

# Cisco DNA Center High Availability Guide, Release 2.3.3

---

**First Published:** 2022-04-26

**Last Modified:** 2023-03-03

## Cisco DNA Center High Availability Guide, Release 2.3.3

Refer to the following guide for a description of Cisco DNA Center's high availability (HA) implementation.



---

**Note** For a description of disaster recovery functionality in Cisco DNA Center, see the "Implement Disaster Recovery" chapter in the [Cisco DNA Center Administrator Guide](#).

---

## High Availability Overview

Cisco DNA Center's HA framework is designed to reduce the amount of downtime that results from failures and make your network more resilient when they take place. It does so by providing the near real-time synchronization of changes across your cluster nodes, giving your network a level of redundancy to deal with any issues that arise. The supported synchronization types include:

- Database changes, such as updates related to configuration, performance and monitoring data.
- File changes, such as report configurations, configuration templates, TFTP-root directory, administration settings, licensing files, and the key store.

This guide covers the requirements that need to be met to use HA, deployment and administration best practices, and the failure scenarios you may encounter (as well as how Cisco DNA Center deals with them and any required user action).



---

**Important** Cisco DNA Center provides HA support for both Automation and Assurance functionality.

---

## High Availability Requirements

To enable HA in your production environment, the following requirements must be met:

- Your cluster consists of three Cisco DNA Center appliances with the same number of cores (three 56-core appliances, for example).
- Your secondary appliances are running the same version of Cisco DNA Center (1.2.8 or later) as the primary appliance.

- Multinode cluster deployments require all of the member nodes to be in the same network and at the same site. The Cisco DNA Center appliance does not support the distribution of nodes across multiple networks or sites.
- Your cluster's Round-Trip Time (RTT) is 10 milliseconds or less.

## High Availability Functionality

Cisco DNA Center supports a three-node cluster configuration, which provides *both* software and hardware high availability. A software failure occurs when a service on a node fails. Software high availability involves the ability of the services on the node or nodes to be restarted. For example, if a service fails on one node in a three-node cluster, that service is either restarted on the same node or on one of the other two remaining nodes. A hardware failure occurs when the appliance itself malfunctions or fails. Hardware high availability is enabled by the presence of multiple appliances in a cluster, multiple disk drives within each appliance's RAID configuration, and multiple power supplies. As a result, a failure by one of these components can be tolerated until the faulty component is restored or replaced.



---

**Note** Cisco DNA Center does not support a cluster with more than three nodes. For example, a multinode cluster with five or seven nodes is not currently supported.

Fault tolerance for a three-node cluster is designed to handle single-node failure. In other words, Cisco DNA Center tries to provide high availability across specific services even if a single node fails. If two nodes fail, the quorum necessary to perform HA operations is lost and the cluster breaks.

---

## Clustering and Database Replication

Cisco DNA Center provides a mechanism for distributed processing and database replication among multiple nodes. Clustering provides both sharing of resources and features, as well as enabling high availability.

## Security Replication

In a multinode environment, the security features of a single node are replicated to the other two nodes, including any X.509 certificates or trustpools. After you join nodes to an existing cluster to form a three-node cluster, the Cisco DNA Center GUI user credentials are shared across the nodes. However, the CLI user credentials are not shared, because they are separate for each node.

## Software Upgrade

In a multinode cluster, you can trigger an upgrade of the whole cluster from the Cisco DNA Center GUI (the GUI represents the entire cluster and not just a single node). An upgrade triggered from the GUI automatically upgrades all the nodes in the cluster.



**Note** After you initiate a system upgrade (which updates Cisco DNA Center's core infrastructure), Cisco DNA Center goes into maintenance mode. In maintenance mode, Cisco DNA Center will be unavailable until the upgrade process completes. You should take this into account when scheduling a system upgrade. Once the system upgrade does complete, you can verify its success in the GUI by accessing **System > Software Updates > Updates** and checking the installed version.

1. Click the menu icon (☰) and choose **System > Software Updates > Updates**.
2. In the **System Update** area, confirm that the latest system package has been installed.

## High Availability Deployment

The topics in this section cover the best practices you should follow when deploying and administering an HA-enabled cluster in your production environment.

