the persistent IPs within its Layer 3 network.

Beginning with Release 2.2(1), the feature is supported even if you deploy the cluster nodes in different Layer 3 networks. In this case, the persistent IPs are advertised out of each node's data links via BGP, which we refer to as "Layer 3 mode". The IPs must not overlap with any of the nodes' management or data subnets. If the persistent IPs are outside the data and management networks, this feature will operate in Layer 3 mode by default; if the IPs are part of those networks, the feature will operate in Layer 2 mode.

Persistent IP Guidelines and Limitations

When configuring persistent IPs for your services:

- Ensure that you check the documentation for the services you plan to deploy as some services do not support this feature or require additional guidelines.

At this time, Persistent IPs are supported for Nexus Dashboard Insights and Nexus Dashboard Fabric Controller. You can find the service-specific documentation at the following links:

- [Nexus Dashboard Fabric Controller](#)
- [Nexus Dashboard Insights](#)
- You can choose which mode you want to operate in as long as the following conditions apply:
 - If you choose to operate in Layer 2 mode, the nodes must be part of the same data and management networks.

- If you choose to operate in Layer 3 mode, all nodes must have BGP configuration provided either during cluster deployment or after as described in [Enabling BGP On All Nodes](#).
- You can switch between the two modes, in which case the existing services of a particular mode must be completely deleted and you will need to configure new persistent IPs corresponding to the new mode.
- If you configure one or more persistent IPs in Layer 3 mode and back up cluster configuration, the BGP settings required for this feature are not saved in the backup.

As such, you must ensure that you configure BGP for all cluster nodes before restoring any cluster configuration that contains Layer 3 persistent IPs in that cluster. If BGP is not configured prior to the configuration import, the import will fail.

Enabling BGP On All Nodes

If you want to operate in Layer 3 mode, you must enable and configure BGP for all nodes in your cluster. If you already configured BGP for each node during cluster deployment or if you want to operate in Layer 2 mode instead, you can skip this section and simply provide one or more persistent IPs from the nodes' management and data subnets, as described in [Configuring Persistent IPs](#). Note that if you choose to operate in Layer 2 mode, all nodes must be part of the same Layer 3 network. If you choose to operate in Layer 3 mode, at least one BGP peer must be configured on all cluster nodes to advertise the IPv4 or the IPv6 persistent IP addresses as described in this section.

Before you begin

- Ensure that the uplink peer routers are capable of exchanging the advertised persistent IPs across the Layer 3 networks of the cluster nodes.
- When a service requests a persistent IP address, the route advertised from the data links via BGP on the node where the service is running is maintained throughout the lifecycle of the service.

To configure BGP on the nodes:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to one of the nodes and choose **Edit**.
4. In the **Edit Node** screen, turn on **Enable BGP**.
5. In the **ASN** field, provide the autonomous system number for the node.
6. Click **+Add IPv4 BGP Peer** or **+Add IPv6 BGP Peer** to provide peer IP address information.
 - a. In the **Peer Address** field, provide the IPv4 or IPv6 address of the peer router for this node.

Multi-hop BGP peering is not supported, so you must ensure that the **Peer Address** is part of the node's data subnet.

- b. In the **Peer ASN** field, provide the autonomous system number of the peer router.

Only EBGp is supported, so you must ensure that the node ASN and Peer ASN are different.

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

7. Repeat these steps for every node in the cluster.

Every node in the cluster must have BGP configured.

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

Configuring Persistent IPs

Before you begin

- For all persistent IPs, you must use either the Layer 2 or Layer 3 approach; a combination of the two is not supported.

If all nodes are in the same Layer 3 network, you can choose to use either the Layer 2 mode or Layer 3 mode for this feature. The two modes are described in [Persistent IP Addresses](#).

If the nodes are in different Layer 3 networks, you must configure the persistent IPs such that they don't overlap with either the management or the data subnets of the nodes.

- If the nodes in your cluster belong to different Layer 3 networks, you must have BGP enabled and configured as described in [Enabling BGP On All Nodes](#).
- There may be a momentary traffic interruption while a service using a persistent IP is relocated to a different node.

The interruption duration depends on the following factors:

- Time to detect the node failure
- Time for the service to get rescheduled to a different node
- Time for the service's external IP to get advertised from the scheduled node via GARP (IPv4) or neighbor discovery (IPv6) in case of Layer 2 mode
- Time for the service's external IP to get advertised from the scheduled node via BGP in case of layer 3 mode

To provide one or more persistent IP addresses:

1. Navigate to your Nexus Dashboard's **Admin Console**.
2. From the left navigation menu, select **Admin > System Settings**.
3. In the **External Service Pools** tile, click the **Edit** icon.
4. In the **External Service Pools** screen that opens, click **+Add IP Address** to add one or more IP addresses for the management or data networks.

When editing persistent IPs, the following rules apply:

- If all nodes in your cluster are part of the same Layer 3 network, you can choose one of the following:
 - Layer 2 mode, in which case the IP addresses you add for management services must be part of the management subnet and the IP addresses for data services must be part of the data subnet.
 - Layer 3 mode, in which case the IP addresses you add must not overlap with the management or the data subnets of the nodes. In this case, adding IPs under "Management Service IPs" is not supported and you must add the IPs to the "Data

Service IPs" category in the GUI.

- You must provide either IPv4 or IPv6 IP addresses, you cannot give both.
- You must add individual IP addresses one by one without any prefix; adding a range of IP addresses is not supported.
- You can remove any previously defined IPs, but you will not be able to remove any IPs that are currently in use by one or more services.

Multi-Cluster Connectivity

You can establish connectivity between multiple Nexus Dashboard clusters for ease of access to all the clusters, as well as access to any of the sites and services running on any of the connected clusters.

