Cisco Digital Network Architecture Center High Availability Guide, Release 1.3

First Published: 2019-05-31

High Availability Overview

Cisco DNA Center's high availability (HA) framework is designed to reduce the amount of downtime that results from failures and make your network more resilient when they take place. When a failure occurs, this framework helps to restore your network to its previous operational state. If this is not possible, Cisco DNA Center will indicate that there is an issue requiring your attention.

Any time Cisco DNA Center's HA framework determines that a change on a cluster node has taken place, it synchronizes this change with the other nodes. 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). As you go through this guide, note the following:

- It uses the terms *seed* and *master* interchangeably. The seed node (master node) is the node where Elasticsearch is running in the NDP namespace.
- In this release, Cisco DNA Center only provides HA support for Automation functionality. HA for Assurance is not supported at this time.

High Availability Requirements

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

- Your cluster consists of three appliances with the same Cisco part number (such as DN2-HW-APL for the 44 core Cisco DNA Center M5 appliance).
- The appliances are running the same version of Cisco DNA Center 1.2.8 or later. For example, if a patch for version 1.2.8 is installed on one cluster node, you must also install the same patch onto the other cluster nodes in order for HA to operate.

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 multi-node 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 multi-node 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 multi-node 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 Cisco DNA Center system upgrade. Once the system upgrade does complete, you can verify its success in the GUI by accessing **System Settings** > **Software Updates** > **Updates** and checking the installed version.

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

We recommend that you set up a cluster consisting of three nodes: one seed node and two non-seed nodes. 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.
- When you are configuring cluster nodes, do *not* specify a link-local subnet (169.x.x.x) as the cluster or services subnet because its addresses are used by the Cisco DNA Center internal network.

Note Subnets must conform with the IETF RFC 1918 specification for private networks. For details, see RFC 1918, Address Allocation for Private Internets, and the Wikipedia article "Private network".)

- 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 procedures:

Procedure

Step 1 Configure Cisco DNA Center on the first node in your cluster using the Maglev Configuration wizard. For information about this procedure, see the "Configure the Master Node" topic in the Cisco Digital Network Architecture Center Installation Guide.

Step 2 Configure Cisco DNA Center on the second node in your cluster using the Maglev Configuration wizard. For information about this procedure, see the "Configure Add-On Nodes" topic in the Cisco Digital Network Architecture Center Installation Guide.

Step 3 Configure Cisco DNA Center on the third node in your cluster using the Maglev Configuration wizard. For information about this procedure, see the "Configure Add-On Nodes" topic in the Cisco Digital Network Architecture Center Installation Guide.

Step 4 Enable high availability on your cluster:

a) Click ⚙ and then choose **System Settings**.

The **System 360** tab is displayed by default.

b) In the **Hosts** area, click **Enable Service Distribution**.

Note

- After you click **Enable Service Distribution** in the GUI, Cisco DNA Center enters into maintenance mode. In this mode, Cisco DNA Center is unavailable until the process completes. 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

In order to run maglev commands successfully on the nodes in your cluster, do the following:

Before you begin

Note You only need to complete this procedure before you run the first maglev command in a session. You do not need to complete it again unless you close the current session and start a new one.

Note When you run a command in an SSH client, you may get an error message that indicates the RSA host key has been changed and prompts you to add the correct key to the ~/.ssh/known_hosts file. This typically happens when an appliance has been reimaged using a different IP address from the one that was specified for the appliance previously. If this happens, do the following:

1. Determine the IP address that is assigned to your appliance: **cat ~/.ssh/known_hosts**

   where ~ represents the directory in which the known_host file resides on your machine.

   The resulting output will look similar to the following example:

   ```
   [192.168.254.21]:2222 ecdsa-sha2-nistp256
   AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBA19/31YV+cQvI1rmIVl/
   CaE/BqCdeg5Xr/pSOtwNnKB6eDrXvLSAUMz+EED339GvbkxT/DdsdGZn2BeWHIifuY=
   ```

2. Remove all of the keys associated with this IP address from the known_hosts file: **ssh-keygen -R** *appliance's-IP-address*

   Continuing our example, you would run the following command: **ssh-keygen -R 192.168.254.21:2222**

   Note Another option is to delete the ~/.ssh/known_hosts file before proceeding to the next step.

3. Run the command you tried to run previously.

Procedure

Step 1 In an SSH client, enter the following command:

ssh *node's-IP-address* **-l maglev -p 2222**

Step 2 If you see a message indicating that the node's authenticity cannot be established, enter **yes** when prompted to continue.

Step 3 Enter the Linux password configured for the node's maglev user.

Step 4 Enter the maglev command that you want to run.

Step 5 Enter the password configured for Cisco DNA Center's default admin superuser.