### Deployment Recommendations

Cisco DNA Center supports three-node clusters. The odd number of nodes provides the quorum necessary to perform any operation in a distributed system such as this. Instead of three separate nodes, Cisco DNA Center views them as one logical entity accessed via a virtual IP address.

When deploying HA, we recommend the following:

- When setting up a three-node cluster, do not configure the nodes to span a LAN across slow links, as this can make the cluster susceptible to network failures. It can also increase the amount of time needed for a service that fails on one of the nodes to recover. When configuring a three-node cluster's cluster interface, also ensure that all of the cluster nodes reside in the same subnet.
- Avoid overloading a single interface with management, data, and HA responsibilities, as this might negatively impact HA operation.
- In the appliance configuration wizards, Cisco DNA Center prepopulates the **Services Subnet** and **Cluster Services Subnet** fields with link-local (169.x.x.x) subnets. We recommend that you use the default subnets, but you can choose to specify different subnets. If you do so, they must conform with the IETF RFC 1918 and 6598 specifications for private networks, which support the following address ranges:
  - 10.0.0.0/8
  - 172.16.0.0/12
  - 192.168.0.0/16
  - 100.64.0.0/10

For details, see RFC 1918, [Address Allocation for Private Internets](#), and RFC 6598, [IANA-Reserved IPv4 Prefix for Shared Address Space](#).


- Enable HA during off-hours, because Cisco DNA Center enters maintenance mode and is unavailable until it finishes redistributing services.

## Deploy a Cluster

To deploy Cisco DNA Center on a three-node cluster with HA enabled, complete the following procedure:

### Procedure

---

- Step 1** Configure Cisco DNA Center on the first node in your cluster:
- If you are configuring a first-generation appliance, see the "Configure the Primary Node" topic in the [Cisco DNA Center First-Generation Appliance Installation Guide](#).
  - If you are configuring a second-generation appliance, see the topic that is specific to the configuration wizard you want to use and your appliance type in the [Cisco DNA Center Second-Generation Appliance Installation Guide](#):
    - If you are configuring a second-generation appliance using the Maglev Configuration wizard, see the "Configure the Primary Node Using the Maglev Wizard" topic.
    - If you are configuring a 44- or 56-core appliance using the browser-based configuration wizard, see the "Configure the Primary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 44/56-Core Node Using the Browser-Based Wizard" chapter.
    - If you are configuring a 112-core appliance using the browser-based configuration wizard, see the "Configure the Primary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 112-Core Node Using the Browser-Based Wizard" chapter.
- Step 2** Configure Cisco DNA Center on the second node in your cluster:
- If you are configuring a first-generation appliance, see the "Configure a Secondary Node" topic in the [Cisco DNA Center First-Generation Appliance Installation Guide](#).
  - If you are configuring a second-generation appliance, see the topic that is specific to the configuration wizard you want to use and your appliance type in the [Cisco DNA Center Second-Generation Appliance Installation Guide](#):
    - If you are configuring a second-generation appliance using the Maglev Configuration wizard, see the "Configure a Secondary Node Using the Maglev Wizard" topic.
    - If you are configuring a 44- or 56-core appliance using the browser-based configuration wizard, see the "Configure a Secondary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 44/56-Core Node Using the Browser-Based Wizard" chapter.
    - If you are configuring a 112-core appliance using the browser-based configuration wizard, see the "Configure a Secondary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 112-Core Node Using the Browser-Based Wizard" chapter.
- Step 3** Configure Cisco DNA Center on the third node in your cluster.
- Refer to the same secondary appliance configuration topic you viewed while completing the preceding step.
- Step 4** Activate high availability on your cluster:
- a) Click the menu icon () and choose **System > Settings > System Configuration > High Availability**.
  - b) Click **Activate High Availability**.

**Note**

- After you click **Activate High Availability** in the GUI, Cisco DNA Center enters into maintenance mode. In this mode, Cisco DNA Center is unavailable until the process completes, which can take several hours. You should take this into account when scheduling an HA deployment.
- Cisco DNA Center also goes into maintenance mode when you restore the database and perform a system upgrade (not a package upgrade).
- To enable external authentication with a AAA server in a three-node cluster environment, you must configure all individual Cisco DNA Center node IP addresses and the virtual IP address for the three-node cluster on the AAA server.

