



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 13](#)
- [Application Framework User Interface, on page 17](#)
- [Catalog, on page 18](#)
- [Compute, on page 24](#)
- [Preferences, on page 25](#)
- [Failure Scenario, on page 26](#)

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 13](#).

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.

Refer to *Installing Cisco DCNM Compute Node* in the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment*.



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.

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 (the maximum qualified scale is 350 switches starting from Cisco DCNM 11.3(1) Release), 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 Fabric Deployment Without Network Insights (NI)



Note For information about various system requirements for proper functioning of Cisco DCNM LAN Fabric deployment, see .

Refer to *Network Insights User guide* for sizing information for Cisco DCNM LAN Deployment with Network Insights (NI).

To see the verified scale limits for Cisco DCNM 11.4(1) for managing LAN Fabric deployments, see .

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	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

Node	CPU Deployment Mode	CPU	Memory	Storage	Network
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 ¹

¹ 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.

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.

Refer to [Installing Cisco DCNM Compute Node](#) in the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment, Release 11.3(1)*.

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

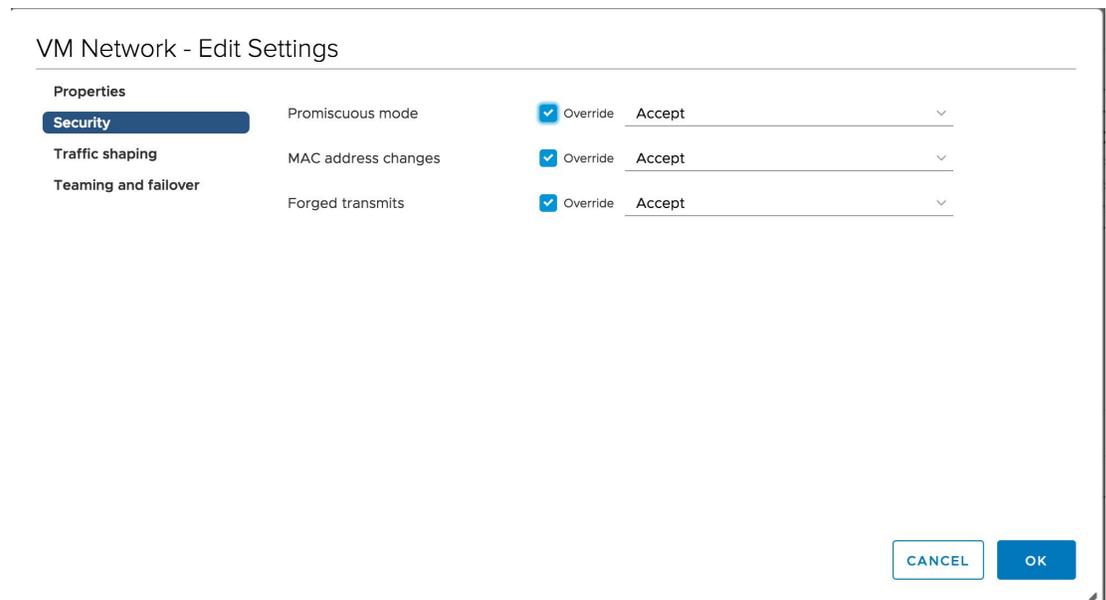


Figure 2: Security Settings for DVSwitch Port-Group

OobFabric - Edit Settings

General		
Advanced	Promiscuous mode	Accept
VLAN	MAC address changes	Accept
Security	Forged transmits	Accept
Teaming and failover		
Traffic shaping		
Monitoring		
Miscellaneous		

CANCEL OK



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

Enabling the Compute Cluster



Note Ensure that you enable Compute Cluster before you install applications. The applications that are installed via the AppCenter will not work if you enable the compute cluster after installing the applications.



Note The services are down until the configuration is complete. Ensure that the session is active while configuration is in progress.



Note If you enable clustered mode while installing Cisco DCNM, you don't need to enable cluster. The compute nodes will be discovered on Cisco DCNM **Web UI > Applications > Compute**. Go to [Compute, on page 24](#) to form a cluster.

If you did not enable clustered mode while installation, use the following command to enable the compute cluster.

```
appmgr afw config-cluster
[--ewpool<InterApp-Subnet>]--oobpool<OutOfBand-Subnet>--ibpool<Inband-Subnet>--computeip<compute-ip>
```

Where:

- **ewpool**: specifies the east-west pool subnet; for inter-service connectivity.

This field is optional, if the inter-application subnet is specified during Cisco DCNM installation for your deployment type. These addresses are not used directly between the computes, or to communicate with another node. These are used by containers to communicate with each other. This subnet must be minimum of /24 (256 addresses) and a maximum of a /20 (4096 addresses).

This field is optional if the Inter-app subnet is specified during Cisco DCNM deployment installation.