When you add a second cluster, a group of clusters is formed. The cluster from which you create the group becomes the "primary" cluster with a number of unique characteristics that do not apply to other clusters in the group:

- You must use the primary cluster to connect all additional clusters.
- You must use the primary cluster to remove any of the clusters from the group.
- When upgrading Nexus Dashboard, you must upgrade the primary cluster before any other clusters in the group.

Establishing multi-cluster connectivity does not create any single databases with information from all clusters in the group. Every cluster continues to maintain its own configuration databases, while simultaneously being able to function as a proxy for all other clusters in the group regardless of which cluster an action or request is originated from or destined to.

Guidelines and Limitations

The following guidelines apply when configuring multi-cluster connectivity:

- This release supports multi-cluster connectivity between clusters deployed using physical or virtual (ESX) form factors only.

In other words, you can join physical Nexus Dashboard clusters with virtual (ESX) clusters, but virtual (KVM) or cloud clusters do not support this feature.

- For supported scale limits, such as number of clusters that can be connected together and number of sites across all clusters, see the [Nexus Dashboard Release Notes](#) for your release.
- Connectivity must be established between all nodes of all clusters, which will be connected via multi-cluster connectivity.
- The names of the sites onboarded in the clusters that you plan to connect together must be unique across those clusters.

Duplicate site names across different clusters may result in DNS resolution failures.

- The primary cluster, which you use to establish multi-cluster connectivity, must be running the same or a later release of Nexus Dashboard than any other cluster in the group.

In other words, you cannot connect a Nexus Dashboard cluster running release 2.3.1 from a primary cluster that is running release 3.0.1.

- If you are upgrading multiple clusters that are connected together, you must upgrade the primary cluster first.
- From any cluster in the connected clusters group, you can view other clusters only if they are running the same or earlier version of Nexus Dashboard.

In other words, if **cluster1** is running release 2.3.1 and **cluster2** is running release 2.2.1, you can view **cluster2** from **cluster1** but not vice versa.

- Multi-Cluster connectivity is supported for remote users only.

If you connect multiple clusters, but then login to one of the clusters as a local **admin** user, you will only be able to view and manage the local cluster into which you logged in.

To view and manage all clusters in the group, you must login as a remote user that is configured on all clusters.

- Nexus Dashboard Insights service in each cluster can view site groups from other Insights services across any cluster in the group.

However, when creating site groups, each Insights service can add sites which are onboarded in the same cluster where the service is installed only.

- Nexus Dashboard Orchestrator service supports managing only sites which are onboarded in the same cluster where the service is installed.

Connecting Multiple Clusters

Before you begin

- You must have familiarized yourself with the information provided in the [Guidelines and Limitations](#) section.
- You must have set up remote authentication and users on all clusters which you plan to connect.

Multi-Cluster connectivity is supported for remote users only, so you must configure the same remote user with **admin** privileges for all clusters. For additional details, see [\[Remote Authentication\]](#).

To connect another cluster:

1. Log in to the Nexus Dashboard GUI of the cluster which you want to designate as the primary.
2. Add second cluster.
 - a. From the main navigation menu, select **Admin > System Settings**.
 - b. In the main pane, select the **Multi-Cluster Connectivity** tab.
 - c. Click **Connect Cluster**.
3. Provide cluster information.
 - a. In the information fields, provide the hostname or IP address and the authentication information for the cluster you are adding.

You only need to provide the management IP address of one of the nodes in the target cluster. Other nodes' information will be automatically synced after connectivity is established.

- b. Then click **Save**.

The user you provide must have administrative rights on the cluster you are adding. The user credentials are used once when you are first establishing connectivity to the additional cluster. After initial connectivity is established, all subsequent communication is done through secure

keys. The secure keys are provisioned to each cluster while adding it to the group.

The cluster you are adding must not be part of an already existing group of clusters.

4. Repeat the procedure for any additional Nexus Dashboard cluster which you want to add to the group.

After multiple clusters are added to the group, you can see their status in the **Cluster Configuration > Multi-Cluster Connectivity** page.

Note that while you can view and manage any cluster from any other cluster as long as they are part of the same multi-cluster group, you can only add and remove clusters from the group when viewing the **primary** cluster.

The **Multi-Cluster Connectivity** page displays all clusters that are part of the multi-cluster group. The **Connect Cluster** button is shown only when viewing the primary cluster. To modify the cluster group, you need to navigate to the primary cluster as described in [\[Navigating Between Clusters\]](#), at which point the **Connect Cluster** button becomes available:

The screenshot shows the Nexus Dashboard interface. The top navigation bar includes the Cisco logo, 'Nexus Dashboard', and 'Admin Console'. The left sidebar has a 'Cluster' dropdown menu with 'nd-cluster-ova' selected, and a blue circle with the number 1 points to it. Below the sidebar are links for Overview, Operate, Analyze, and Admin. The main content area is titled 'System Settings' and has a 'Multi-Cluster Connectivity' tab selected. A 'Filter by attributes' input field and a 'Connect Cluster' button are at the top right. A table lists clusters with columns for Connectivity Status, Name, and URL. The first cluster is 'cluster-dev164' with status 'Up' and URL '10.195.255.164'. The second cluster is 'nd-c' with status 'Up' and three URLs: '172.23.141.129', '172.23.141.130', and '172.23.141.131'. The 'nd-c' cluster is labeled 'Primary' and 'Local'. A blue circle with the number 2 points to the 'Primary' label, and a blue circle with the number 3 points to the 'Local' label. At the bottom, there is a '10 Rows' selector and a pagination bar showing 'Page 1 of 1'.

