



Deploying Applications in Unclustered Mode



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 compute cluster after installing the applications.

- In a Cisco DCNM Native HA setup, ensure both active and standby nodes are synchronized, i.e. HA status is OK.
- Apply the vSwitch networking policies, see [Networking Policies for OVA Installation - Unclustered mode](#).
- To upload an application, see [Download app from the App Store](#). To start, stop, or delete applications, see [Starting Application](#) or [Stop and Delete Applications](#).
- [Deploying Applications in Clustered Mode, on page 1](#)
- [Application Framework User Interface, on page 8](#)
- [Replication of Containers for DCNM Services, on page 13](#)
- [Disaster Recovery, on page 14](#)

Deploying Applications in Clustered Mode

Refer to the following sections for deploying applications in cluster mode:

- Requirements for a clustered setup, see [Requirements for Cisco DCNM Clustered Mode, on page 2](#).
- To install a DCNM compute, see [Installing a DCNM Compute, on page 4](#).
- To ensure that the vSwitch Networking Policies are properly configured, see [Networking Policies for OVA Installation - Clustered mode, on page 4](#).
- To form the compute cluster, see [Enabling the Compute Cluster](#).
- Once cluster is formed, the compute nodes will appear as Discovered in the Applications > Compute tab. Then compute Nodes need to be added by an user with admin privileges on the DCNM. To add them compute nodes from the UI, see [Adding Computes into the Cluster Mode, on page 6](#).
- Now compute cluster is ready for applications to be deployed. To upload an application, see [Add a new application to DCNM](#). To start/stop/delete application, see [Stop and Delete Applications](#).

For additional compute/application related operations, see [Application Framework User Interface](#), on page 8.

For advanced compute and application health monitoring, see [Watch Tower](#).

Cisco DCNM Cluster Mode

Starting from Cisco DCNM 11.1, in a DCNM HA setup (active + standby), 80 switches with Endpoint Locator, Virtual Machine Manager, config compliance are validated. For a network exceeding 80 switches, with these features in a given DCNM instance, (maximum qualified scale is 256 switches), it is recommended to add 3 compute nodes. This is called the clustered mode in Cisco DCNM.

Compute nodes are scale out application hosting nodes that run resource-intensive services to provide services to larger fabric. When compute nodes are added, all containerized services run only on these nodes. This includes Config Compliance, Endpoint Locator, Virtual Machine Manager. The Elasticsearch time series database for these features run on compute nodes in case of clustered mode.

While DCNM core functionalities only run on the DCNM HA nodes. Addition of compute nodes beyond 80 switches is to build a scale out model for 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.

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 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 cluster mode have additional requirements. That is, 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. This IP address needs to remain sticky during all failure scenarios. For this purpose, an IP pool needs to 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 need to be time synchronized (NTP is recommended). 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 always time synchronized. To setup telemetry network configuration, see .

How Do the Compute Nodes Communicate with DCNM Nodes?

The compute nodes communicate with DCNM nodes using APIs.

How Do the Compute Nodes Communicate with Each Other?

The compute nodes communicate with each other using Eth0 IP addresses. When in clustered mode, it is recommended to have at least 16 IP addresses in the pool for the containers.

Installing a 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 required to provision and monitor applications.

If you have a Cisco DCNM LAN Fabric deployment, refer to [Installing Cisco DCNM Compute Node](#) in the *Cisco DCNM Installation and Upgrade Guide for Classic LAN Deployment, Release 11.2(1)*.

If you have a Cisco DCNM LAN Fabric deployment, refer to [Installing Cisco DCNM Compute Node](#) in the *Cisco DCNM Installation and Upgrade Guide for LAN Fabric Deployment, Release 11.2(1)*.



Note Compute nodes and Cluster modes are supported only on these two Cisco DCNM Deployments.