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 or preparing a node for Return Merchandise Authorization (RMA) and rebooting a node (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.

Single Cluster Node

- If a node is being shut down for RMA, drain it and then remove the node from the cluster:
  1. **maglev node drain** *node's-IP-address*
  2. **maglev node remove** *node's-IP-address*
- Reboot a cluster node that requires a reboot:

  sudo shutdown -r now

Note You do not need to drain the node.

All Cluster Nodes

- To shut down all of the nodes in a cluster for planned maintenance, run the following command on the three nodes simultaneously:

  sudo shutdown -h now

  When you are ready to bring the nodes back up, power on all of the nodes simultaneously.
- To reboot all of the nodes in a cluster after making a hardware or software configuration change, run the following command on the three nodes simultaneously:

  sudo shutdown -r now

Recover a Failed Cluster Node

If a node that belongs to a three-node cluster fails, it usually takes 30 minutes for the cluster to recover: five minutes to detect that the node is down and 25 minutes to move services to another node. After five minutes, the following banner message is displayed: `Automation and Assurance services are currently down. Connectivity with node` *node_details* `has been lost.` To recover the failed node, do the following:

Procedure

Step 1 Log in to a healthy cluster node and run the following command: **maglev node remove** *failed-node's-IP-address*

This will remove the faulty node from the cluster.

Step 2 Enter the **maglev package status** command on the active node.

Step 3 Contact the Cisco TAC, give them the output of that command, and ask for an ISO that matches your version.

Step 4 To add back the removed node, you must reinstall it using the Maglev Configuration wizard's **Join a Cisco DNA Center Cluster** option.

For information about using the Maglev Configuration wizard, see the "Configure Add-On Nodes" topic in the Cisco Digital Network Architecture Center Installation Guide.

Step 5 Redistribute services among the cluster nodes to optimize HA operation. To do so, click ⚙ and then choose **System Settings**. In the **System 360** tab > **Hosts** area, click **Enable Service Distribution**.

Replace a Failed Seed Node

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

1. Remove the failed node from your cluster.

   See Remove the Failed Seed Node.

2. Replace the failed node with another node.

   See Add a New Seed Node.

Remove the Failed Seed Node

If a seed node fails, you must remove it so that you can replace it with a working node. Removing the seed node takes about 30 minutes.

This section applies only if the failure is due to a hardware failure.

Note When you remove a seed node, its existing data is lost, but the remaining nodes begin collecting new data. This data is lost as soon as the new seed node is added and begins collecting data.

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.
- Allocate at least 30 minutes to perform this procedure.

Procedure

Step 1 (Optional) If you need to remove an Assurance seed node, complete the following actions to identify it:

1. Run the following command: **magctl appstack status ndp | grep elastic**
2. Locate the `elasticsearch-0` entry.

   The seed node's IP address is listed in the **Node** column.

Step 2 Shut down the node that you want to remove.

The shutdown process takes about 10 minutes.

Step 3 Verify that the node is down:

magctl node display

The node status should be `NOT_READY`.

Step 4 Check the appstack status:

magctl appstack status

The pods for the node that was shut down should show `NODE LOST` or `Pending` as their status.

Step 5 Log in to one of the nodes that you are *not* removing (a non-seed node):

maglev login -u admin -p *admin-password* **-c** *node's-IP-address***:443**

Step 6 Remove the failed seed node from the cluster:
maglev node remove *node's-IP-address*

The node removal process takes about 30 minutes to complete.

Step 7 Check that all services are running on the remaining two nodes:

magctl node display

magctl appstack status

Add a New Seed Node

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

Before you begin

Make sure that you complete the following tasks:

- Remove the failed seed node. For information, see Remove the Failed Seed Node, on page 7.
- Allocate at least 30 minutes to perform this procedure.

Procedure

Step 1 On the new node, install the same software version that the other nodes in the cluster are running.

During the installation, choose the **Join a Cisco DNA Center Cluster** option and enter the required configuration information using the Maglev Configuration wizard. For information, see the Cisco Digital Network Architecture Center Installation Guide.

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

magctl node display

The new node should show the `Ready` status.

Step 3 From the new node, do the following:

a) Enter the following command:

maglev node allow *node's-IP-address*

b) Redistribute services to the new node:

maglev service nodescale refresh

c) Verify that services have been redistributed:

magctl appstack status

The new node should show a `Running` status.

Step 4 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
- As soon as the old data is restored, the new node begins gathering new data, and any data that the remaining nodes gathered while the new node was offline is lost.
- If you are adding a new Assurance seed node, configure the same IP address that was used by the Assurance seed node you are replacing.

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.

For information about known HA bugs and workarounds, see "Open Bugs—HA" in the Release Notes for Cisco Digital Network Architecture Center.

