



Managing Applications in a High-Availability Environment

This chapter describes how to configure a high-availability (HA) environment in your Cisco DCNM Open Virtual Appliance deployment for your Cisco Programmable Fabric solution. It also includes details about the HA functionality for each of the applications bundled within the Cisco DCNM Open Virtual Appliance.



Note Ensure that the NTP server is synchronized between active and standby peers is essential for proper HA functioning in DCNM

This chapter contains the following sections:

- [Information About Application Level HA in the Cisco DCNM Open Virtual Appliance, on page 1](#)
- [Native HA Failover and Troubleshooting, on page 2](#)
- [Application High Availability Details, on page 4](#)

Information About Application Level HA in the Cisco DCNM Open Virtual Appliance

To achieve HA for applications that are run on the Cisco DCNM Open Virtual Appliance, you can run two virtual appliances. You can run one in Active mode and the other in Standby mode.



Note This document refers to these appliances as OVA-A and OVA-B, respectively.

In this scenario:

1. All applications run on both appliances.
The application data is either constantly synchronized or applications share a common database as applicable.
2. Only one of the applications running on the two appliances serves the client requests. Initially this would be the applications running on OVA-A. The application continues to do so until one of the following happens:

- The application on OVA-A crashes.
 - The operating system on OVA-A crashes.
 - OVA-A is powered off for some reason.
3. At this point, the application running on the other appliance (OVA-B) takes over.
For DHCP, when the first node fails, the second node starts serving the IP addresses.
 4. The existing connections to OVA-A are dropped and the new connections are routed to OVA-B.
This scenario demonstrates why one of the nodes (OVA-A) is initially referred to as the Active node and OVA-B is referred as the Standby node.

Automatic Failover

The application-level and virtual machine (VM)-level and switchover process is as follows.

- If any of the applications managed by the load-balancing software (DCNM/AMQP) goes down on OVA-A, the Active node that handles the client requests detects the failure and redirects subsequent requests to the Standby node (OVA-B). This process provides an application-level switchover.
- If the Active node (OVA-A) fails or is powered-off for some reason, the Standby node (OVA-B) detects the failure and enables the VIP address for Cisco DCNM/AMQP on OVA-B. It also sends a gratuitous ARP to the local switch to indicate the new MAC address that is associated with the IP address. For applications not using VIP, the DHCPD running on OVA-B detects the failure of DHCPD on OVA-A and activates itself; whereas LDAP running on OVA-B continues running as LDAP is deployed Active-Active. Consequently, a VM-level failover is accomplished for all four applications (DCNM/AMQP/DHCP/LDAP).

Manually Triggered Failovers

An application-level failover can also be triggered manually. For instance, you might want to run AMQP on OVA-B and the rest of the applications on OVA-A. In that case, you can log in to the SSH terminal of OVA-A and stop AMQP by using the **appmgr stop amqp** command.

This failover triggers the same process that is described in the [Automatic Failover, on page 2](#); subsequent requests to the AMQP Virtual IP address are redirected to OVA-B.

Native HA Failover and Troubleshooting

Due to the nature of Native HA, the role of the host might alternate from Active to Standby or from Standby to Active.

The following sections provide information on troubleshooting in different use cases.

Native HA Failover from Active Host to Standby Host

Perform the following steps when the Native HA failover occurs from Active to Standby host:

1. Log on to DCNM Web UI, and navigate to **Administrator > Native HA**.

2. Verify the status of HA. If the DCNM HA status is not in **OK** mode, you cannot perform Failover operation. Click **Failover**. The Cisco DCNM server will shutdown and the DCNM Standby appliance will be operational.
3. Refresh the Cisco DCNM Web UI.
After the DCNM server is operational, you can log on to the DCNM Web UI.



Note We recommend that you do not run **appmgr stop all** or **appmgr stop ha-apps** commands on the Active host to trigger failover. If Cisco DCNM HA status is not in **OK** mode, a failover may cause loss of data, as the Standby DCNM appliance is not synchronized with the Active appliance before failover.

Issue with DCNM Application Framework

If DCNM Web UI is not accessible, and a failover operation is necessary, execute one of the following commands under Linux console:

appmgr failover—This command triggers the HA heartbeat failover.