## Administer a Cluster

The topics in this section cover the administrative tasks you will need to complete when HA is enabled in your production environment.

### Run maglev Commands

To make any changes to the IP address, static route, DNS server, or **maglev** user password that are currently configured for a Cisco DNA Center appliance, you'll need to run the `sudo maglev-config update` CLI command.


### Typical Cluster Node Operations

The following operations are the ones you will typically need to complete for the nodes in your cluster, such as shutting down a cluster node (which you would do before performing planned maintenance), preparing a node for Return Merchandise Authorization (RMA), or rebooting (which you would do to restore a node that has been down or save configuration changes).



**Note** You cannot simultaneously reboot or shut down two nodes in an operational three-node cluster, as this breaks the cluster's quorum requirement.

Operation	Required Actions
From the CLI, shut down all of the nodes in a three-node cluster.	Run the <b>sudo shutdown -h now</b> command on all of the nodes at the same time.  When powering nodes back on, be sure to power on all nodes at the same time through Cisco IMC.
Reboot one or more nodes after making any change that may require a reboot.	Run the <b>sudo shutdown -r now</b> command on the relevant nodes.

Operation	Required Actions
Shut down or disconnect one node for maintenance (in situations where you are not just rebooting the node).	<p>Run the following commands:</p> <ol style="list-style-type: none"> <li>1. <b>maglev node drain</b> <i>node's-IP-address</i></li> <li>2. <b>maglev node drain_history</b> (to confirm that the node drained successfully)</li> <li>3. <b>sudo shutdown -h now</b> (run on the node you are shutting down)</li> </ol> <p>After performing maintenance on the node, complete the following steps:</p> <ol style="list-style-type: none"> <li>1. Log in to the Cisco IMC GUI as the Cisco IMC user.</li> <li>2. From the hyperlinked menu, choose <b>Host Power &gt; Power On</b> to power on the node. It should take 30–45 minutes for the node to come back up.</li> <li>3. Run the <b>magctl node display</b> command and wait for the node's status to display as <i>Ready</i>.</li> <li>4. Run the <b>maglev node allow</b> <i>node's-IP-address</i> command.</li> <li>5. Run the <b>magctl workflow status</b> command and wait until its output indicates that the task you initiated in the previous step completed successfully before you proceed.</li> <li>6. Run the <b>maglev service nodescale refresh</b> command, which puts the node in maintenance mode.</li> </ol> <p><b>Note</b> Instead of running the command, you can also do the following:</p> <ol style="list-style-type: none"> <li>a. From the Cisco DNA Center GUI, click the menu icon (  ) and choose <b>System &gt; Settings &gt; System Configuration &gt; High Availability</b>.</li> <li>b. Click <b>Activate High Availability</b>.</li> </ol>

Operation	Required Actions
Prepare a node for RMA.	<ol style="list-style-type: none"> <li>1. Drain the node: <b>maglev node drain</b> <i>node-IP-address</i> To confirm that the node drained successfully, run the <b>maglev node drain_history</b> command.</li> <li>2. Shut down the node: <b>sudo shutdown -h now</b></li> <li>3. Confirm that the node's status is listed as <i>NotReady, SchedulingDisabled</i>: <b>magctl node display</b></li> <li>4. Remove the node from the cluster: <b>maglev node remove</b> <i>node-IP-address</i></li> <li>5. Install the same Cisco DNA Center version that's already installed on the cluster's other two nodes.</li> <li>6. Add the node back to the cluster by configuring it as a secondary node (see the <a href="#">Cisco DNA Center Second-Generation Appliance Installation Guide</a>).</li> <li>7. Enable service distribution, which puts the node in maintenance mode: <b>maglev service nodescale refresh</b></li> </ol> <p><b>Note</b> Instead of running the command, you can also do the following:</p> <ol style="list-style-type: none"> <li>a. In the Cisco DNA Center GUI, click the <b>Menu</b> icon (☰) and choose <b>System &gt; Settings &gt; System Configuration &gt; High Availability</b>.</li> <li>b. Click <b>Activate High Availability</b>.</li> </ol>

## Replace a Failed Node

If a node fails, complete the following tasks in order to replace it:

1. Remove the failed node from your cluster.  
See [Remove the Failed Node, on page 7](#).
2. Replace the failed node with another node.  
See [Add a Replacement Node, on page 9](#).