Figure 2. Primary vs Non-primary Clusters

1. The **Cluster: <name>** dropdown in the main navigation menu shows the cluster you are currently viewing.

You can select a different cluster from this dropdown, which opens a new window allowing you to navigate to another cluster in the same group.



While the 2.x releases of Nexus Dashboard allowed you to view and manage any cluster from any other cluster as long as they were part of the same multi-cluster group, release 3.0.1 changed this behavior. You can now easily navigate between clusters by picking a specific cluster from the **Cluster** dropdown in the main navigation pane, but you cannot manage or configure another cluster directly from the one where you are logged in.

2. The **Primary** label indicates the group's primary cluster.

You must be viewing this cluster to make any changes to the cluster group, such as adding or removing clusters.

3. The **Local** label indicates the cluster you logged into.

This is the cluster whose address is displayed in the browser's URL field. If you navigate to a different cluster as mentioned above, the browser URL and the **Local** label will not change.

Disconnecting Clusters

To disconnect a cluster from an existing group:

1. Log in to the Nexus Dashboard GUI of the primary cluster.

Adding and removing clusters from the group must be done from the primary cluster.

2. From the main navigation menu, select **Admin > System Settings**.
3. In the main pane, select the **Multi-Cluster Connectivity** tab.
4. From the **Actions (...)** menu for the cluster you want to remove, select **Disconnect Cluster**
5. In the confirmation window, click **Ok**.

Deploying Additional Physical Nodes

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following sections describe how to deploy an additional physical node so you can add it as a **worker** or **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (such as physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster based on its role:

- For more information about **worker** nodes, see [Managing Worker Nodes](#).
- For more information about **standby** nodes, see [Managing Standby Nodes](#).

Prerequisites and Guidelines for Physical Nodes

- Ensure that you have reviewed and completed the general prerequisites described in the [\[Platform Overview\]](#), especially the network and fabric connectivity sections.
- Ensure that you have Reviewed and complete any additional prerequisites described in the *Release Notes* for the services you have deployed.

Some services may have additional caveats for **worker** and **standby** nodes. You can find the service-specific documents at the following links:

- [Nexus Dashboard Fabric Controller Release Notes](#)
- [Nexus Dashboard Insights Release Notes](#)
- [Nexus Dashboard Orchestrator Release Notes](#)
- For maximum number of **worker** and **standby** nodes in a single cluster, see the [Nexus Dashboard Release Notes](#) for your release.
- Ensure you are using the supported hardware and the servers are racked and connected.

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

Table 1. Supported UCS-C220-M5 Hardware

PID	Hardware
SE-NODE-G2	<ul style="list-style-type: none"> - UCS C220 M5 Chassis - 2x 10 core 2.2G Intel Xeon Silver CPU - 256 GB of RAM - 4x 25G Virtual Interface Card 1455 - 4x 2.4TB HDDs - 400GB SSD - 1.2TB NVMe drive - 1050W power supply

Table 2. Supported UCS-C225-M6 Hardware

PID	Hardware
ND-NODE-G4	<ul style="list-style-type: none"> - UCS C225 M6 Chassis - 2.8GHz AMD CPU - 256 GB of RAM - 4x 2.4TB HDDs - 960GB SSD - 1.6TB NVME drive - Intel X710T2LG 2x10 GbE (Copper) - Intel E810XXVDA2 2x25/10 GbE (Fiber Optic) - 1050W power supply



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

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

The minimum supported and recommended versions of CIMC are listed in the "Compatibility" section of the [Release Notes](#) for your Nexus Dashboard release.

- Ensure the hardware is running the same Nexus Dashboard release as your existing cluster.

If the new node is running an earlier release, you must manually upgrade to the current release, as described in [\[Manual Upgrades\]](#).

If for any reason you are unable to run the manual upgrade, you can reinstall the software, as described in [\[Re-Imaging Nodes\]](#).

Deploying Physical Nodes

Once you have completed all prerequisites described above, simply connect the node and power it own.

Once the node is deployed, you can add it to the cluster:

- To add the node as a **worker** node, see [Managing Worker Nodes](#).
- To add the node as a **standby** nodes, see [Managing Standby Nodes](#).

Deploying Additional Virtual Nodes in VMware ESX

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following sections describe how to deploy an additional node in VMware ESX so you can add it as a **worker** or **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster based on its role:

- For more information about **worker** nodes, see [Managing Worker Nodes](#).
- For more information about **standby** nodes, see [Managing Standby Nodes](#).

Prerequisites and Guidelines for ESX Nodes

- Ensure that you reviewed and completed the general prerequisites described in the [\[Platform Overview\]](#), especially the network and fabric connectivity sections.
- When deploying in VMware ESX, you can choose to deploy using a vCenter or directly in the ESXi host.

For detailed information, see one of the following sections.

- When deploying in VMware ESX, you can deploy two types of nodes:
 - Data node—node profile designed for data-intensive applications, such Nexus Dashboard Insights
 - App node—node profile designed for non-data-intensive applications, such Nexus Dashboard Orchestrator