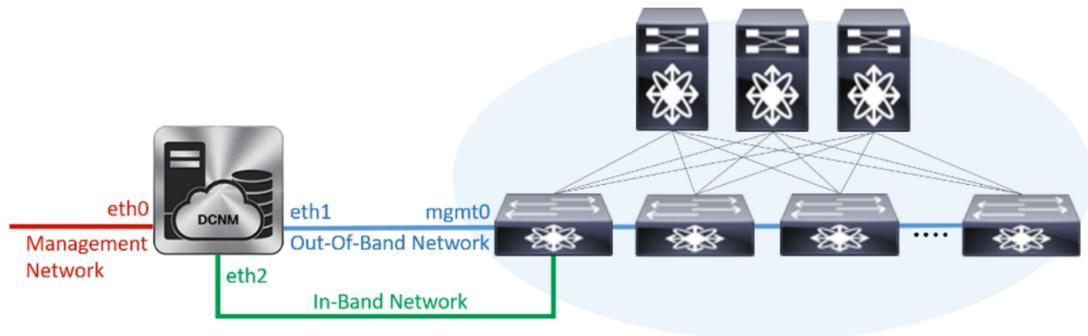
- **oobpool**: specifies the out-of-band pool; a smaller prefix of available IP addresses from the eth1 subnet. For example: Use 10.1.1.240/28 if the eth1 subnet was configured as 10.1.1.0/24 during installation.

This subnet must be a minimum of /28 (16 addresses) and maximum of /24 (256 addresses). It should also be longer than the east-west pool. This subnet is assigned to containers, to communicate with the switches.

- **ibpool**: specifies the in-band pool; a smaller prefix of available IP addresses eth2 subnet. For example: Use 11.1.1.240/28 if the eth2 subnet was configured as 11.1.1.0/24 during installation.

This subnet must be a minimum of /28 (16 addresses) and maximum of /24 (256 addresses). It should also be longer than the east-west pool. This subnet is assigned to containers, to communicate with the switches.

- **computeip**: specifies the dcnm-mgmt network (eth0) interface IP address of the first compute node added to the cluster. This compute is added into the cluster as part of this command process and is used to migrate application data from DCNM servers to computes.



Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime
<input type="radio"/> 172.28.12.205	eth2	eth1	Joined	80%	80%	-- 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			

The other two computes are Discovered automatically, and is displayed on the Cisco DCNM Web UI > **Applications > Compute**.

The In-Band or out-of-band pools are used by services to connect with switches as required. The IP addresses from these pools must be available for configuration.



Note To add computes to the cluster mode, see [Adding Computes into the Cluster Mode, on page 8](#).

Managing Application Network Pools

When you alter the eth1 or eth2 interface subnets, the corresponding oob pool and inband pool must be modified to match the new configuration. Network Insights and Endpoint Locator applications use the IP addresses from the Out-of-Band and In-Band pools.

To modify the IP addresses that are assigned to services running in the compute cluster, use the following command:



Note The inband or out-of-band pools are used by applications to connect with Cisco Nexus Switches. Hence, the IP addresses from these pools must be available and free.

```
appmgr afw config-pool [--ewpool <InterApp-Subnet>] --oobpool <OutOfBand-Subnet> --ibpool
<Inband-Subnet>--compute<compute-IP>
```

Where:

- **ewpool**: specifies the east west pool subnet; for inter-service connectivity.
The network mask ranges from 20 to 24. These addresses aren't used directly between the computes, or to communicate with another node. These are used by containers to communicate with each other.
- **oobpool**: specifies the out-of-band pool; a smaller prefix of available IP Addresses from eth1 subnet.
The network mask ranges from 24 to 28.
- **ibpool**: specifies the inband pool; a smaller prefix of available IP addresses from eth2 subnet.
The network mask ranges from 24 to 28.
- **ipv6oobpool**: specifies the out-of-band IPv6 pool; a smaller prefix of available IPv6 addresses from eth1 subnet.
If IPv6 is enabled, these addresses are required on both inband and out-of-band subnet.
The network mask ranges from 112 to 124.
- **ipv6ibpool**: specifies the inband IPv6 pool; a smaller prefix of available IPv6 addresses from eth2 subnet.
If IPv6 is enabled, these addresses are required on both inband and out-of-band subnet.
The network mask ranges from 112 to 124.

Adding Computes into the Cluster Mode

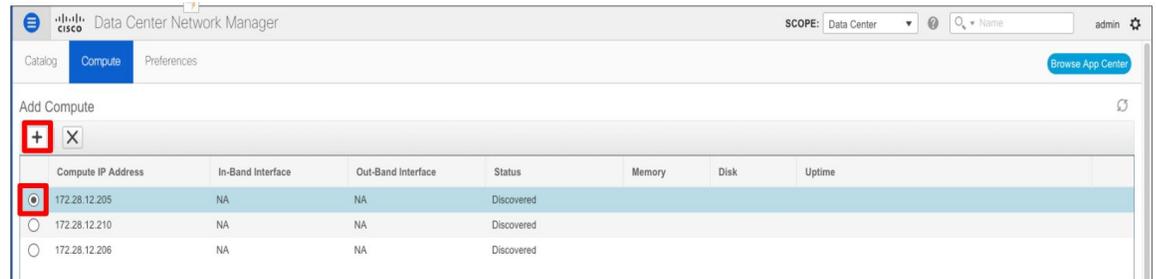
To add computes into the cluster mode from Cisco DCNM Web UI, perform the following steps:

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 Health Monitor application provides more detailed statistics.
- Cisco DCNM 3 node cluster is resilient to single node failure only.
- 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.

