



# Applications

---

Cisco Data Center Network Manager (DCNM) uses the application framework to host various plugins and microservices to support operations and related features in Cisco DCNM.

The Applications Framework provides the following features:

- An infrastructure for hosting applications that require more system resources as the scale of the network increases.
- An independent application development-deployment-management lifecycle for applications.

Cisco DCNM Applications Framework supports two modes namely clustered mode and unclustered mode. In clustered mode, the compute nodes are clustered together whereas in the latter only the DCNM server nodes namely the active/standby exist. Most of the applications for ex: Network Insights require clustered setup to be ready before they can be uploaded and deployed using DCNM Applications Framework.

- [Cisco DCNM in Unclustered Mode, on page 1](#)
- [Cisco DCNM in Clustered Mode, on page 2](#)
- [Installing and Deploying Applications, on page 9](#)
- [Application Framework User Interface, on page 13](#)
- [Compute, on page 15](#)
- [Preferences, on page 16](#)
- [Enabling the Compute Cluster, on page 16](#)
- [Failure Scenario, on page 18](#)
- [Converting from Unclustered to Clustered Mode with Existing Elasticsearch Data, on page 18](#)

## Cisco DCNM in Unclustered Mode

From Cisco DCNM Release 11.0(1), the unclustered mode is the default deployment mode in both Standalone and Native HA environment. In this mode, the Cisco DCNM runs some of its internal services as containers, also.

- Endpoint Locator is running as a container application, from Cisco DCNM Release 11.1(1).
- Configuration Compliance service is a container application, from Cisco DCNM Release 11.0(1).
- Virtual Machine Manager (VMM) is also a container application, from Cisco DCNM Release 11.0(1)

Cisco DCNM leverages resources from the Standby node for running some containers applications. The Cisco DCNM Active and Standby nodes work together to extend resources to the overall functionality and deployment of DCNM and its applications. However, it has limited resources to run some of the advanced applications and to extend the system to deploy more applications delivered through the Cisco AppCenter. For example, you cannot deploy the Network Insights applications that are downloaded from the Cisco AppCenter, for production, in unclustered mode.

To install and deploy applications, see [Installing and Deploying Applications, on page 9](#).

For best practices and recommended deployments for IP address configurations of all interfaces of the Cisco DCNM and Compute nodes, see *Best Practices for Deploying Cisco DCNM and Computes* in *Cisco DCNM Installation Guide*, for your deployment type.

## Cisco DCNM in Clustered Mode

By default, the clustered mode is not enabled on the Cisco DCNM deployments. Enable the cluster mode after you deploy the Cisco DCNM Server. In a clustered mode, the Cisco DCNM Server with more compute nodes provides an architecture to expand resources, as you deploy more applications.

Compute nodes are scale out application hosting nodes that run resource-intensive services to provide services to the larger Fabric. When compute nodes are added, all services that are containers, run only on these nodes. This includes Config Compliance, Endpoint Locator, and Virtual Machine Manager. The Elasticsearch time series database for these features runs on compute nodes in clustered mode. In the clustered mode deployment, the DCNM Servers do not run containerized applications. All applications that work in unclustered mode works in the clustered mode, also.



---

**Note** The clustered mode is not supported on Cisco DCNM for Media Controller deployment.

---

From Cisco DCNM Release 11.1(1), in a Native HA setup, 80 switches with Endpoint Locator, Virtual Machine Manager, config compliance are validated in the unclustered mode. For a network exceeding 80 switches, with these features in a given Cisco DCNM instance (maximum qualified scale is 256 switches), we recommend that you enable clustered mode.

While the Cisco DCNM core functionalities only run on the Native HA nodes, addition of compute nodes beyond 80 switches is to build a scale-out model for Cisco DCNM and related services.

From Release 11.2(1), you can configure IPv6 address for Network Management for compute clusters. However, DCNM does not support IPv6-address for containers, and must connect to DCNM using only IPv4 address only.

For best practices and recommended deployments for IP address configurations of all interfaces of the Cisco DCNM and Compute nodes, see *Best Practices for Deploying Cisco DCNM and Computes* in *Cisco DCNM Installation Guide*, for your deployment type.

## Requirements for Cisco DCNM Clustered Mode



---

**Note** We recommend that you install the Cisco DCNM in the Native HA mode.

---

### Cisco DCNM LAN Deployment Without Network Insights (NI)

**Table 1: Upto 80 Switches**

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes	NA	—	—	—	—

**Table 2: 81–250 Switches**

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes x 3	OVA/ISO	16 vCPUs	64G	500G HDD	3xNIC

### Cisco DCNM LAN Deployment with NIA and NIR Software Telemetry



**Note** We recommend that you install the Cisco DCNM in the Native HA mode.

**Table 3: Upto 80 Switches**

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes x 3	OVA/ISO	16 vCPUs	64G	500G HDD	3xNIC

**Table 4: 81–250 Switches**

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
DCNM	OVA/ISO	16 vCPUs	32G	500G HDD	3xNIC
Computes x 3	ISO	32 vCPUs	256G	2.4-TB HDD	3xNIC <sup>1</sup>

<sup>1</sup> Network card: Quad-port 10/25G

### Subnet Requirements

In general, Eth0 of the Cisco DCNM server is used for Management, Eth1 is used to connect Cisco DCNM Out-Of-Band with switch management, and eth2 is used for In-Band front panel connectivity of Cisco DCNM. The same concept extends into compute nodes as well. Some services in clustered mode have other requirements. Some services require a switch to reach into Cisco DCNM. For example, Route Reflector to Endpoint Locator connection or switch streaming telemetry into the Telemetry receiver service of the application require a switch to reach DCNM. This IP address needs to remain sticky during all failure scenarios. For this purpose, an IP pool must be provided to Cisco DCNM at the time of cluster configuration for both out-of-band and In-Band subnets.