Table 3. Supported Hardware

Nexus Dashboard Version	Data Node Requirements	App Node Requirements
Release 3.0.1	<p>VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3</p> <p>VMware vCenter 7.0.1, 7.0.2, 7.0.3 if deploying using vCenter</p> <p>Each VM requires the following:</p> <ul style="list-style-type: none"> 32 vCPUs with physical reservation of at least 2.2GHz 128GB of RAM with physical reservation 3TB SSD storage for the data volume and an additional 50GB for the system volume <p>Data nodes must be deployed on storage with the following minimum performance requirements:</p> <ul style="list-style-type: none"> The SSD must be attached to the data store directly or in JBOD mode if using a RAID Host Bus Adapter (HBA) The SSDs must be optimized for Mixed Use/Application (not Read-Optimized) 4K Random Read IOPS: 93000 4K Random Write IOPS: 31000 <p>We recommend that each Nexus Dashboard node is deployed in a different ESXi server.</p>	<p>VMware ESXi 7.0, 7.0.1, 7.0.2, 7.0.3</p> <p>VMware vCenter 7.0.1, 7.0.2 if deploying using vCenter</p> <p>Each VM requires the following:</p> <ul style="list-style-type: none"> 16 vCPUs with physical reservation of at least 2.2GHz 64GB of RAM with physical reservation 500GB HDD or SSD storage for the data volume and an additional 50GB for the system volume <p>Some services require App nodes to be deployed on faster SSD storage while other services support HDD. Check the Nexus Dashboard Capacity Planning tool to ensure that you use the correct type of storage.</p> <p>We recommend that each Nexus Dashboard node is deployed in a different ESXi server.</p>

Deploying ESX Node Using vCenter

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for ESX Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in VMware ESX using vCenter.

1. Obtain the Cisco Nexus Dashboard OVA image.

- a. Browse to the Software Download page.

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

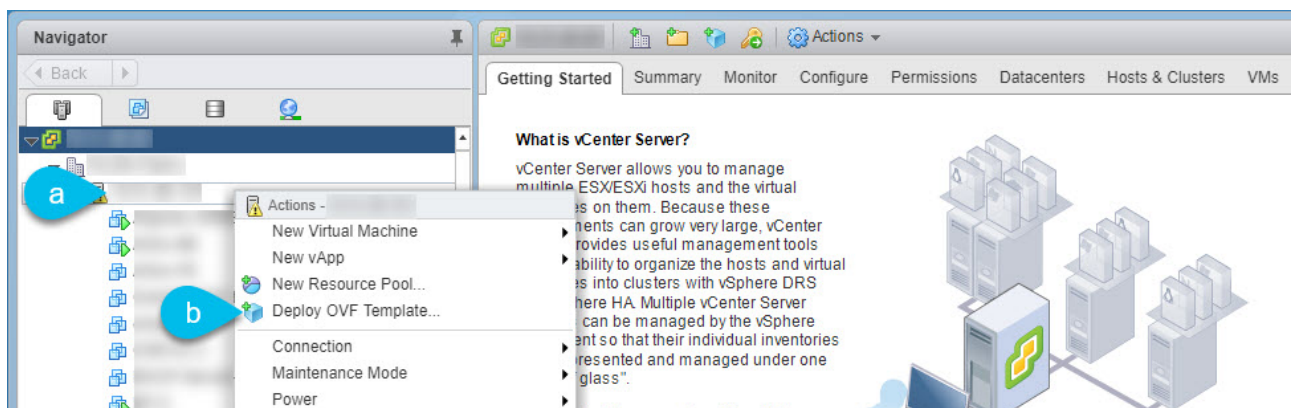
- b. Choose the Nexus Dashboard version you want to download.

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

2. Log in to your VMware vCenter.

Depending on the version of your vSphere client, the location and order of configuration screens may differ slightly.

3. Start the new VM deployment.



- a. Right-click the ESX host where you want to deploy.

- b. Then select "Deploy OVF Template..." .

The **Deploy OVF Template** wizard appears.

4. In the **Select an OVF template** screen, provide the OVA image, then click **Next**.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

☒ URL

http:// /nd-dk9.2.2.0.83.ova

☐ Local file

Choose Files

No file chosen

CANCEL

BACK

NEXT

a. Provide the image.

If you hosted the image on a web server in your environment, select **URL** and provide the URL to the image.

If your image is local, select **Local file** and click **Choose Files** to select the OVA file you downloaded.

b. Click **Next** to continue.

5. In the **Select a name and folder** screen, provide a name and location for the VM.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select a name and folder

Specify a unique name and target location

Virtual machine

nd-cluster-vm1

Select a location for the virtual machine.

172.31.141.49

Datacenter1

CANCEL

BACK

NEXT

a. Provide the name for the virtual machine.

b. Select the location for the virtual machine.

c. Click **Next** to continue

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

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- 3 Select a compute resource**
- 4 Review details
- 5 Select storage
- 6 Ready to complete

Select a compute resource

Select the destination compute resource for this operation

▼ Datacenter1

> 172.23.136.84

> 172.23.136.86

> 172.23.136.87

> 172.23.136.88

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK **NEXT**

a. Select the vCenter datacenter and the ESX host for the virtual machine.

b. Click **Next** to continue

7. In the **Review details** screen, click **Next** to continue.

8. In the **Configuration** screen, select the node profile you want to deploy.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 Configuration**
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Configuration

Select a deployment configuration

<div><div><input checked="" type="radio"/> App</div><div><input type="radio"/> Data</div></div>	Description Use this deployment profile to configure an App OVA with 16vCPUs, 64GB RAM, and 500GB Disk.
---	---

2 Items

CANCEL

BACK **NEXT**

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

b. For more information about the node profiles, see [Prerequisites and Guidelines for ESX Nodes](#).

c. Click **Next** to continue

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

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
✓ 5 Configuration
6 Select storage
7 Select networks
8 Customize template
9 Ready to complete

Select storage
Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thick Provision Lazy Zeroed

VM Storage Policy: Datastore Default

Name	Capacity	Provisioned	Free	Type	Cluster
datastore1 (3)	925.25 GB	225.74 GB	707.7 GB	VMFS 5	

Compatibility
✓ Compatibility checks succeeded.

CANCEL BACK **NEXT**

- From the **Select virtual disk format** drop-down, select **Thick Provision Lazy Zeroed**.
- Select the datastore for the virtual machine.

We recommend a unique datastore for each node.