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.

Step 4 Click **OK**.

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

Add Compute							
Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime	
<input type="radio"/> 172.28.12.205	NA	NA	Joining				
<input type="radio"/> 172.28.12.210	NA	NA	Discovered				
<input type="radio"/> 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	62%	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	48%	99%	183 Hrs : 15 Min : 41 Sec	
<input type="radio"/> 172.28.12.210	eth2	eth1	Joined	97%	99%	-- Hrs : 4 Min : 9 Sec	
<input type="radio"/> 172.28.12.206	eth2	eth1	Joined	98%	99%	-- 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.

Transitioning Compute Nodes

Transitioning Compute nodes from VM to Service Engine

To transition Cisco DCNM Compute Nodes from VMs to Applications Services Engine using the Cisco DCNM Web Client, perform the following steps:

Before you begin

- Ensure that Cisco DCNM Web Client is functioning.
- On the Cisco DCNM **Web Client** > **Applications** > **Compute**, all the Compute nodes must be in **Joined** state.

Procedure

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

For example, let us indicate the three Compute nodes as **compute1** , **compute2** , and **compute3** .

- Step 2** Open the vCenter Server application and connect to the vCenter Server with your vCenter user credentials.
- Step 3** Navigate to **Home > Inventory > Hosts and Clusters** and identify the VM on which the DCNM Compute nodes are deployed.
- Step 4** For **compute1**, make a note of the configurations and setup details provided during installation.
- Step 5** Turn off **compute1**. Right click on the VM, select **Power off**.
On the **Web UI > Applications > Compute**, the status of **compute1** shows **Offline**.
- Step 6** Using the configuration details of the compute node VM, install the compute node on Cisco Applications Services Engine.
For instructions, refer to *Installing DCNM Compute Node on Cisco ASE*.
- Step 7** Launch the Web UI, and choose **Applications > Compute**.
The newly added compute automatically joins the cluster. The status of **compute1** changes from **Offline** → **Joining** → **Joined**.
- Step 8** Repeat Steps [Step 4, on page 11](#) to [Step 7, on page 11](#) on **compute2** and **compute3** compute nodes.
After completion, all the Compute nodes on **Web UI > Applications > Compute** are in the **Joined** state.
All are Compute nodes are successfully hosted on the Cisco Applications Services Engine.
-

Transitioning Compute nodes from Service Engine to VM

To transition Cisco DCNM Compute Nodes from Applications Services Engine to VMs using the Cisco DCNM Web Client, perform the following steps:

Before you begin

- Ensure that Cisco DCNM Web Client is functioning.
- On the Cisco DCNM **Web Client > Applications > Compute**, all the Compute nodes must be in **Joined** state.

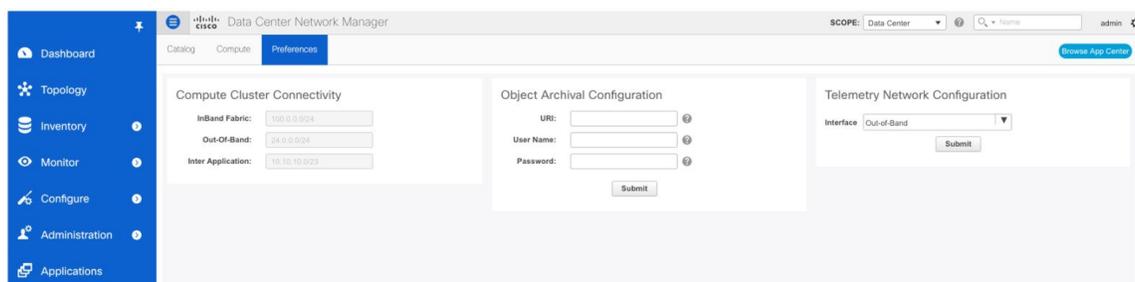
Procedure

- Step 1** Choose **Applications > Compute**.
For example, let us indicate the three Compute nodes as **compute1** , **compute2** , and **compute3** .
- Step 2** On the Cisco Applications Server console, for **compute1**, make a note of the configurations and setup details provided during installation.
- Step 3** Power off the Applications Service Engine to turn off **compute1**.
On the Cisco DCNM **Web UI > Applications > Compute**, the status of **compute1** shows **Offline**.
- Step 4** Using the configuration details of the compute node on Applications Service Engine, install the compute node on the VM.

- Step 5** Launch the Web UI, and choose **Applications > Compute**.
The newly added compute automatically joins the cluster. The status of **compute1** changes from **Offline** → **Joining** → **Joined**.
- Step 6** Repeat Steps 3 to 5 on **compute2** and **compute3** compute nodes.
After completion, all the Compute nodes on **Web UI > Applications > Compute** are in the **Joined** state.
All are Compute nodes are successfully hosted on the VMs.