Networking Policies for OVA Installation - Clustered mode

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

- Login to vCenter.
- Click on the Host where compute 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** window 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

VM Network - Edit Settings

Properties			
Security	Promiscuous mode	<input checked="" type="checkbox"/> Override	Accept
Traffic shaping	MAC address changes	<input checked="" type="checkbox"/> Override	Accept
Teaming and failover	Forged transmits	<input checked="" type="checkbox"/> Override	Accept

Figure 2: Security settings for DVSwitch Port-group

OobFabric - Edit Settings

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

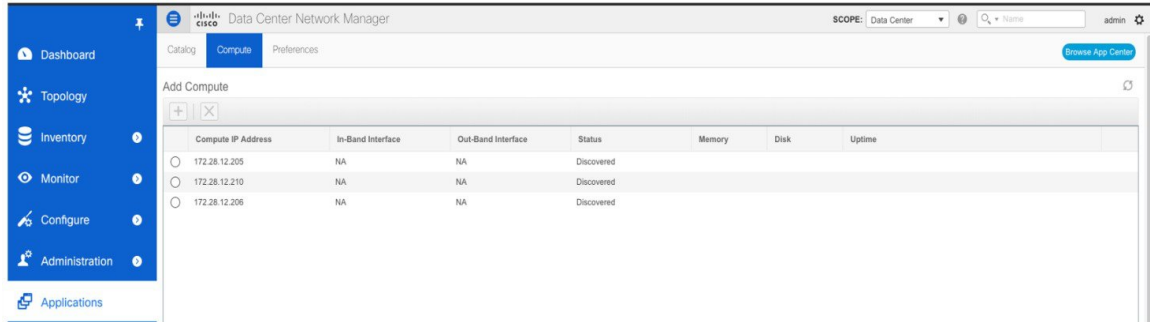


Note Ensure that you repeat this procedure on all the hosts, where a Compute OVA is running.

Adding Computes into the Cluster Mode

Compute is an additional installation mode with Cisco DCNM Release 11.2(1). It is supported with both small and large installations. Cisco DCNM supports a maximum of three Computes.

When a Compute is installed with correct parameters, it appears as **Joined** in the Status column. However, the other two computes will appear as **Discovered**.



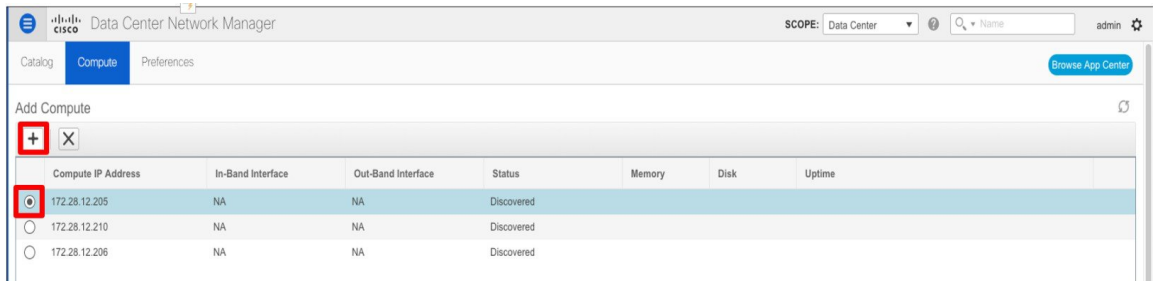
To add computes into the cluster mode from Cisco DCNM Web UI, perform the following steps:

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 that some connectivity issues between Nodes (on Eth0).
- **Failed** indicates that the compute node failed to join the cluster.
- **Health only** indicates the amount of Memory left (Just like UnClustered mode).

The Compute window can also be used to monitor the health of computes. The health essentially indicates the amount of memory 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 will be 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, verify the **Compute IP Address**, **In-Band Interface**, and the **Out-Band Interface** values.

Step 4 Click **OK**.

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

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			

You must wait until the Compute IP status shows **Joined**.

Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime
172.28.12.205	eth2	eth1	Joined	60%	99%	-- Hrs : 4 Min : 17 Sec
172.28.12.210	NA	NA	Discovered			
172.28.12.206	NA	NA	Discovered			

- Note**
- Offline indicates some connectivity issues, therefore no applications will be running on Offline Computes.
 - Failed indicates that the compute node failed to join the cluster.
 - Health indicates the amount of free memory and disk on the Compute node. The Watchtower application provides more detailed statistics.
 - 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.
 - 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.

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

All the Computes appear as

Compute IP Address	In-Band Interface	Out-Band Interface	Status	Memory	Disk	Uptime
172.28.12.205	eth2	eth1	Joined	48%	99%	183 Hrs : 15 Min : 41 Sec
172.28.12.210	eth2	eth1	Joined	51%	99%	-- Hrs : 4 Min : 9 Sec
172.28.12.206	eth2	eth1	Joined	55%	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 should allow for packets associated with Mac address other than eth1 and eth2 to be forwarded.