### Telemetry NTP Requirements

For telemetry to work correctly, the Cisco Nexus 9000 switches and Cisco DCNM must be time that is synchronized. Cisco DCNM telemetry manager does the required NTP configuration as part of enablement. If there is a use-case to change the NTP server configuration manually on the switches ensure that the DCNM and the switches are time synchronized, always. To set up telemetry network configuration, see .

## Installing a Cisco DCNM Compute



---

**Note** With Native HA installations, ensure that the HA status is **OK** before DCNM is converted to cluster mode.

---

A Cisco DCNM Compute can be installed using an ISO or OVA of a regular Cisco DCNM image. It can be deployed directly on a bare metal using an ISO or a VM using the OVA. After you deploy Cisco DCNM, using the DCNM web installer, choose **Compute** as the install mode for Cisco DCNM Compute nodes. On a Compute VM, you will not find DCNM processes or postgres database; it runs a minimum set of services that are required to provision and monitor applications.

## Networking Policies for OVA Installation

For each compute OVA installation, ensure that the following networking policies are applied for the corresponding vSwitches of host:

- Log on to the vCenter.
- Click on the Host on which the computes OVA is running.
- Click **Configuration > Networking**.
- Right click on the port groups corresponding to the eth1 and eth2, and select **Edit Settings**.  
The **VM Network - Edit Settings** is displayed.
- In Security settings, for **Promiscuous** mode, select **Accepted**.
- If a DVS Port-group is attached to the compute VM, configure these settings on the **vCenter > Networking > Port-Group**. If a normal vSwitch port-group is used, configure these settings on **Configuration > Networking > port-group** on each of the Compute's hosts.

Figure 1: Security Settings for vSwitch Port-Group

The screenshot shows the 'VM Network - Edit Settings' dialog box with the 'Security' tab selected. The 'Properties' list on the left includes 'Security', 'Traffic shaping', and 'Teaming and failover'. The 'Security' section contains three rows of settings:

Property	Setting	Override	Value
Promiscuous mode		<input checked="" type="checkbox"/>	Accept
MAC address changes		<input checked="" type="checkbox"/>	Accept
Forged transmits		<input checked="" type="checkbox"/>	Accept

Buttons for 'CANCEL' and 'OK' are located at the bottom right of the dialog.

Figure 2: Security Settings for DVSwitch Port-Group

The screenshot shows the 'OobFabric - Edit Settings' dialog box with the 'Security' tab selected. The 'General' list on the left includes 'General', 'Advanced', 'VLAN', 'Security', 'Teaming and failover', 'Traffic shaping', 'Monitoring', and 'Miscellaneous'. The 'Security' section contains three rows of settings:

Property	Setting	Value
Promiscuous mode		Accept
MAC address changes		Accept
Forged transmits		Accept

Buttons for 'CANCEL' and 'OK' are located at the bottom right of the dialog.



**Note** Ensure that you repeat this procedure on all the hosts, where a Compute OVA is running.

## Adding Computes into the Cluster Mode

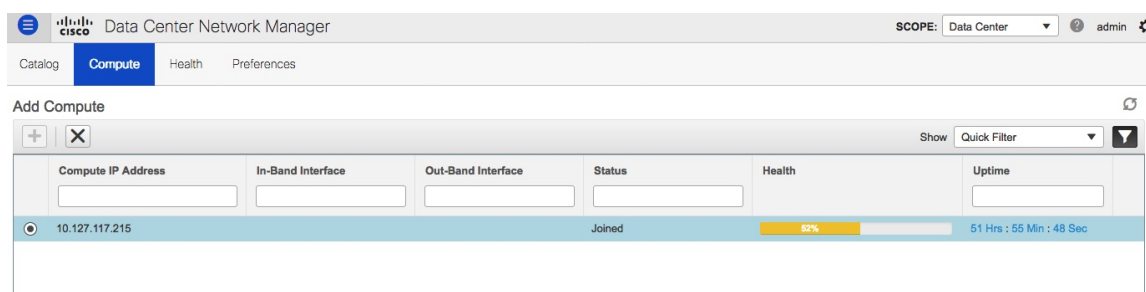
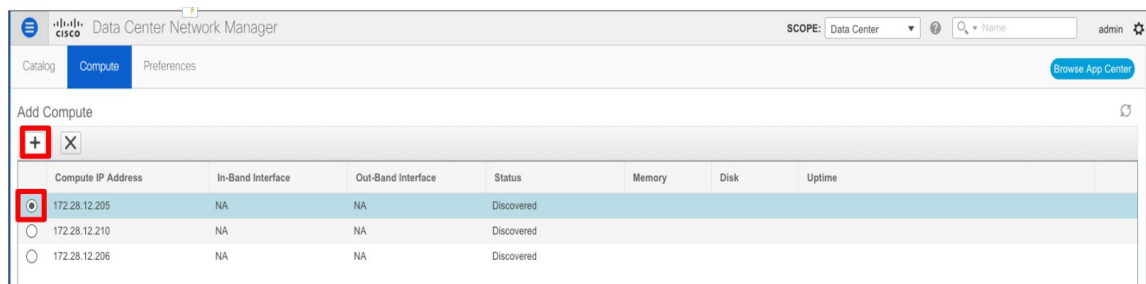
To add computes into the cluster mode from Cisco DCNM Web UI, perform the following steps:

### Procedure

**Step 1** Choose **Applications > Compute**.

The Compute tab displays the computes enabled on the Cisco DCNM.

**Step 2** Select a Compute node which is in **Discovered** status. Click the **Add Compute (+)** icon.

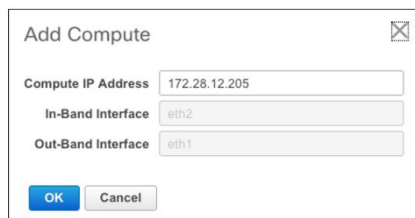


- While using Compute, ensure that Cisco DCNM GUI shows nodes as Joined.
- Offline indicates some connectivity issues, therefore no applications are running on Offline Computes.
- Failed indicates that the compute node could not join the cluster.
- Health indicates the amount of free memory and disk on the Compute node. The application provides more detailed statistics.
- Most applications do not function properly if there are less than three computes, while a loss of a single Compute node is mostly fine. In such cases, refer to the requirements of the individual applications.
- If the Performance Manager was stopped during or after the inline upgrade and after all the computes have changed to Joined, you must restart the Performance Manager.

The Compute window allows you to monitor the health of computes. The health essentially indicates the amount of memory that is left in the compute, this is based on applications that are enabled. If a Compute is not properly communicating with the DCNM Server, the status of the Compute appears as Offline, and no applications are running on Offline Computes. Most applications do not function properly if there are less than three computes, while a short loss of a single Compute node is mostly fine. In such cases, refer to the requirements of the individual applications.

**Step 3** In the **Add Compute** dialog box, view the **Compute IP Address**, **In-Band Interface**, and the **Out-Band Interface** values.