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.



Compute Cluster Connectivity

The fields show the IP address that is used to configure the connectivity interfaces for the cluster node. The IP addresses for in-band fabric, out-of-band fabric, and Inter-Application are displayed.

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.

For the Cisco Network Insights for Resources (NIR) Release 2.1 and later, and flow telemetry, **feature lldp** command is one of the required configuration.

Cisco DCNM pushes **feature lldp** on the switches only for the Easy Fabric deployments, that is, for the eBGP routed fabric or VXLAN EVPN fabric. Therefore, NIR users need to enable **feature lldp** on all the switches in the following scenarios:

- External fabric in Monitored or Managed Mode

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
```

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

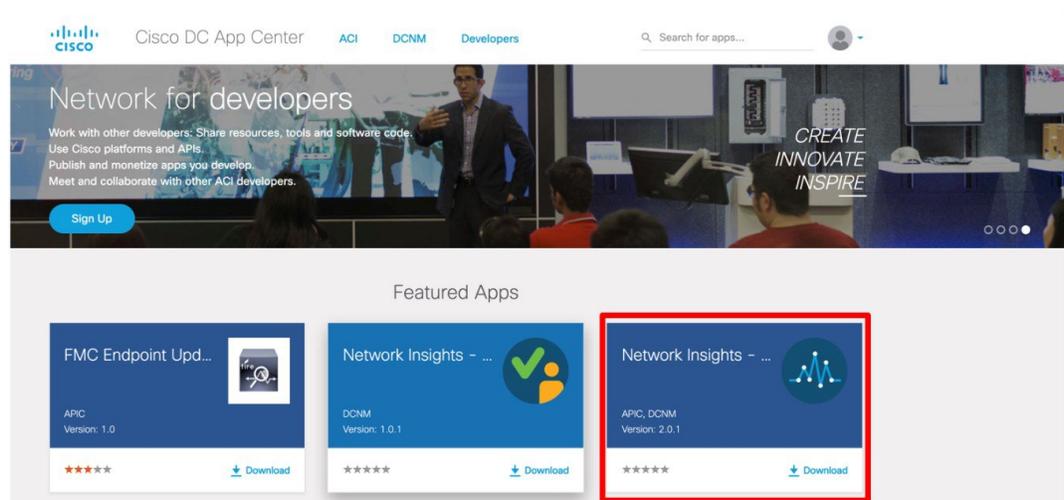
To download new applications from the Cisco DCNM Web UI, perform the following steps:

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.

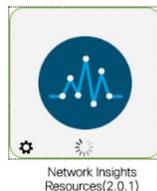


Note Ensure that the Compute Cluster is enabled before you install applications. A few applications may not work if the compute cluster is configured after installing the applications.

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.



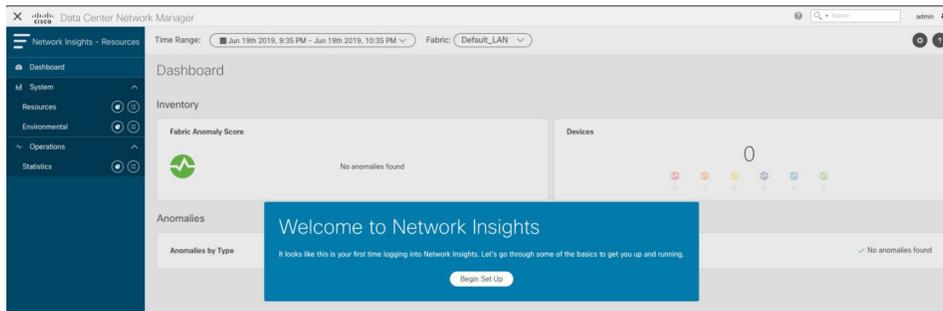
The green icon on the left-top corner indicates that the application is launched successfully and is operational.



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.

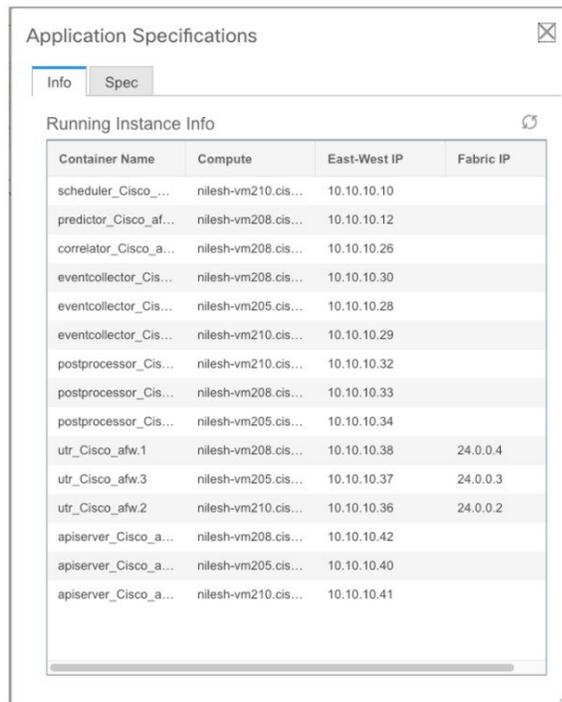
If the application has a user interface, after the application is successfully started the UI redirects to the index page served by the application.



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". The table lists various container instances and their associated compute nodes and IP addresses.

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 17](#).

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 24](#).



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 12](#).

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.
- **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.
- **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.

Catalog

The Catalog allows you to view all the applications that you have installed or enabled on the Cisco DCNM. Few applications are installed and are operational by default, when you install the Cisco DCNM.