Or

reboot -h now—This command triggers the Linux host to reboot, which causes a failover.

However, we recommend that you use DCNM Web UI to perform failover, as all other methods carry a risk of data loss when both HA peers are not in sync.

Stop and Restart DCNM

To completely stop DCNM and restart it, perform the following:

1. On the Standby appliance, stop all the applications by using the **appmgr stop all** command.
2. Check if all the applications have stopped, using the **appmgr status all** command.
3. On the Active appliance, stop all the applications using the **appmgr stop all** command.
4. Verify if all the applications are stopped using the **appmgr status all** command.
5. On the deployed Active host, start all the applications using the **appmgr start all** command.
Verify if all the applications are running. Log on to the DCNM Web UI to check if it is operational.
6. On the deployed Standby host, start all the applications using the **appmgr start all** command.
On the Web UI, navigate to **Administration > Native HA** and ensure that the HA status displays **OK**.

Restart Standby Host

Perform this procedure to restart only the Standby host:

1. On the Standby host, stop all the applications using the **appmgr stop all** command.
2. Verify if all the applications have stopped using the **appmgr status all** command.
3. Start all the applications using the **appmgr start all**.

On the Web UI, navigate to **Administration > Native HA** and ensure that the HA status displays **OK**.

Application High Availability Details

This section describes all of the Cisco Programmable Fabric HA applications.

Cisco DCNM Open Virtual Appliance has two interfaces: one that connects to the Open Virtual Appliance management network and one that connects to the enhanced Programmable Fabric network. Virtual IP addresses are defined for both interfaces.

- From the Open Virtual Appliance management network, the DCNM-REST API, DCNM interface, and AMQP are accessed through the VIP address
- From the enhanced fabric management network, LDAP and DHCP are accessed directly.

Only three Virtual IPs are defined:

- DCNM REST API (on dcnm management network)
- DCNM REST API (on enhanced fabric management network)
- AMQP (on dcnm management network)



Note Although DCNM Open Virtual Appliance in HA sets up a VIP, the VIP is intended to be used for the access of DCNM, REST API. For GUI access, we still recommend that you use the individual IP addresses of the DCNM HA peers and use the same to launch DCNM SAN Java clients, etc.

See the following table for a complete list of Programmable Fabric applications and their corresponding HA mechanisms.

Programmable Fabric Application	HA Mechanism	Use of Virtual IPs	Comments
Data Center Network Manager	DCNM Clustering/Federation	Yes	Two VIPs defined, one on each network
RabbitMQ	RabbitMQ Mirrored Queues	Yes	One VIP defined on the OVA management network
Repositories	—	—	External repositories have to be used

Data Center Network Management

The data center network management function is provided by the Cisco Data Center Network Manager (DCNM) server. Cisco DCNM provides the setup, visualization, management, and monitoring of the data center infrastructure. Cisco DCNM can be accessed from your browser at [http://\[host/ip\]](http://[host/ip]).



Note For more information about Cisco DCNM, see <http://cisco.com/go/dcnm>.

HA Implementation

Cisco DCNMs that run on both OVAs are configured in clustering and federated modes for HA. Cisco DCNM federation is the HA mechanism for SAN devices. Groups of SAN devices can be managed by each node in the DCNM federated setup. All the devices can be managed using a single client interface.

You can enable automatic failover in the Cisco DCNM UI by choosing: **Admin > Federation**. If you enable an automatic failover and the Cisco DCNM that is running on OVA-A fails, the automatic failover moves only the fabrics and shallow-discovered LANs that are managed by OVA-A to OVA-B automatically.

DCNM Virtual IP Usage

An Open Virtual Appliance HA setup has two VIP addresses (one for each network) for the Cisco DCNM at the default HTTP port. These VIPs can be used for accessing the DCNM RESTful services on the Open Virtual Appliance management network and the enhanced fabric management network. For example, external systems such as Cisco UCS Director can point to the VIP in the Open Virtual Appliance management network and the request gets directed to the active Cisco DCNM. Similarly, the switches in an enhanced fabric management network access the VIP address on the enhanced fabric management network during the POAP process.

You can still directly connect to Cisco DCNM real IP addresses and use them as you would in a DCNM in a cluster/federated set up.