**Note** The interface value for each compute node is configured by using the `appmgr afw config-cluster` command.



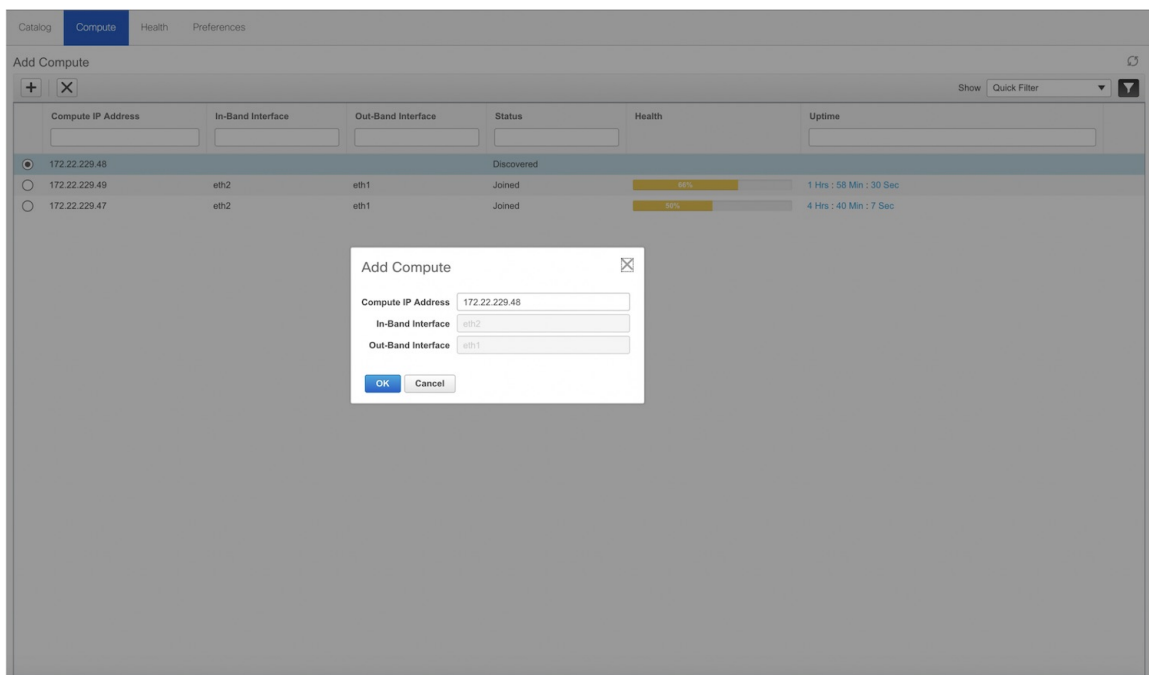
Add Compute

Compute IP Address: 172.28.12.205

In-Band Interface: eth2

Out-Band Interface: eth1

OK Cancel



Add Compute

Compute IP Address: 172.22.229.48

In-Band Interface: eth2

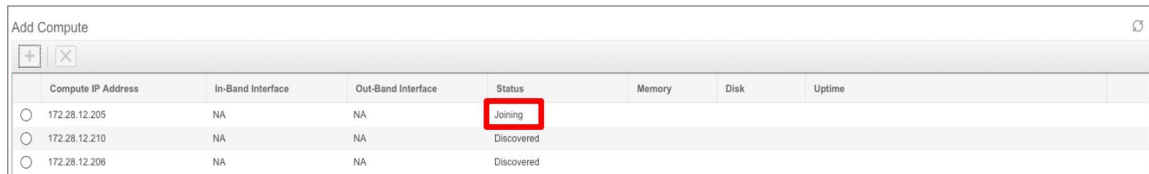
Out-Band Interface: eth1

OK Cancel

Compute IP Address	In-Band Interface	Out-Band Interface	Status	Health	Uptime
172.22.229.48			Discovered		
172.22.229.49	eth2	eth1	Joined	60%	1 Hrs : 56 Min : 30 Sec
172.22.229.47	eth2	eth1	Joined	50%	4 Hrs : 40 Min : 7 Sec

**Step 4** Click **OK**.

The Status for that Compute IP changes to **Joining**.



Add Compute

Compute IP Address: 172.28.12.205

In-Band Interface: NA

Out-Band Interface: NA

OK Cancel

Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime
172.28.12.205	NA	NA	Joining			
172.28.12.210	NA	NA	Discovered			
172.28.12.206	NA	NA	Discovered			

Wait until the Compute IP status shows **Joined**.

Add Compute							
Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime	
<input type="radio"/> 172.28.12.205	eth2	eth1	Joined	60%	99%	-- Hrs : 4 Min : 17 Sec	
<input type="radio"/> 172.28.12.210	NA	NA	Discovered				
<input type="radio"/> 172.28.12.206	NA	NA	Discovered				

**Step 5** Repeat the above steps to add the remaining compute node.

All the Computes appear as **Joined**.

Add Compute							
Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime	
<input type="radio"/> 172.28.12.205	eth2	eth1	Joined	40%	99%	183 Hrs : 15 Min : 41 Sec	
<input type="radio"/> 172.28.12.210	eth2	eth1	Joined	57%	98%	-- Hrs : 4 Min : 9 Sec	
<input type="radio"/> 172.28.12.206	eth2	eth1	Joined	58%	98%	-- Hrs : 2 Min : 18 Sec	

**Note** When you install compute as a virtual machine on the VMware platform, vSwitch or DV switch port groups associated eth1 and eth2 must allow for packets that are associated with Mac address other than eth1 and eth2 to be forwarded.

## Preferences

This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute cluster connectivity and configure the Cluster Connectivity preferences.

### Object Archival Configuration

The NIA application collects tech support logs for all switches in Fabric, and determines the advisory, based on the data. The logs are saved on the Cisco DCNM server for further analysis or troubleshooting. If you need to download these logs before their life span ends or to create some space on the DCNM server, you can move the logs to a remote server.

In the **URI** field, enter the relative path to the archive folder, in the format `host[:port]/[path to archive]`. Enter the username and password to access the URI, in the **username** and **Password** field. Click **Submit** to configure the remote server.



## Telemetry Network and NTP Requirements

For the Network Insights Resource (NIR) application, a UTR micro-services running inside the NIR receives the telemetry traffic from the switches either through Out-Of-Band (Eth1) or In-Band (Eth2) interface. By default, the telemetry is configured, and is streaming via the Out-Of-Band interface. You can choose to change it to In-Band interface as well.

### Telemetry Using Out-of-band (OOB) Network

By default, the telemetry data is streamed through the management interface of the switches to the Cisco DCNM OOB network eth1 interface. This is a global configuration for all fabrics in Cisco DCNM LAN Fabric Deployment, or switch-groups in Cisco DCNM Classic LAN Deployment. After the telemetry is enabled via NIR application, the telemetry manager in Cisco DCNM will push the necessary NTP server configurations to the switches by using the DCNM OOB IP address as the NTP server IP address. The following example is sample output for `show run ntp` command.

```
switch# show run ntp

!Command: show running-config ntp
!Running configuration last done at: Thu Jun 27 18:03:07 2019
!Time: Thu Jun 27 20:32:18 2019

version 7.0(3)I7(6) Bios:version 07.65
ntp server 192.168.126.117 prefer use-vrf management
```

### Telemetry Using In-Band (IB) Network:

The switches stream telemetry data through their front panel ports to Cisco DCNM assuming the connectivity from the switches to the Cisco DCNM In-Band network eth2 interface.

## Installing and Deploying Applications

The following sections describes how to download, add, start, stop, and delete applications from the Cisco DCNM Web UI.