The following applications appears based on the Cisco DCNM Deployments:

- Compliance (2.0)
- Debug plug-in (1.0)
- Kibana (1.1)
- Vmmplugin (3.0)
- Health Monitor (2.0)
- Endpoint Locator 2 (2.0)
- Kubernetes Visualizer(1.0)



Note The applications started by default, or also installed on the DCNM utilizes infrastructure services are operational, by default.

You can install more applications from the App Center, via the Web UI.

For instructions about downloading, adding, starting, stopping, and deleting applications from the Cisco DCNM Web UI, see [Installing and Deploying Applications, on page 13](#).

Health Monitor

The Health Monitor helps you to monitor the infrastructure health and status. You can monitor the Alerts, Service Utilization, and Compute Utilization using the Health Monitor application. When you install or upgrade to 11.2(1), the Health Monitor application is installed and operational, by default.

To launch the Health Monitor app, on the Cisco DCNM Web UI, choose **Applications**. On the Catalog tab, click on **Health Monitor** to launch the application.



Note Health Monitor application is installed by default in Cisco DCNM cluster mode.

Health Monitor app broadly monitors and alerts on the following metrics for Services, Computes and DCNM server:

- CPU utilization
- Memory utilization
- Network I/O (eth0)
- Disk I/O



Note Health Monitor metrics can also be viewed in iReasoning or any other MIB browser/tool.

You can monitor the following using the Health Monitor application:

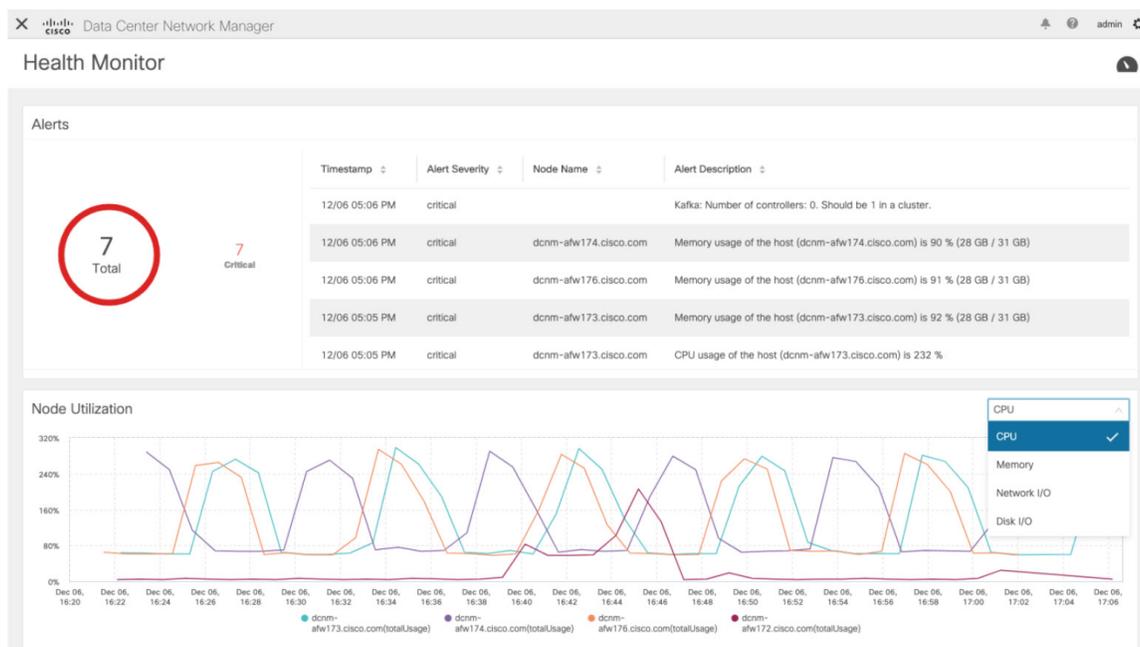
Alerts

The Alerts window provides information about the number of alerts that have occurred, from the specified date and time. You can view the alerts, based on the following categories, in the graphical view and the list view.

In the graphical view, the categories are:

- **Severity** displays the alerts, based on the severity: Critical/Major/Minor/Info.
- **Type** displays the alerts, based on the cluster type.
- **Compute** displays the alerts, for each compute node.
- **Service** displays the alerts, for all the services running on Cisco DCNM.

Click on the Refresh icon to refresh the alerts. Click on the list view icon to view the alerts in list format.



In the List View, alerts are displayed in tabular format with the following categories:

- **Timestamp** displays the time when the alert triggers. Format is MM/DD HH:MM AM/PM.