Note Cisco recommends that you must use VIP addresses only for accessing DCNM REST API. To access the Cisco DCNM Web or SAN client, you must connect using the IP address of the server.

Licenses

For Cisco DCNM, we recommend that you have licenses on the first instance and a spare matching license on the second instance.

Application Failovers

Enable an automatic failover option in the Cisco DCNM UI when an Open Virtual Appliance HA pair is set up by choosing: **Administration > DCNM Server > Native HA**. This process ensures that if the DCNM that is running on OVA-A fails, all the fabrics and shallow-discovered LANs managed by DCNM-A are managed by DCNM-B automatically after a given time interval (usually about 5 minutes after the failure of DCNM on OVA-A).

The Cisco DCNM VIP address still resides on OVA-A. The Representational State Transfer Web Services (REST) calls initially hit the VIP addresses on OVA-A and get redirected to the Cisco DCNM that is running on OVA-B.

Application Failbacks

When the Cisco DCNM on OVA-A comes up, the VIP address automatically redirects the REST requests to DCNM-A.

Virtual-IP Failovers

The VIP address that is configured for Cisco DCNM REST API on OVA-A can fail due to two reasons:

- The load-balancing software running on OVA-A fails.
- OVA-A fails.

The VIP address of Cisco DCNM automatically migrates to OVA-B. The only difference is which DCNM will be used after the failover.

- If a load-balancing software failure occurs, the VIP address on OVA-B directs the requests to DCNM-A.
- If an OVA-A failure occurs, the VIP address on OVA-B directs the requests to DCNM-B.

The automatic failover ensures that the ownership of all of the fabrics and shallow-discovered LANs managed by DCNM-A automatically change to DCNM-B.

Virtual-IP Failbacks

When OVA-A is brought up and Cisco DCNM is running, the VIP addresses keep running on the Standby node. The failback of Virtual IP addresses from OVA-B to OVA-A occurs only in the following sequence.

1. OVA-A comes up.
2. Cisco DCNM runs on OVA-A.
3. OVA-B goes down or the load-balancing software fails on OVA-B.

RabbitMQ

RabbitMQ is the message broker that provides the Advanced Messaging Queuing Protocol (AMQP).



Note You need to stop and restart AMQP on both DCNM's server in HA within 30 seconds, otherwise AMQP may not start. For more information about RabbitMQ, go to <https://www.rabbitmq.com/documentation.html>.

HA Implementation

Enabling the HA on the Open Virtual Appliance creates a VIP address in the Open Virtual Appliance management network. Orchestration systems such as vCloud Director, set their AMQP broker to the VIP address.

Enabling the HA on the Open Virtual Appliance also configures the RabbitMQ broker that runs on each node to be a duplicate of the broker that is running on the other node. Both OVAs act as “disk nodes” of a RabbitMQ cluster, which means that all the persistent messages stored in durable queues are replicated. The RabbitMQ policy ensures that all the queues are automatically replicated to all the nodes.

Application Failovers

If RabbitMQ-A fails, the VIP address on OVA-A redirects the subsequent AMQP requests to RabbitMQ-B.

Application Failbacks

When RabbitMQ-A comes up, the VIP address automatically starts directing the AMQP requests to RabbitMQ-A.

Virtual-IP Failovers

The VIP address configured for the AMQP broker on OVA-A can fail due to two reasons:

- The load-balancing software running on OVA-A fails.
- OVA-A fails.

In both cases, the VIP address of the AMQP automatically migrates to OVA-B. The only difference is which AMQP broker will be used after the failover.

- In a load-balancing software failure, the VIP address on OVA-B directs the requests to RabbitMQ-A.
- In an OVA-A failure, the VIP address on OVA-B directs the requests to RabbitMQ-B.

Virtual-IP Failbacks

When OVA-A is brought up and AMQP-A is running, the VIP addresses keep running on the OVA-B (directing the requests to AMQP-A). The failback of the RabbitMQ VIP from OVA-B to OVA-A occurs only in the following sequence.

1. OVA-A comes up.
2. RabbitMQ runs on OVA-A.
3. OVA-B goes down or the load-balancing software fails on OVA-B.

Repositories

All repositories must be remote.