### Download App from the App Store

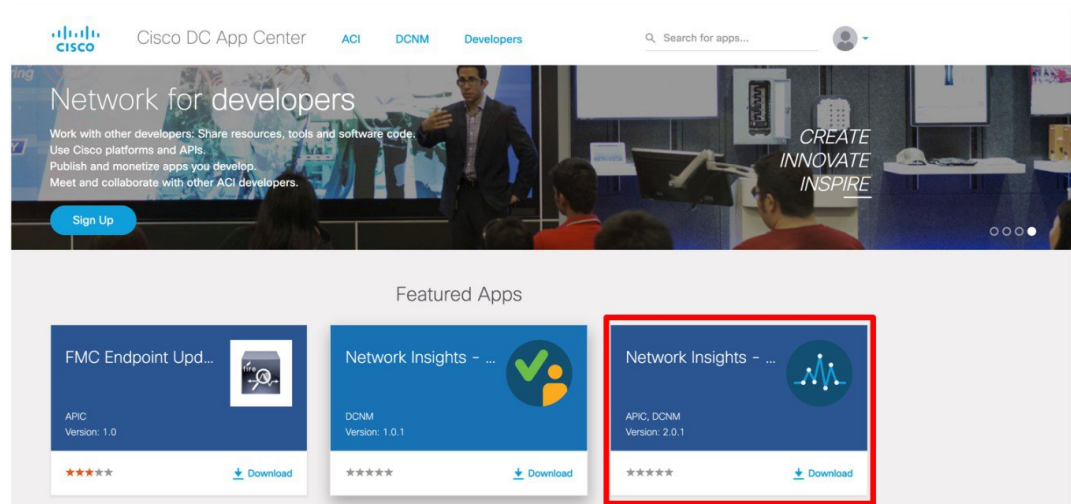
To download new applications from the Cisco DCNM Web UI, perform the following steps:

1. Choose **Applications**.

By default, the **Catalog** tab displays.

2. Click **Browse App Center** on the top-right corner on the window.

On the Cisco ACI App Center, locate the required application and click the download icon.



3. Save the application executable file on your local directory.

### Add a New Application to DCNM

To add new applications from the Cisco DCNM Web UI, perform the following steps:

1. Choose **Applications**.  
By default, the **Catalog** tab displays.
2. Click **Add Application (+)** icon.



On the Application Upload window, from the Type drop-down field, choose one of the following to upload the application.

From the Type drop-down list, select one of the following:

- If the file is located in a local directory, select **Local-file**.

In the Upload field, click **Select files...** Navigate to the directory where you have stored the application file.

Select the application file and click **Open**.

Click **Upload**.

- If the application is located on a remote server, select **Secure copy**.




---

**Note** Ensure that the remote server must be capable of serving Secure-copy (SCP).

---

In the URI field, provide the path to the application file. The path must be in `{host-ip}:{filepath}` format.

In the Username field, enter the username to access the URI.

In the Password field, enter the appropriate password for accessing the URI.

Click **Upload**.

After the application successfully uploaded, it is displayed in the Catalog window.

The green icon on the left-top corner indicates that the application is launched successfully and is operational. If there is no green icon on the application, it indicates that the application is not running. Click the application to Launch it.




---


**Note** Ensure that the Compute Cluster is enabled before you install applications. A few applications may not work if the compute cluster is configured after installing the applications.

---

Click the gear icon on the left-bottom on the application icon to view the Application Specifications. The Info tab displays the running container information. The Specs tab displays the configuration.

### Starting Application

After the application is installed on the Cisco DCNM server, you must deploy the application. Click on the Application to begin deployment. Cisco DCNM starts all the services in the backend that are required for the application.



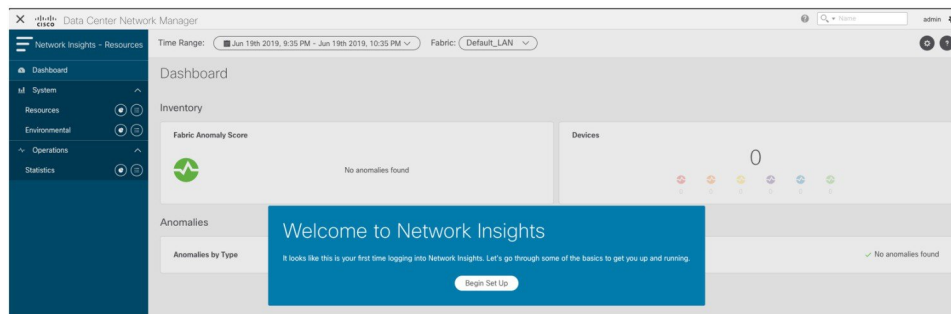
The green icon on the left-top corner indicates that the application is launched successfully and is operational.



Network Insights  
Resources(2.0.1)

The applications utilizing the Kafka infrastructure services require three actively joined compute nodes, when you begin the application. For example, NIR and NIA applications. If the application has a user interface, after the application is successfully started the UI redirects to the index page served by the application.

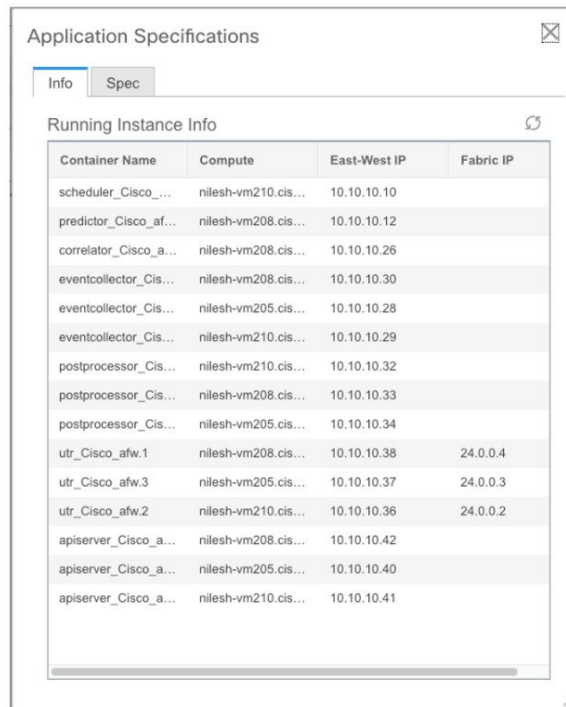
If the application has a user interface, after the application is successfully started the UI redirects to the index page served by the application.