## Remove the Failed Node

If a node fails, you'll need to do two things:

- Drain the node in order to remove its production workload.
- Remove the node from the cluster.




---

**Note** If you have any questions or concerns about removing your cluster's faulty node, contact the Cisco TAC for assistance.

---




---

**Note** A two-node cluster (a transient configuration that's not supported for normal use) results when one of the following situations occur:

- During the initial formation of a three-node cluster, only two of the cluster nodes are available.
- In an existing three-node cluster, one of the nodes has failed or is currently down.

While a two-node cluster is active, you will not be able to remove either of its nodes.

---

### Before you begin

Make sure that you have a backup of your data. If you are performing this procedure due to a node failure, you cannot create a backup now. Instead, you must rely on backups that you have been routinely creating.

### Procedure

---

**Step 1** Drain the faulty node:

- a) Drain the node: **maglev node drain** *<node's-Cluster-interface-IP-address>*
- b) Monitor the progress of the node drain: **maglev node drain\_progress**

**Step 2** Run the following commands:

- **sudo shutdown -h now**—Shuts down the node that you want to remove. The process takes about 10 minutes.
- **magctl node display**—Provides confirmation that the node is down. The node status should be `NotReady`.

**Warning** This step is irreversible. Only complete it if absolutely necessary.

**Step 3** Remove the faulty node from your cluster:

- a) After the node drain completes, remove the node: **maglev node remove** *<node's-Cluster-interface-IP-address>*
- b) Monitor the progress of the node's removal: **maglev node remove\_progress**
- c) Check that all services are running on the remaining two nodes:
  - **magctl node display**
  - **magctl appstack status**
- d) After the node has been successfully removed, reconfigure the node and add it back to the cluster.



See [Add a Replacement Node, on page 9](#) for more information.

---

## Add a Replacement Node

After removing the failed node, you can add a replacement node to the cluster.

### Before you begin

Make sure that you complete the following tasks:

- Remove the failed node. For information, see [Remove the Failed Node, on page 7](#).
- Allocate at least 30 minutes to perform this procedure.

### Procedure

- 
- Step 1** On the replacement node, install the same software version that the other nodes in the cluster are running.
- If you are configuring a first-generation appliance, use the Maglev Configuration wizard's **Join a Cisco DNA Center Cluster** option. See the "Configure a Secondary Node" topic in the [Cisco DNA Center First-Generation Appliance Installation Guide](#).
  - If you are configuring a second-generation appliance using the Maglev Configuration wizard, use the wizard's **Join a Cisco DNA Center Cluster** option. See the "Configure a Secondary Node Using the Maglev Wizard" topic in the [Cisco DNA Center Second-Generation Appliance Installation Guide](#).
  - If you are configuring a second-generation appliance using the browser-based configuration wizard, use the wizard's **Join an existing cluster** option. See one of the following topics in the [Cisco DNA Center Second-Generation Appliance Installation Guide](#):
    - 44 or 56-core appliance: See the "Configure a Secondary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 44/56-Core Appliance Using the Browser-Based Wizard" chapter.
    - 112-core appliance: See the "Configure a Secondary Node Using the Advanced Install Configuration Wizard" topic in the "Configure the 112-Core Appliance Using the Browser-Based Wizard" chapter.

**Important** In the Maglev Cluster Details screen (Maglev Configuration wizard) or Primary Cluster Details screen (Advanced Install configuration wizard), enter the IP address that's configured for the Cluster port on either of the nodes that are still active.

- Step 2** After the installation is complete, enter the following command:

```
magctl node display
```

The replacement node should show the `Ready` status.

- Step 3** Redistribute services to the replacement node by activating high availability on your cluster:

- a. Click the menu icon (☰) and choose **System > Settings > System Configuration > High Availability**.
- b. Click **Activate High Availability**.