- Click **Next** to continue

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

There are two networks required by the Nexus Dashboard cluster:

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

For more information about these networks, see "Network Connectivity".

- In the **Customize template** screen, provide the required information.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Configuration
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values ✕

Resource Configuration	1 settings	
a 1. Data Disk Size (GB)	Data disk size (min 500GB, max 1536GB (1.5TB)) <div style="border-bottom: 1px solid #ccc; margin-top: 5px;">500</div>	
Node Configuration	3 settings	
b 1. Password	Local "rescue-user" password <div style="display: flex; justify-content: space-between; align-items: flex-start;"><div style="width: 80%;"><div style="border-bottom: 1px solid #ccc; margin-bottom: 5px;">Password</div><div style="border-bottom: 1px solid #ccc; margin-bottom: 5px;">Confirm Password</div></div><div style="width: 15%; text-align: center;"><div style="border-bottom: 1px solid #ccc; margin-bottom: 5px;">Password</div><div style="border-bottom: 1px solid #ccc; margin-bottom: 5px;">Confirm Password</div></div><div style="width: 5%; text-align: center;"><div>?</div><div>?</div></div></div>	
c 2. Management Network Address and subnet	Management network address. Enter IP/subnet <div style="border-bottom: 1px solid #ccc; margin-top: 5px;">172.31.140.46/24</div>	
d 3. Management Gateway IP	Management network gateway IP address. Enter IP only <div style="border-bottom: 1px solid #ccc; margin-top: 5px;">172.31.140.1</div>	

CANCEL

BACK **e** NEXT

- a. Provide the sizes for the node's data disks.

We recommend using the default values for the required data volume.

The default values will be pre-populated based on the type of node you are deploying, with App node having a single 500GB disk and Data node having a single 3TB disk.

Note that in addition to the data volume, a second 50GB system volume will also be configured but cannot be customized.

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

This password is used for the **rescue-user** account on each node. We recommend configuring the same password for all nodes, however you can choose to provide different passwords for the second and third node.

- c. Provide the **Management Network** IP address, netmask.

- d. Provide the **Management Network** IP gateway.

- e. Click **Next** to continue.

12. In the **Ready to complete** screen, verify that all information is accurate and click **Finish** to begin deploying the node.
13. Once the VM deployment is finished, power on the VM.
14. Add the node as **primary** or **standby**.

Once the node is deployed, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#).
- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#).

Deploying ESX Node Directly in ESXi

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for ESX Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in VMware ESX using vCenter.

1. Obtain the Cisco Nexus Dashboard OVA image.

- a. Browse to the Software Download page.

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

- b. Choose the Nexus Dashboard version you want to download.

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

2. Log in to your VMware ESXi.

Depending on the version of your ESXi server, the location and order of configuration screens may differ slightly.

3. Right-click the host and select Create/Register VM.

4. In the Select creation type screen, choose Deploy a virtual machine from an OVF or OVA file, then click Next.

5. In the Select OVF and VMDK files screen, provide the virtual machine name (for example, **nd-node-worker1**) and the OVA image you downloaded in the first step, then click Next.

6. In the Select storage screen, choose the datastore for the VM, then click Next.

7. In the Select OVF and VMDK files screen, provide the virtual machine name (for example, **nd-node-worker1**) and the OVA image you downloaded in the first step, then click Next.

8. In the Deployment options screen, choose Disk Provisioning: Thick, uncheck the Power on automatically option, then click Next to continue.

There are two networks, fabric0 is used for the data network and mgmt0 is used for the management network.

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

10. Wait for the VM to finish deploying, ensure that the VMware Tools periodic time synchronization is disabled, then start the VM.

To disable time synchronization:

- a. Right-click the node's VM and select Edit Settings.

- b. In the Edit Settings window, select the VM Options tab.

- c. Expand the VMware Tools category and uncheck the Synchronize guest time with host option.

11. Open the node's console and configure the node's basic information.

- a. Begin initial setup.

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

```
[ OK ] Started atomix-boot-setup.  
Starting Initial cloud-init job (pre-networking)...  
Starting logrotate...  
Starting logwatch...  
Starting keyhole...  
[ OK ] Started keyhole.  
[ OK ] Started logrotate.  
[ OK ] Started logwatch.  
Press any key to run first-boot setup on this console...
```

b. Enter and confirm the admin password

This password will be used for the rescue-user SSH login and for adding this node to the cluster.

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

c. Enter the management network information.

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

d. Review and confirm the entered information.

You will be asked if you want to change the entered information. If all the fields are correct, choose **n** to proceed.

If you want to change any of the entered information, enter **y** to re-start the basic configuration script.

```
Please review the config  
Management network:  
Gateway: 192.168.9.1  
IP Address/Mask: 192.168.9.172/24  
Re-enter config? (y/N): n
```

12. Add the node as **primary** or **standby**.

Once the node is deployed, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#).

- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#).

Deploying Additional Virtual Nodes in Linux KVM

Initial cluster deployment is described in [Nexus Nexus Dashboard Deployment Guide](#). The following sections describe how to deploy an additional node in Linux KVM so you can add it as a **standby** node.



When adding nodes to an existing cluster, the additional nodes must be of the same form factor (physical or virtual) as the rest of the nodes in the cluster. This release does not support clusters with nodes of different form factors.

After you deploy an additional node, you can add it to the cluster as a **standby** node as described in [Managing Standby Nodes](#).