To check the services that are running go back to **Applications > Catalog**. Click the gear icon on the left-bottom on the application icon to view the Application Specifications. The Info tab displays the running container information and the Specs tab displays the configuration as shown in the figures below.



Network Insights  
Resources(2.0.1)



The screenshot shows a window titled "Application Specifications" with two tabs: "Info" and "Spec". The "Info" tab is active, displaying "Running Instance Info" with a refresh icon. Below this is a table with four columns: "Container Name", "Compute", "East-West IP", and "Fabric IP".

Container Name	Compute	East-West IP	Fabric IP
scheduler_Cisco_...	nilesh-vm210.cis...	10.10.10.10	
predictor_Cisco_af...	nilesh-vm208.cis...	10.10.10.12	
correlator_Cisco_a...	nilesh-vm208.cis...	10.10.10.26	
eventcollector_Cis...	nilesh-vm208.cis...	10.10.10.30	
eventcollector_Cis...	nilesh-vm205.cis...	10.10.10.28	
eventcollector_Cis...	nilesh-vm210.cis...	10.10.10.29	
postprocessor_Cis...	nilesh-vm210.cis...	10.10.10.32	
postprocessor_Cis...	nilesh-vm208.cis...	10.10.10.33	
postprocessor_Cis...	nilesh-vm205.cis...	10.10.10.34	
utr_Cisco_afw.1	nilesh-vm208.cis...	10.10.10.38	24.0.0.4
utr_Cisco_afw.3	nilesh-vm205.cis...	10.10.10.37	24.0.0.3
utr_Cisco_afw.2	nilesh-vm210.cis...	10.10.10.36	24.0.0.2
apiserver_Cisco_a...	nilesh-vm208.cis...	10.10.10.42	
apiserver_Cisco_a...	nilesh-vm205.cis...	10.10.10.40	
apiserver_Cisco_a...	nilesh-vm210.cis...	10.10.10.41	

For information on how to remove computes from the cluster, stopping or deleting the applications, see [Application Framework User Interface, on page 13](#).

### Stop and Delete Applications

To delete the applications from the Catalog on the Cisco DCNM Web UI, perform the following steps:

1. Choose **Applications**.

By default, the **Catalog** tab displays, showing all the installed applications.

2. Click the red icon on the right-bottom corner to stop the application.

3. Check the **Wipe Volumes** check box to erase all the data that is related to the application.

4. Click **Stop** to stop the application from streaming data from Cisco DCNM.

The Green icon disappears after the application is successfully stopped.

5. After you stop the application, click the **Waste Basket** icon to remove the application from the Catalog.

## Application Framework User Interface

To use the Applications Framework feature, in the Cisco DCNM home page's left pane, click **Applications**.

The Applications window displays the following tabs:

- **Catalog**—This tab lists the applications that are used by Cisco DCNM. These applications for performing various functions within Cisco DCNM. For more information, see *Catalog*.

- **Compute**—This tab displays the existing compute nodes. The tab shows nodes that are part of the hosting infrastructure. The uptime indicates how long they have been part of the infrastructure. In a High Availability (HA) setup, both the active and the standby nodes appear as joined. For more information, see [Compute, on page 15](#).



---

**Note** In the cluster mode, the Cisco DCNM servers will not appear under the Compute tab.

---

- **Preferences**—This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute the cluster connectivity and configure the Cluster Connectivity preferences. For more information, see [Preferences, on page 8](#).

Cisco DCNM uses the following applications:

- **Compliance**: This application helps in building fabrics for the Easy Fabric installation. The Compliance application runs as one instance per fabric. It is enabled when fabric is created. Similarly, it is disabled when fabric is deleted.
- **ElasticCluster**: This application is used for storing various database information.
- **Kibana**: This is an open-source data-visualization plug-in for Elasticsearch, which provides visualization capabilities. Cisco DCNM uses the Kibana application for the Media Controller, and Endpoint Locator, and Telemetry.
- **UTR**: The Universal Telemetry Receiver (UTR) receives and decodes the streaming data from the switches, and feeds the data to a Kafka storage. The UTR can dynamically scale and instantiate multiple instances of UTR. It is a containerized application, that initiates using the Cisco DCNM Application Hosting Framework.
- **vmmplugin**: The Virtual Machine Manager (VMM) plug-in stores all the computes and the virtual machine information that connects to the fabric or the switch groups that are loaded into Cisco DCNM. VMM gathers compute repository information and displays the VMs, VSwitches/DVS, hosts in the topology view.
- **Endpoint Locator**: The Endpoint Locator (EPL) feature allows real-time tracking of endpoints within a data center. The tracking includes tracing the network life history of an endpoint and getting insights into the trends that are associated with endpoint additions, removals, moves, and so on. An endpoint is anything with an IP and MAC address. In that sense, an endpoint can be a virtual machine (VM), container, bare-metal server, service appliance and so on.