- **Alert Severity** displays the severity of alert.
- **Alert Type** displays the cluster alert type.
- **Node Name** displays the node name where the alert triggers.
- **Alert Description** displays the summary of the alert.

Click on the right or left navigation arrows to move to the next or the previous page.

You can also choose to set the number of items to view on page. Select a suitable number from the **Objects Per Page** drop-down list.

Click on the **Graphical representation** icon to go to the graphical view. Click on **Download Data** icon to download alerts information for troubleshooting purposes.

Health Monitor generates alerts for the following metrics:

- CPU utilization $\geq 65\%$
- Memory utilization $\geq 65\%$
- Disk utilization $\geq 65\%$
- Elasticsearch cluster status: Red/yellow
- Elasticsearch unassigned shards > 0
- Elasticsearch JVM heap used $\geq 65\%$
- Kafka partitions without leader: Controller offline partitions count > 0

- Kafka controllers count: Controller active controller count != 1
- Kafka partition leader: Controller unclear leader elections count > 0

Service Utilization

You can monitor all the services running on the Cisco DCNM on this window. Based on the time range and the service, the graphical view shows the CPU and Memory utilization for service. Click on the **Service Utilization** icon on the top-right corner to launch the CPU utilization graphical view.

From the **Time Range** drop-down list, choose the time range for which you want to view the utilization. You can select a specific time interval to view the metrics during that time interval. Click the fields showing the date and time to select the required date and time interval. You can also click the date on the calendar to set range. Click **Apply** to confirm the time range.

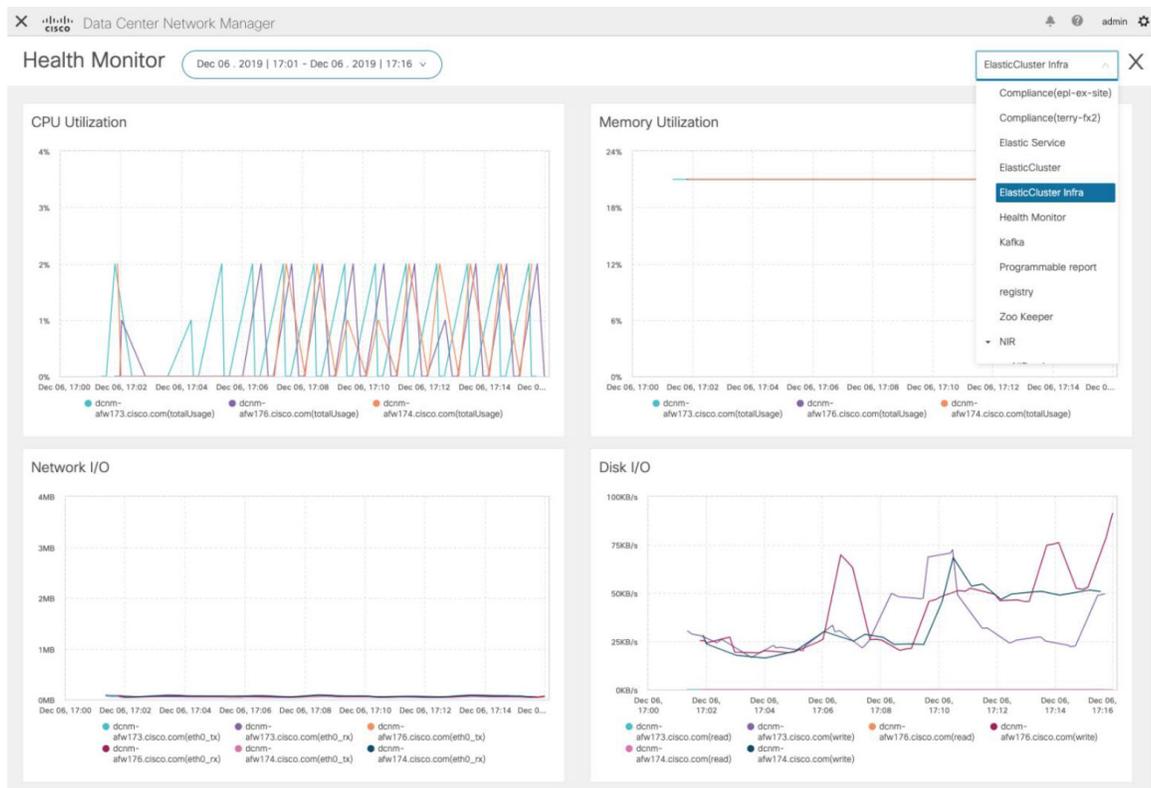
From the **Services** drop-down list, choose the service to view its Service utilization. This list comprises of all the services that are currently running on the Cisco DCNM.

Select the Time Range to view the **Service**, the **Cpu Utilization**, and **Memory Utilization** graphs. You can hover over specific points on the respective graphs for more information on CPU and Memory utilization at specific time.

The memory utilization graphical view depicts the actual memory consumption (RAM) in Gigabytes (GB). Click [**X**] icon on the top-right corner to close the Service Utilization window and revert to the Alerts window.