Prerequisites and Guidelines for KVM Nodes

- Ensure that you reviewed and complete the general prerequisites described in the [\[Platform Overview\]](#), especially the network and fabric connectivity sections.
- Ensure that your VM has sufficient resources:

Table 4. Supported Hardware

Nexus Dashboard Version	VM Requirements
Release 2.2.x	<ul style="list-style-type: none"> Supported Linux distribution: <ul style="list-style-type: none"> For Nexus Dashboard Orchestrator, you must deploy in CentOS Linux For Nexus Dashboard Fabric Controller, you must deploy in CentOS or Red Hat Enterprise Linux Supported versions of Kernel and KVM: <ul style="list-style-type: none"> Kernel 3.10.0-957.el7.x86_64 or later KVM libvirt-4.5.0-23.el7_7.1.x86_64 or later 16 vCPUs 64 GB of RAM 500 GB disk <p>Each node requires a dedicated disk partition</p> <ul style="list-style-type: none"> The disk must have I/O latency of 20ms or less. You can verify the I/O latency using the following command: <pre>fio --rw=write --ioengine=sync --fdatasync=1 --directory=test -data_with_se --size=22m --bs=2300 --name=mytest And confirm that the 99.00th=[<value>] in the fsync/fdatasync/sync_file_range section is under 20ms.</pre> We recommend that each Nexus Dashboard node is deployed in a different KVM server.

Deploying KVM Nodes

Before you begin

Ensure that you meet the requirements and guidelines described in [Prerequisites and Guidelines for KVM Nodes](#).

This section describes how to deploy an additional Cisco Nexus Dashboard node in Linux KVM.

- Download the Cisco Nexus Dashboard image.
 - Browse to the **Software Download** page.

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

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

- c. Download the Cisco Nexus Dashboard image for Linux KVM (`nd-dk9.<version>.qcow2`).
2. Copy the image to the Linux KVM servers where you will host the nodes.

If you have already copied the image, for example when initially deploying the cluster, you can use the same base image and skip this step. The following steps assume you copied the image into the `/home/nd-base` directory.

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

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

3. Create the required disk images for the node.

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

- a. Log in to your KVM host as the `root` user.
- b. Create a directory for the node's snapshot.

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

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

- c. Create the snapshot.

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

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

The following steps assume you are adding `nd-node4`.

- d. Create the additional disk image for the node.

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

```
# qemu-img create -f qcow2 /home/nd-node1/nd-node4-disk2.qcow2 500G
```

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

- `/home/nd-node4/nd-node4-disk1.qcow2`, which is a snapshot of the base `qcow2` image you downloaded in Step 1.
- `/home/nd-node4/nd-node4-disk2.qcow2`, which is a new 500GB disk you created.

4. Create the first node's VM.

You can use CLI or KVM GUI to create the VM with the following configuration:

- o 16 vCPUs
- o 64GB of RAM
- o Operating system type set to **linux2020**
- o Network device model set to **virtio**
- o Management interface mapped to bus **0x00** and slot **0x03** and Data interface mapped to bus **0x00** and slot **0x04**



Nexus Dashboard expects the Management interface to be connected to bus **0x00** and slot **0x03** and the Data interface to bus **0x00** and slot **0x04**. If this is not the case, the cluster will not have network connectivity.

For example, to create the VM using CLI:

```
# virt-install --name <node-name> \  
  --vcpus 16 --ram 64000 --osinfo linux2020 \  
  --disk path=/home/nd-node4/nd-node4-disk1.qcow2 \  
  --disk path=/home/nd-node4/nd-node4-disk2.qcow2 \  
  --network bridge:br-  
oob,model=virtio,address.type=pci,address.domain=0,address.bus=0,address.slot=3 \  
  --network bridge:br-  
vnd,model=virtio,address.type=pci,address.domain=0,address.bus=0,address.slot=4 \  
  --noautoconsole --import
```

5. Open the node's console and configure the node's basic information.

- a. Press any key to begin initial setup.

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

```
[ OK ] Started atomix-boot-setup.  
Starting Initial cloud-init job (pre-networking)...  
Starting logrotate...  
Starting logwatch...  
Starting keyhole...  
[ OK ] Started keyhole.  
[ OK ] Started logrotate.  
[ OK ] Started logwatch.  
Press any key to run first-boot setup on this console...
```

- b. Enter and confirm the **admin** password.

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

Admin Password:
Reenter Admin Password:

- c. Enter the management network information.

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

- d. When asked if the node is the "Cluster Leader", choose "no".

Since you are adding a **worker** or **standby** node, do not designate it as the cluster leader

Is cluster leader?: n

- e. Review and confirm the entered information.

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

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

6. Wait for the initial bootstrap process to complete.

After you provide and confirm management network information, wait for the initial bootstrap process to complete:

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

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

7. Add the node to the cluster as **primary** or **standby**.

Once the bootstrap process is finished, you can add it to the cluster:

- o To add the node as a **worker** node, see [Managing Worker Nodes](#).
- o To add the node as a **standby** nodes, see [Managing Standby Nodes](#).

Managing Worker Nodes

You can add a number of worker nodes to an existing 3-node cluster for horizontal scaling to enable application co-hosting.

For additional information about application co-hosting and cluster sizing, see the [\[Platform Overview\]](#) section of this document.



Worker nodes are not supported for cloud form factors of Nexus Dashboard clusters deployed in AWS or Azure.

Adding Worker Nodes

This section describes how to add a worker node to your cluster to enable horizontal scaling

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node as described in [Deploying Additional Physical Nodes](#), [Deploying Additional Virtual Nodes in VMware ESX](#), [Deploying ESX Node Directly in ESXi](#), or [Deploying Additional Virtual Nodes in Linux KVM](#).
- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a worker node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.
 - a. Provide the **Name** of the node.
 - b. From the **Type** dropdown, select **Worker**.
 - c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

- d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

- e. Provide the **Data Network** information.

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

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

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

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

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

1. If you are running Nexus Dashboard Insights application, disable and re-enable the application.

After you add the new worker node, you must disable and re-enable the application for its services to be properly distributed to the new node.