---

**Note** In Release 11.1(1), the applications are view-only in the Cisco DCNM UI. You cannot control the applications from the user interface. For example, you cannot start or stop these applications from the **Catalog** tab of the **Applications** window.

---

# Compute

This tab displays the existing compute nodes. The tab shows nodes that are part of the hosting infrastructure. The uptime indicates how long they have been part of the infrastructure. In a High Availability (HA) setup, both the active and the standby nodes appear as joined. In clustered mode, the compute nodes status indicate if the nodes are joined or discovered.

Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime
172.28.12.205	NA	NA	Discovered			
172.28.12.210	NA	NA	Discovered			
172.28.12.206	NA	NA	Discovered			



**Note** If the NTP server for compute nodes is not synchronized with the NTP server for DCNM Servers (Active and Standby) and Computes, you cannot configure a cluster.

The certificates are generated with a timestamp. If you configure the Compute nodes using a different NTP server, the mismatch in timestamp will not allow to validate the certificates. Therefore, if the compute cluster is configured despite of a mismatch of NTP server, the applications will not function properly.



**Note** In clustered mode, the Cisco DCNM servers will not appear under the Compute tab.

The following table describes the fields that appear on **Applications > Compute**.

**Table 5: Field and Description on Compute Tab**

Field	Description
Compute IP Address	Specifies the IP Address of the Compute node.
In-Band Interface	Specifies the in-band management interface.
Out-Band Interface	Specifies the out-band management interface.
Status	Specifies the status of the Compute node. <ul style="list-style-type: none"> <li>• Joined</li> <li>• Discovered</li> <li>• Failed</li> <li>• Offline</li> </ul>

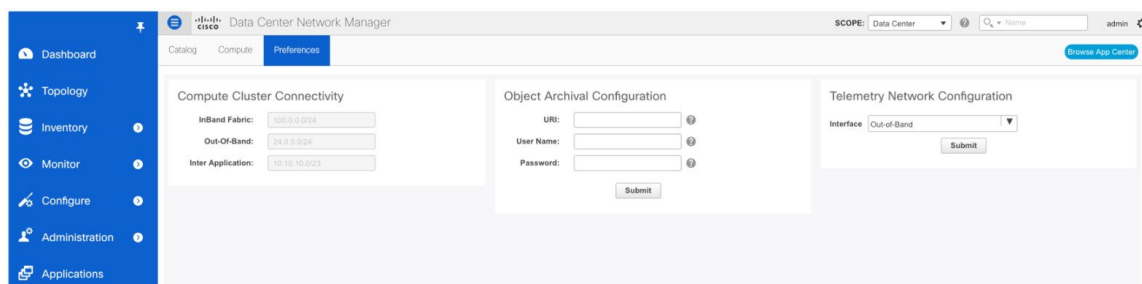
Field	Description
Memory	Specifies the memory that is consumed by the node.
Disk	Specifies the disk space that is consumed on the compute node.
Uptime	Specifies the duration of the uptime for a compute node.

When you install a compute node with correct parameters, it appears as **Joined** in the Status column. However, the other two computes appears as **Discovered**. To add computes to the cluster mode from Cisco DCNM Web UI, see [Adding Computes into the Cluster Mode, on page 6](#).

To configure or modify the Cluster Connectivity preferences, see [Preferences, on page 8](#).

## Preferences

This tab is relevant to the cluster mode of deployment, where the application instances are placed. This tab enables you to compute cluster connectivity and configure the Cluster Connectivity preferences.



### Object Archival Configuration

The NIA application collects tech support logs for all switches in Fabric, and determines the advisory, based on the data. The logs are saved on the Cisco DCNM server for further analysis or troubleshooting. If you need to download these logs before their life span ends or to create some space on the DCNM server, you can move the logs to a remote server.

In the **URI** field, enter the relative path to the archive folder, in the format `host[:port]/[path to archive]`. Enter the username and password to access the URI, in the **username** and **Password** field. Click **Submit** to configure the remote server.

## Enabling the Compute Cluster

The Compute cluster feature provides a scale mechanism to serve large fabrics and advanced features such as Telemetry, Advisories, and so on. The default status of Cisco DCNM in any of the install modes does not support having computes as a cluster.





**Note** Cisco DCNM LAN Fabric deployment supports compute clustered mode. The Classic LAN and the IP Fabric for Media (IPFM) installations do not support the Compute cluster functionality.

The values that you specify under the **Preferences** tab are not the subnets with their gateway. They are IP prefix pools of addresses that are used within the services that run inside a container. Hover your cursor over the ? to view the acceptable input. Cisco DCNM constantly communicates with DCNM controllers. This ensures that the required applications are running on only the Computes that are functioning. In a native HA environment, only the active DCNM initiates the compute cluster; the standby DCNM takes over if the active one fails. In Cisco DCNM 11.1(1), each Cisco DCNM server installation supports a maximum of three compute nodes.