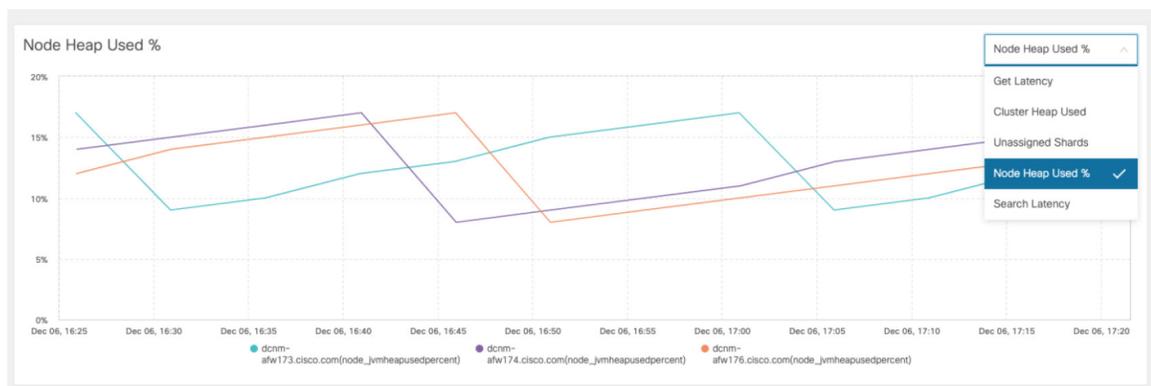
Guidelines and Limitations for Health Monitor in Service Utilization

- The CPU utilization for applications without a CPU limit, like Kafka, Elasticsearch, FMserver, and so on, may show 100% utilization in the graphs. 100% utilization is because this application uses one or more cores.
- The following alerts are triggered for the CPU utilization of applications:
 - Minor alert: 200-400 %
 - Major alert: 400-600%
 - Critical: > 600%
- The transient message for Kafka controller counts appears as a severe alert sometimes. You can ignore the alert if it clears within two minutes after refresh.
- The **Disk I/O** and **Memory Utilization** metrics are not available for Kafka and Elastic Service.
- The **Network I/O** metric is not available for **DCNM: FMServer** and **DCNM: Postgres**.
- The metrics does not auto-refresh. Navigate between different windows using the options in the drop-down list to refresh the metrics. Additionally, you can change the time range to refresh the metrics for a selected period.
- There might be duplicate alerts for the same feature.



The following additional metrics are collected for Elastic Cluster:

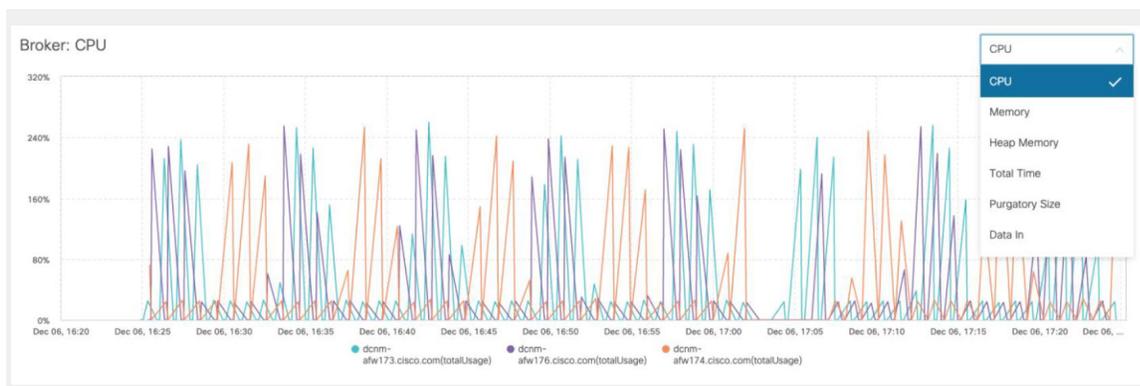
- Get latency: Latency for getting a single record by id
- Cluster heap used: Heap memory used by the cluster
- Unassigned shards: Count of unassigned shards
- Node heap used percentage: Percentage heap memory used by the node
- Search latency: Latency for getting a collection of records



The following additional metrics are collected for Kafka broker:

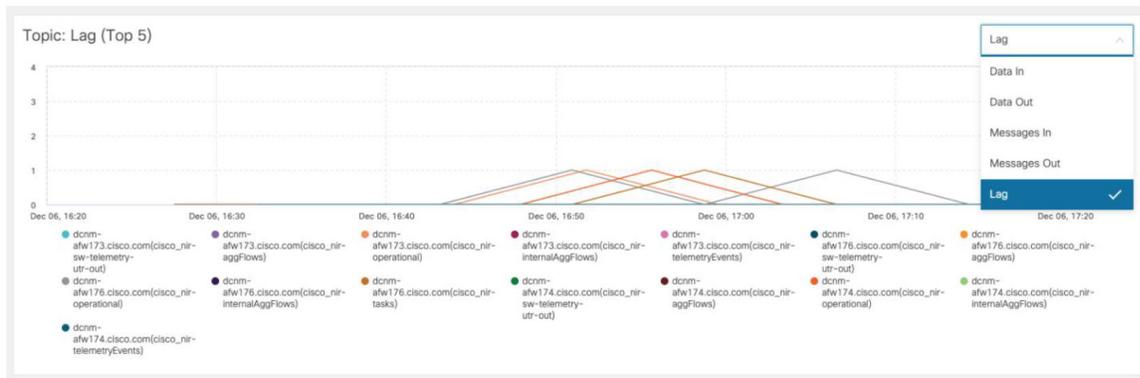
- CPU: CPU utilization of broker

- Memory: Memory utilization of broker
- Heap memory: Heap memory utilized by broker
- Total time: Network produce, network fetch follower, network fetch consumer time
- Purgatory size: Server fetch purgatory size, server produce purgatory size of broker
- Data in: Bytes in for the broker
- Data out: Bytes out for the broker
- Messages in: Messages received by the broker
- Fetch request: Total fetch requests for the broker
- ISR: In-sync-replicas expands and shrinks for the broker



The following additional metrics are collected for top 5 Kafka topics:

- Data in: Bytes in for the topic
- Data out: Bytes out for the topic
- Messages in: Message in count for topic
- Messages out: Message out count for topic
- Lag: Lag per topic



Compute Utilization