Telemetry Network and NTP Requirements

For the Network Insights Resource (NIR) application, a UTR micro-service 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 to be 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 DCCNM LAN Fabric Deployment, or switch-groups in Cisco DCNM Classic LAN Deployment. After the telemetry is enabled via Network Insights Resources (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, as shown below:

```
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 the 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.

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 9](#).



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 13](#).

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 a fabric is created. Similarly, it is disabled when a 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.

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 a cluster mode, the compute nodes are also displayed along with the indication of whether they are joined or discovered.

The screenshot shows the Cisco DCNM interface with the 'Compute' tab selected. The table below displays the data for discovered compute nodes.

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 in sync 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 the Compute nodes are configured using a different NTP server, the mismatch in timestamp will not allow the certificates to be validated. Therefore, if the compute cluster is configured despite of a mismatch of NTP server, the applications will not function properly.



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

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

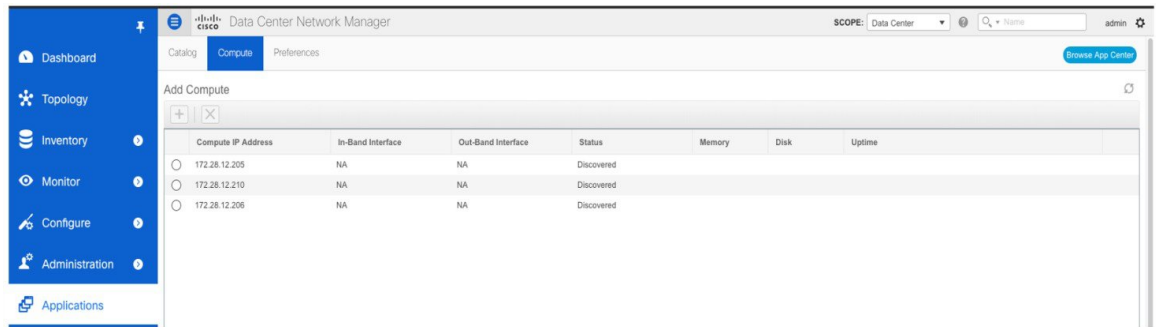
Table 5: Field and Description on Compute Tab

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 consumed for that node.
Disk	Specifies the disk space consumed on the compute node.
Uptime	Specifies the duration of the uptime for compute node.

Adding Computes into the Cluster Mode

Compute is an additional installation mode with Cisco DCNM Release 11.2(1). It is supported with both small and large installations. Cisco DCNM supports a maximum of three Computes.

When a Compute is installed with correct parameters, it appears as **Joined** in the Status column. However, the other two computes will appear as **Discovered**.



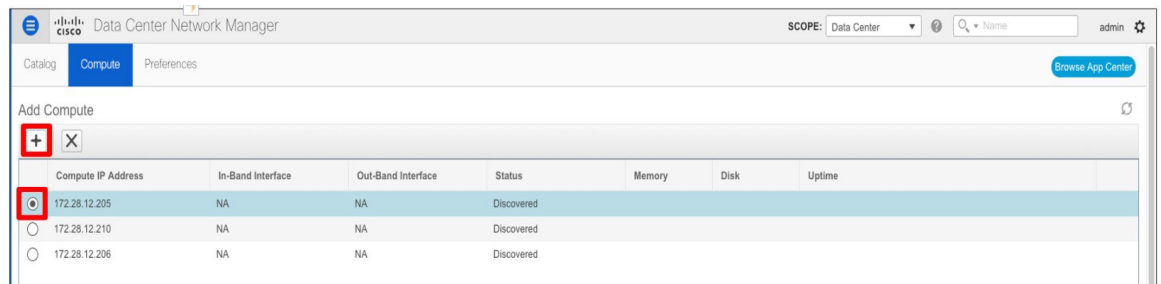
To add computes into the cluster mode from Cisco DCNM Web UI, perform the following steps:

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 that some connectivity issues between Nodes (on Eth0).
- Failed indicates that the compute node failed to join the cluster.
- Health only indicates the amount of Memory left (Just like UnClustered mode).

The Compute window can also be used to monitor the health of computes. The health essentially indicates the amount of memory 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 will be 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, verify the **Compute IP Address**, **In-Band Interface**, and the **Out-Band Interface** values.

Add Compute ✕

Compute IP Address

In-Band Interface

Out-Band Interface

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				

You must 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				

Note

- Offline indicates some connectivity issues, therefore no applications will be running on Offline Computes.
- Failed indicates that the compute node failed to join the cluster.
- Health indicates the amount of free memory and disk on the Compute node. The Watchtower application provides more detailed statistics.
- 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.
- 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.

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

All the Computes appear as

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%	99%	→ Hrs : 4 Min : 9 Sec	
<input type="radio"/> 172.28.12.206	eth2	eth1	Joined	55%	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 should allow for packets 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.

Compute Cluster Connectivity

The fields show the IP address that are 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 thereby, determines the advisory based on the data. The logs can be 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 user name and password to access the URI, in the **User Name** and **Password** field. Click **Submit** to configure the remote server.

Replication of Containers for DCNM Services