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 this guide.
No	A node fails, is unreachable, or experiences a service failure for less than 5 minutes.	• The UI is not accessible for 5 minutes after a node fails. • Services that were running on the failed node are not migrated to other nodes. • The northbound interface (NBI) remains usable on the remaining two nodes when using the VIP. • VIP connectivity will be restored after failover, and API calls recover after services are up and running. After the node is restored: • Data on the restored node is synched with other cluster members. • Pending UI and NBI calls that have not timed out complete.

Requires User Action	Failure Scenario	HA Behavior
No	A non-seed node fails, is unreachable, or experiences a service failure for longer than 5 minutes.	• After 5 minutes, Cisco DNA Center displays a status message indicating that connectivity with a node has been lost. • The UI remains usable on the remaining two nodes when using the VIP. • Services that were running on the failed node are migrated to other nodes. • The NBI on the failed node is not accessible, while the NBI on the remaining two nodes remain operational. After the node is restored, and before the node rejoins the cluster: • Cisco DNA Center provides a status message indicating that cluster operation has resumed. • Pending UI calls that have not timed out complete. • Service requests that were pending on the failed node are completed on the node that the service was migrated to. After the node rejoins the cluster: • Data on the restored node is synched with other cluster members. • Services that were running on the failed node are stopped. • All service requests that were pending on the failed node are stopped.

Requires User Action	Failure Scenario	HA Behavior
No	A seed node fails, is unreachable, or experiences a service failure for longer than 5 minutes.	• Cisco DNA Center displays a status message indicating that connectivity with a node has been lost. • The UI remains usable on the remaining two nodes when using the VIP. • Services that were running on the failed node are migrated to other nodes. • The status of services running on the failed node may be set to `waiting`. • The NBI on the failed node is not accessible, while the NBI on the remaining two nodes remain operational. • When Assurance is running only on the seed node, the status of Assurance UI selections is set to `pending`. • When Assurance has multiple instances running, Assurance UI selections continue to operate. After the node is restored, and before the node rejoins the cluster: • Cisco DNA Center provides a status message indicating that cluster operation has resumed. • Pending UI calls that have not timed out complete. • When Assurance is running only on the seed node, the status of Assurance UI selections is set to `pending`. • Service requests that were pending on the failed node are completed on the node that the service was migrated to. • When Assurance has multiple instances running, Assurance UI selections continue to operate. After the node rejoins the cluster: • Data on the restored node is synched with other cluster members. • Services that were running on the failed node are stopped. • All service requests that were pending on the failed node are stopped. • Assurance UI selections operate as expected.

Requires User Action	Failure Scenario	HA Behavior
Yes	Two nodes fail or are unreachable.	The cluster is broken and the UI is not accessible until connectivity has been restored. • If the nodes recover, operations resume and the data shared by cluster members is synced. • If the nodes do not recover, you will need to do the following: 1. Restore the seed node. See "Configure the Master Node" in the *Cisco Digital Network Architecture Center Appliance Installation Guide*. 2. Restore the other cluster nodes. See "Configure Add-On Nodes" in the *Cisco Digital Network Architecture Center Appliance Installation Guide*.
Yes	A node fails and needs to be removed from a cluster.	Do the following: 1. Remove the node that failed. 2. Add a new node to replace the one that failed. 3. Redistribute services to the new node, making it an official cluster member. For more information, see Recover a Failed Cluster Node.
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 seed node goes down due to a hardware failure.	Do the following: 1. Remove all of the cluster nodes by running the following command on each node: **maglev node remove** *node's-IP-address* 2. Contact the Cisco TAC for a replacement node, and for assistance with joining the new node to the cluster.
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 seed node is down, and Assurance data has been lost. If the seed node comes back up, your Assurance functionality is restored. But if the failure is related to a hardware failure, do the following: 1. Remove the seed node that failed. See Remove the Failed Seed Node, on page 7. 2. Add a new node to replace the one that failed. See Add a New Seed Node, on page 8.

Requires User Action	Failure Scenario	HA Behavior
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.	Do the following: 1. Remove all of the cluster nodes by running the following command on each node: **maglev node remove** *node's-IP-address* 2. Restore the seed node. See "Configure the Master Node" in the *Cisco Digital Network Architecture Center Appliance Installation Guide*. 3. Restore the other cluster nodes. See "Configure Add-On Nodes" in the *Cisco Digital Network Architecture Center Appliance Installation Guide*.
No	An appliance port fails.	• 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. • Enterprise port: Cisco DNA Center might not be able to reach and manage your network. • Management port: Any upgrades and image downloads that are currently in progress will fail and northbound interface operations will also be affected.
Yes	Appliance hardware fails.	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. 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.