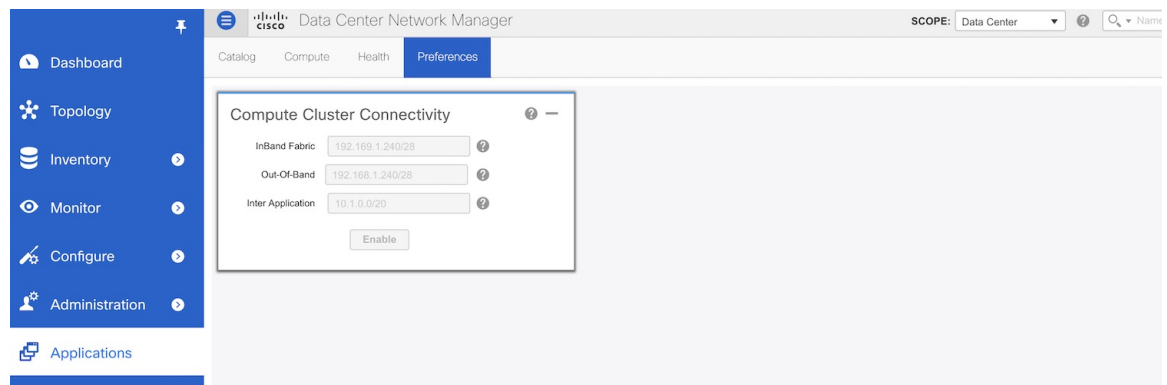
For more information about the Cisco DCNM Cluster mode, see the *Cisco DCNM Cluster Mode* section.

To enable the Compute Cluster Connectivity, perform the following steps:

## Procedure

**Step 1** In the Cisco DCNM home page's left pane, click **Applications**.

**Step 2** Click the **Preferences** tab.



**Step 3** In the **In-Band Fabric** field, specify the IP prefix pool to connect applications through the in-band path.

In-Band Fabric is required for applications that connect with switches over the Underlay or Front Panel network. For example, EPL and Telemetry over In-Band. These IP addresses are visible outside of the compute. The In-Band Pool Mask must be in range between /24-/28.

**Step 4** In the **Out-Of-Band** field, specify the IP prefix pool to connect applications with the fabric management network.

Out-Of-Band is required for applications that connect with switches over the Management network of the switches, e.g, Telemetry. These IP addresses are visible outside of the compute. The out-of-band pool range must be between /24 - /28.

**Step 5** In the **Inter Application** field, specify the IP prefix pool for inter-application communication.

The micro-services use this IP pool to communicate with each other. For example, Endpoint Locator trying to add entries to Elasticsearch. This network is not directly visible externally, but it is overlaid on top of the eth0 interfaces of the computes. The applications also communicate with the DCNM server using this network.

You can specify a pool that does not fall in either of the above subnets. Inter-Application Network is not visible outside of the DCNM cluster, it is an internal network for communication between multiple services. But, the network must not conflict with the endpoints that the services try to reach. The Inter-Application Pool mask must be between /20 to /24 and, it must be wider than the out-of-band or in-band pool mask.

**Step 6** Click **Enable**.

---

## Failure Scenario

Recommendation for minimum redundancy configuration with a DCNM OVA install is as follows:

- DCNM Active Node(Active) and compute node 1 in server1.
- DCNM Standby Node and compute node 2 in server2.
- Compute node 3 in server3.

When DCNM Active node is down, the Standby node takes full responsibility of running the core functionality.

When a compute node is down, the applications may continue to function with limited functionality. If this situation persists for a longer duration, it affects the performance and reliability of the applications. When more than one node is down, it affects the applications functionality and most of the applications fail to function.

You must maintain 3 compute nodes at any time. If a compute node goes down, rectify the issue as soon as possible, for the services to function as expected.

## Compute Node Disaster Recovery

When a compute node is lost due to a disaster and is irrecoverable, you must install another compute node with the same parameters. This will essentially appear as a reboot of the compute with lost data and it tries to join the cluster automatically. After it joins the cluster, all the data will synchronize from the other two compute nodes.

## Converting from Unclustered to Clustered Mode with Existing Elasticsearch Data



---

**Note** When you convert Cisco DCNM from unclustered mode to clustered mode, perform the following procedure to rebuild the database schema and preserve the Elasticsearch data.

---

Read through the complete procedure before proceeding with the Cluster mode conversion.

To convert from uncluster mode to clustered mode, perform the following steps.

## Procedure

---

- Step 1** From the left pane, choose **Preferences**, and then click the **Preferences** tab.
- Step 2** Take backup of the Elasticsearch data by executing the **appmgr afw backup-es** command.
- If the setup is HA, repeat the step on DCNM Standby.
- Step 3** Copy the tarball file produced from Step 2 to Compute 1.
- If the setup is HA, copy the tarball file on DCNM Standby to Compute 2.
- Step 4** Restore the Elasticsearch data by executing the **appmgr afw restore-es \$tarball** command on the Compute node, where \$tarball is the location of the tarball file copied in Step 3.
- If the setup is HA, repeat the step on Compute 2.
- Step 5** Proceed with adding the Compute nodes. Ensure that Compute 1 is added before Compute 2.
- If the setup is HA, ensure that you add Compute 2 before the other the computes (except for Compute 1).
-