The replication and deployment of the containers is dependent on how the services are configured. There are three modes.

- Applications that do not require replication (single container applications, such as config compliance or Endpoint Locator of LAN fabric installation). If a compute node that is hosting this container is lost, a new container is brought up on the available computes.
- Applications that get one container per compute. If a new compute is added, automatic creation of container for the new compute happens on runtime (such as Elasticsearch infra application). If a compute node is lost, the corresponding container is lost as well and is unavailable until the compute comes back.

- Applications that require their containers be replicated based on a certain count (defined by the application). These containers are generally spread evenly across the computes. When a compute goes down, new instances are created to match the required count and spread evenly on the rest of the computes. In Cisco DCNM Release 11.1(1), when the computes come back, they are not rebalanced. In Cisco Release 11.2(1), they are rebalanced if the number of computes are 3.



Note Docker default algorithm tries to round robin spinning up of containers for the services. All computes will be of uniform resource.

To replicate the containers on the application on the Cisco DCNM Web UI, perform the following steps:

1. Click **Applications**.

All the applications installed on the Cisco DCNM deployment is displayed under the **Catalog** tab.

2. Move the cursor on the left-top corner of the application that needs replication of containers.

A + icon appears indicating that you can replicate the containers on the application.

3. Click **Replicate** to add containers to the application.

Click **Cancel** to abort.

Disaster Recovery

The **appmgr backup** operation on a compute gathers all the data that is required to re-install the compute. Also, this operation preserves all the application data. Using the tar ball generated by the backup command, the **appmgr restore** command restores all the data into the compute. This is similar to how you restore Cisco DCNM from a backup data.

When a compute needs to be reinstalled in a disaster recovery mode, you need to restore the application data into a new installation. It is also possible that the Cisco DCNM servers need to be restored into a new server. You may find the following scenarios:

- Cisco DCNM Controllers need to be recovered
- Cisco DCNM Computes need to be recovered
- Cisco DCNM Controllers and Computes need to be recovered

Scenario 1

You can use SSH to access the computes as root, and enter the **appmgr stop afw** command on each of the compute nodes. Now, the DCNM controllers can be powered off and restored onto a new DCNM Installation.

After the restore of the DCNM controllers is complete, verify that DCNM controller is completely up and the Applications screen is loading. Verify that all computes are showing up as offline. Now, enter the **appmgr start afw** command on each of the computes. After a while ensure all the applications are running and Computes are showing as “Joined”.

Scenario 2

In this case, enter the **appmgr stop afw** command on the compute that is being restored, after the compute shows offline in the Compute tab. Restore the compute on a new installation.

Perform one restore after the other.

Scenario 3

In this case, first perform scenario 1, and then perform scenario 2.

Failure Scenario

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

- DCNM Active Node(Active) and compute node 1 in server 1
- DCNM Standby Node and compute node 2 in server 2
- Compute node 3 in server 3

When one DCNM node is lost, the standby DCNM node takes full responsibility of running the core functionality.

Applications may continue to function at loss of one compute node, sometimes with limited functionality. If this situation persists for a longer duration, performance of the applications will be affected. When more than 1 node is lost, services which write data to Elasticsearch will be affected until the 2 nodes are brought back up. For example, Virtual Machine Manager, Endpoint Locator etc. All of the 250 switches' Config compliance will run on a single compute and hence may notice relatively low performance.

Hence, you need to maintain 3 compute nodes at any time, and at an event any of these are lost, they must be brought up for the services to function as expected.