**Step 4** Verify that services have been redistributed:

**magctl appstack status**

The replacement node should show a `Running` status.

**Step 5** If you previously backed up Assurance data, restore it.

For information, see the "Restore Data from Backups" topic in the [Cisco Digital Network Architecture Center Administrator Guide](#).

**Important** After you add the failed node back to your cluster, it serves as an add-on node. The node does not resume its previous role as the primary node.

---

## Minimize Failure and Outage Impact

In a typical three-node Cisco DNA Center cluster, each node is connected to a single cluster switch via the node's cluster port interface. Connectivity with the cluster switch requires two transceivers and a fiber optic cable, any of which can fail. The cluster switch itself can also fail (due to things like a loss of power or manual restart), which can result in an outage of your Cisco DNA Center cluster and loss of all controller functionality. To minimize the impact of a failure or outage on your cluster, do one or more of the following:

- Perform management operations such as software upgrades, configuration reloads, and power cycling during non-critical time periods, as these operations can result in a cluster outage.
- Connect your cluster nodes to a switch that supports the in-service software upgrade (ISSU) feature. This feature allows you to upgrade system software while the system continues to forward traffic, using nonstop forwarding (NSF) with stateful switchover (SSO) to perform software upgrades with no system downtime.
- Connect your cluster nodes to a switch stack, which allows you to connect each cluster node to a different member of the switch stack joined via Cisco StackWise. As the cluster is connected to multiple switches, the impact of one switch going down is mitigated.

## High Availability Failure Scenarios

Nodes can fail due to issues in one or more of the following areas:

- Software
- Network access
- Hardware

When a failure occurs, Cisco DNA Center normally detects it within 5 minutes and resolves the failure on its own. Failures that persist for longer than 5 minutes might require user intervention.

The following table describes failure scenarios your cluster might encounter and how Cisco DNA Center responds to them. Pay attention to the table's first column, which indicates the scenarios that require action from you in order to restore the operation of your cluster.



**Important** For a cluster to operate, Cisco DNA Center's HA implementation requires at least two cluster nodes to be up at any given time.

Requires User Action	Failure Scenario	HA Behavior
Yes	Any node in the cluster goes down.	Perform an Automation backup immediately. See the "Backup and Restore" chapter in the <i>Cisco Digital Network Architecture Center Administrator Guide</i> .
No	A node fails, is unreachable, or experiences a service failure for less than 5 minutes.	<ul style="list-style-type: none"> <li>• The UI is not accessible for 5 minutes after a node fails.</li> <li>• Services that were running on the failed node are not migrated to other nodes.</li> <li>• The northbound interface (NBI) remains usable on the remaining two nodes when using the VIP.</li> <li>• VIP connectivity will be restored after failover, and API calls recover after services are up and running.</li> </ul> <p>After the node is restored:</p> <ul style="list-style-type: none"> <li>• Data on the restored node is synched with other cluster members.</li> </ul> <p><b>Note</b> Historical Assurance data is restored, but data that was modified or updated during the failover process is not.</p> <ul style="list-style-type: none"> <li>• Pending UI and NBI calls that have not timed out complete.</li> </ul>