Deleting a Worker node

Before you begin

- Ensure that the primary nodes and the cluster are healthy.

To delete an existing worker node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the worker node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Managing Standby Nodes

You can add up to two standby nodes, which you can use to quickly restore the cluster functionality in case one or more primary nodes fail by replacing the failed primary node with the standby node.

Standby nodes are similar to worker nodes in deployment, initial configuration, and upgrades. However, unlike worker nodes, the cluster will not use the standby nodes for any workloads.



Standby nodes are not supported for single-node clusters or clusters deployed in AWS or Azure.

The following two cases are supported:

- Single primary node failure

You can use the UI to convert the standby node into a new primary node.

- Two primary nodes failure

You will need to perform manual failover of one of the nodes to restore cluster functionality. Then fail over the second node using standard procedure.

Adding Standby Nodes

This section describes how to add a standby node to your cluster for easy cluster recover in case of a primary node failure.

Before you begin

- Ensure that the existing primary nodes and the cluster are healthy.
- Prepare and deploy the new node as described in [Deploying Additional Physical Nodes](#), [Deploying Additional Virtual Nodes in VMware ESX](#), [Deploying ESX Node Directly in ESXi](#), or [Deploying Additional Virtual Nodes in Linux KVM](#).

You can failover only between nodes of identical types (physical or virtual), so you must deploy the same type of node as the nodes in your cluster which you may need to replace. In case of virtual nodes deployed in VMware ESX, which have two node profiles (**OVA-app** and **OVA-data**), you can failover only between nodes of the same profile.

- Ensure that the node you are adding is powered on.
- If you are adding a physical node, ensure that you have the new node's CIMC IP address and login information.

You will need to use the CIMC information to add the new node using the Nexus Dashboard GUI.

- If you are adding a virtual node, ensure that you have the node's management IP address and login information.

To add a standby node:

1. Log in to the Cisco Nexus Dashboard GUI.

2. From the main navigation menu, select **System Resources > Nodes**.
3. In the main pane, click **Add Node**.

The **Add Node** screen opens.

4. In the **Add Node** screen, provide the node information.
 - a. Provide the **Name** of the node.
 - b. From the **Type** dropdown, select **Standby**.
 - c. Provide the **Credentials** information for the node, then click **Verify**.

For physical nodes, this is the IP address, username, and password of the server's CIMC. The CIMC will be used to configure the rest of the information on the node.

For virtual nodes, this is the IP address and **rescue-user** password you defined for the node when deploying it.

- d. Provide the **Management Network** information.

For virtual nodes, the management network information will be pre-populated with the information pulled from the node based on the IP address and credentials you provided in the previous sub-step.

For physical nodes, you must provide the management network IP address, netmask, and gateway now.

- e. Provide the **Data Network** information.

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

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

Starting with release 2.1.1, Nexus Dashboard supports dual stack IPv4/IPv6 for the management and data networks.

If you want to provide IPv6 information, you must do it when adding the node.

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

5. Click **Save** to add the node.

The configuration will be pushed to the node and the node will be added to the list in the GUI.

Replacing Single Primary Node with Standby Node

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

Before you begin

- Ensure that at least 2 primary nodes are healthy.

If two of the primary nodes are unavailable, you will need to manually restore the cluster as described in [Replacing Two Primary Nodes with Standby Nodes](#)

- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **primary** node you want to replace is powered off.



You cannot re-add the **primary** node you are replacing back to the cluster after the failover is complete. If the **primary** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-image it as described in [Troubleshooting](#) and add it as a **standby** or **primary** node only.

To failover a single primary node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Click the **Actions (...)** menu next to the **Inactive** primary node that you want to replace.
4. Choose **Failover**.

Note that you must have a standby node already configured and added or the **Failover** menu option will not be available.

5. In the **Fail Over** window that opens, select a standby node from the dropdown.
6. Click **Save** to complete the failover.

The failed primary node will be removed from the list and replaced by the standby node you selected. The status will remain **Inactive** while the services are being restored to the new primary node.

It can take up to 10 minutes for all services to be restored, at which point the new primary node's status will change to **Active**.

Replacing Two Primary Nodes with Standby Nodes

This section describes failover using a pre-configured **standby** node. If your cluster does not have a standby node, follow the steps described in one of the sections in [Troubleshooting](#) instead.

If only one of your primary nodes failed, you can use the GUI to replace it with a standby node as described in [Replacing Single Primary Node with Standby Node](#).

However, when two primary nodes are unavailable, the cluster goes offline. In this case, most operations including the UI are disabled and no configuration changes can be made to the cluster. You can still SSH into the remaining primary node as the **rescue-user**, which is used to recover the cluster by manually failing over one of the failed primary nodes to a standby node. Once two **primary** nodes are available again, the cluster can resume normal operation, at which point you can recover the second primary node using the normal procedure.

Before you begin

- Ensure that you have at least one **standby** node available in the cluster.

Setting up and configuring **standby** nodes is described in [Adding Standby Nodes](#).

- Ensure that the **primary** nodes you want to replace are powered off.



You cannot re-add the **primary** node you are replacing back to the cluster after the failover is complete. If the **primary** node you replace is still functional and you want to re-add it to the cluster after the failover, you must factory reset or re-image it as described in [Troubleshooting](#) and add it as a **standby** or **primary** node only.

- If you had installed the Nexus Dashboard Fabric Controller (NDFC) service in the cluster, you must have a configuration backup available to restore after you recover the cluster.

The Fabric Controller service cannot recover from a two **primary** node failure of the Nexus Dashboard cluster where it is running. After your recover the cluster, you must re-install the NDFC service and restore its configuration from a backup.