You can monitor all the computes installed with the Cisco DCNM. Based on the time range and the service, the graphical view shows the CPU and Memory utilization for service. Click on the **Compute Utilization** icon on the top-right corner to launch the CPU utilization graphical view.

From the **Time Range** drop-down list, choose the time range for which you want to view the utilization. You can select a specific time interval to view the metrics during that time interval. Click the fields showing the date and time to select the required date and time interval. You can also click the date on the calendar to set range. Click **Apply** to confirm the time range.

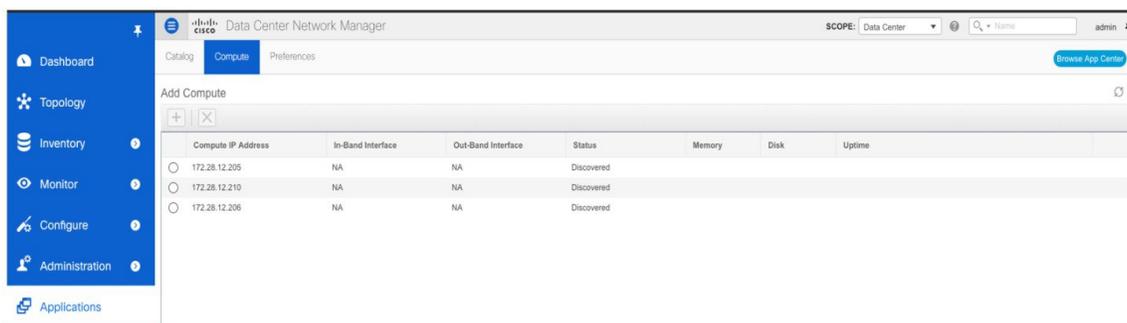
Select the Time Range to view the **Service**, the **Cpu Utilization**, and **Memory Utilization** graphs. You can hover over specific points on the respective graphs for more information on CPU and Memory utilization at specific time.

The memory utilization graphical view depicts the actual memory consumption (RAM) in Gigabytes (GB).

Click [**X**] icon on the top-right corner to close the Service Utilization window and revert to the Alerts 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"> • Joined • Discovered • Failed • Offline
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 8](#).

To configure or modify the Cluster Connectivity preferences, see [Preferences, on page 12](#).

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.

The screenshot shows the Cisco Data Center Network Manager interface. The left sidebar contains navigation options: Dashboard, Topology, Inventory, Monitor, Configure, Administration, and Applications. The main content area is titled 'Compute Cluster Connectivity' and includes three configuration panels:

- Compute Cluster Connectivity:** Contains three input fields for IP addresses: InBand Fabric (100.0.0.0/24), Out-Of-Band (24.0.0.0/24), and Inter Application (10.10.10.0/23).
- Object Archival Configuration:** Contains three input fields: URI, User Name, and Password, each with a help icon.
- Telemetry Network Configuration:** Contains a dropdown menu for 'Interface' set to 'Out-of-Band' and a 'Submit' button.

Compute Cluster Connectivity

The fields show the IP address that is used to configure the connectivity interfaces for the cluster node. The IP addresses for in-band fabric, out-of-band fabric, and Inter-Application are displayed.

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.

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.