Requires User Action	Failure Scenario	HA Behavior
No	A node fails, is unreachable, or experiences a service failure for longer than 5 minutes.	<ul style="list-style-type: none"> <li>• Cisco DNA Center displays a status message indicating that connectivity with a node has been lost.</li> <li>• The UI remains usable on the remaining two nodes when using the VIP.</li> <li>• Services that were running on the failed node are migrated to other nodes.</li> <li>• The status of services running on the failed node may be set to either <code>NodeLost</code> or <code>Unknown</code>.</li> <li>• The NBI on the failed node is not accessible, while the NBI on the remaining two nodes remain operational.</li> </ul> <p>After the node is restored, and before the node rejoins the cluster:</p> <ul style="list-style-type: none"> <li>• Cisco DNA Center provides a status message indicating that cluster operation has resumed.</li> <li>• Pending UI calls that have not timed out complete.</li> <li>• Service requests that were pending on the failed node are completed on the node that the service was migrated to.</li> </ul> <p>After the node rejoins the cluster:</p> <ul style="list-style-type: none"> <li>• Data on the restored node is synched with other cluster members.</li> <li>• Services that were running on the failed node are stopped.</li> <li>• All service requests that were pending on the failed node are stopped.</li> <li>• Assurance UI selections operate as expected.</li> </ul>
Yes	Two nodes fail or are unreachable.	<p>The cluster is broken and the UI is not accessible until connectivity has been restored.</p> <ul style="list-style-type: none"> <li>• If the nodes recover, operations resume and the data shared by cluster members is synced.</li> <li>• If the nodes do not recover, contact the Cisco TAC for assistance.</li> </ul>
Yes	A node fails and needs to be removed from a cluster.	Contact the Cisco TAC for assistance.
No	All nodes lose connectivity with one another.	The UI is not accessible until connectivity has been restored. Once connectivity has been restored, operations resume and the data shared by cluster members is synced.
Yes	A backup is scheduled and a node goes down due to a hardware failure.	Contact the Cisco TAC for a replacement node, as well as assistance with joining the new node to the cluster and restoring services on the two remaining nodes.

Requires User Action	Failure Scenario	HA Behavior
Yes	A red banner in the UI indicates that a node is down: "Assurance services are currently down. Connectivity with host <IP-address> has been lost."	The banner indicates that the node is down. As a result, Assurance data collection and processing stops and data will not be available. If the node comes back up, your Assurance functionality is restored. If the failure is related to a hardware failure, do the following: <ol style="list-style-type: none"> <li>1. Remove the node that failed. See <a href="#">Remove the Failed Node, on page 7</a>.</li> <li>2. Add a new node to replace the one that failed. See <a href="#">Add a Replacement Node, on page 9</a>.</li> </ol>
Yes	A red banner in the UI indicates that a node is down, but eventually changes to yellow with this message: "This IP address is down."	The system is still usable. Investigate why the node is down, and bring it back up.
Yes	A failure occurs while upgrading a cluster.	Contact the Cisco TAC for assistance.
No	An appliance port fails.	<ul style="list-style-type: none"> <li>• Cluster port: Cisco DNA Center detects the failure within 5 minutes and times the user out. After 5 minutes, you should be able to log back in. A banner then appears, indicating the services that are currently unavailable. Service failover completes within 10 minutes. The areas of the UI you can access will depend on which services have been restored. After the services that were unavailable are fully restored, the banner closes.</li> <li>• Enterprise port: Cisco DNA Center might not be able to reach and manage your network.</li> <li>• Management port: Any upgrades and image downloads that are currently in progress will fail and northbound interface operations will also be affected.</li> </ul>
Yes	Appliance hardware fails.	<p>Replace the hardware component (such as a fan, power supply, or disk drive) that failed. Because multiple instances of these components are found in an appliance, the failure of one component can be tolerated temporarily.</p> <p>As the RAID controller syncs a newly added disk drive with the other drives on the appliance, there might be a degradation in performance on the I/O system while this occurs.</p>

## Explanation of Pending State During a Failover

A pod that is in Pending state behaves as follows:

- Stateful set: The pod has some type of data storage. These pods are node bound using [local persistent volume \(LPV\)](#)—when the node is down, all stateful sets on that node move to Pending state. Stateful examples are MongoDB, Elasticsearch, and Postgres.

- DaemonSet: By design, the pod is strictly **node bound**. DaemonSet examples are agent, broker-agent, and keepalived.
- Stateless/**deployment**:
  - While the pod doesn't have a data to store of its own, it uses a stateful set to store and/or retrieve data.
  - Deployment scale varies. Some deployments have 1x pod instance (such as spf-service-manager-service); some have 2x pod instances (such as apic-em-inventory-manager-service); some have 3x pod instances (such as kong, platform-ui, collector-snmp).
  - The 1x stateless pods are free to move across nodes based on the current state of the cluster.
  - The 2x stateless pods have flexibility to move across nodes, but no two instances of stateless pods can run on the same node.
  - The 3x stateless pods have node anti-affinity, meaning no two instances can run on the same node.

---

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2022–2023 Cisco Systems, Inc. All rights reserved.