To fail over two primary nodes:

1. Log in to the remaining primary node via CLI as **rescue-user**.
2. Execute the failover command.

In the following command, replace **<node1-data-ip>** and **<node2-data-ip>** with the data network IP addresses of the failed nodes:

```
# acs failover --failedIP <node1-data-ip> --failedIP <node2-data-ip>
```



Even though only the first node is failed over, the second failed node you provide is required internally to recover the cluster.

By default, the healthy primary node will automatically pick an available standby node and fail over the first failed node you provide (**<node1-data-ip>**) to it.

If you would like to provide a specific standby node, you can add **<standby-node-data-ip>** to the above command:

```
# acs failover --failedIP <node1-data-ip> --failedIP <node2-data-ip> \  
--standbyIP <standby-node1-data-ip>
```

3. Confirm that you want to proceed.

Warning: Failover can be a disruptive operation and should only be performed as last resort option to recover cluster from disasters using standby where two primary nodes have lost their state due to hardware faults.

Proceed? (y/n): y

The primary node will copy the configuration state to the standby node and both nodes will restart. It may take up to 30 minutes for the nodes to come up and the cluster to be restored. You can check the progress by navigating to the primary node's UI.

4. Wait for the failover to complete.

The healthy primary node will copy the configuration state to the standby node and both nodes will restart. It may take up to 30 minutes for the nodes to come up and the cluster to become functional. You can check the progress by navigating to the master node's UI:

5. If necessary, re-install the NDFC service and restore NDFC configuration.

We recommend fully recovering the cluster by replacing the 3rd **primary** node before any configuration changes. However, if you must recover your production NDFC configuration as soon as possible, you can do so now:

- a. Using a browser, log in to one of the two active **primary** nodes of the ND cluster.
- b. Disable the NDFC service.

This is described in [\[Disabling Services\]](#).

- c. Delete the NDFC service.

This is described in [\[Uninstalling Services\]](#).

- d. Re-install the NDFC service and enable it.

This is described in [\[Services Management\]](#).

This step assumes all installation pre-requisites are completed from the initial service deployment. For detailed information on all NDFC requirements, see the [Nexus Dashboard Fabric Controller Installation Guide](#) for your release.

- e. Restore NDFC configuration from a backup.

This is described in the "*Operations > Backup and Restore*" chapter of the [NDFC Fabric Controller Configuration Guide](#) for your release.

- f. Validate that NDFC service is up and running before moving to the next step.

6. After the cluster is back up, fail over the second failed primary node.

At this point, you can use the standard procedure described in [Replacing Single Primary Node with Standby Node](#). If you do not have a second **standby** node, you can add it to the cluster while it has only 2 **primary** nodes, as described in [Adding Standby Nodes](#).

If you want re-add the same 2 **primary** nodes that had failed, you must factory reset or re-image them as described in [Troubleshooting](#) and only then add them as **standby** or **primary** nodes using the following steps:

- a. Ensure that the failed primary nodes are disconnected before they can be re-added to the cluster.

Connectivity must be disabled on both ND management and data interfaces. For virtual ND

deployments, the VM vNICs can be disabled/disconnected. For physical ND deployments, the interfaces connected to the ND management and data networks can be shut down.

- b. In Nexus Dashboard UI, navigate to system resources and nodes page and note the listed inactive master node.

The page displays the failed primary nodes as “inactive”. Note down one of the nodes as that’s the first node you will re-add back to cluster.

- c. From the console of the node you noted in the previous substep, run **acs reboot factory-reset** command.
- d. Wait for node’s console to display " Press any key to run first-boot setup on this console" .
- e. Ensure that network connectivity for data and management interfaces of the node is restored.

For example, in case of a virtualized ND deployment, you may need to enable the respective vnics if they were disabled/disconnected.

- f. In the Nexus Dashboard GUI of the functional 2-node cluster that you restored in the previous steps, choose “Register” for the inactive node.

It may take several minutes for the node to fully join the cluster and for it to become a fully functional 3-node cluster.

- g. After the node is registered, run the post-recovery command.

Log in to any of the primary nodes as the **rescue-user** and run the following command:

```
curl -k `curl -k https://dcnm-fm.cisco-ndfc.svc:9443/fm/internal/resetobjectstore
```

The process takes up to 10 minutes to complete and you will see the following message when it is done:

Minio is reset, Please disable and enable NDFC to make it effective.

- h. Finally, disable and re-enable the NDFC service from the Nexus Dashboard’s **Admin Console** to ensure the NDFC service is fully operational.



If any image management policies were created after Step 5, images uploaded for the policies must be re-uploaded from the NDFC’s **Image Management** UI. Policies will reconcile with the newly uploaded images and do not need to be re-created.

7. After the 3-node cluster is fully operational, re-add the last node to the cluster.

After the cluster has been restored to a full 3-node cluster as described in the previous steps, add the last node as a **standby** node back to the cluster.

If you want to re-add the old primary node back to the cluster, simply factory reset it and add it as a **standby** node as described in [Adding Standby Nodes](#).

Deleting Standby Nodes

Before you begin

- Ensure that the primary nodes and the cluster are healthy.

To delete an existing standby node:

1. Log in to the Cisco Nexus Dashboard GUI.
2. From the main navigation menu, select **System Resources > Nodes**.
3. Select the checkbox next to the standby node you want to delete.
4. From the **Actions** menu, choose **Delete** to delete the node.

Trademarks

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

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

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2017–2023 Cisco Systems, Inc. All rights reserved.

First Published: 2023-01-31

Last Modified: 2023-04-11

Americas Headquarters

Cisco Systems, Inc.

170 West Tasman Drive

San Jose, CA 95134-1706

USA

<http://www.cisco.com>

Tel: 408 526-4000

800 553-NETS (6387)

Fax: 408 527-0883