CHAPTER 7

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 includes the following sections:

Note For instruction about installing these applications with the Cisco DCNM Open Virtual Appliance, see the "DCNM Open Virtual Appliance Installation in Programmable Fabric mode" section on page 3-2.

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 DCNM REST API and AMQP, this transition is done by a load-balancing software that hides the interface address of the appliances using a Virtual IP (VIP) address.

   For LDAP, both nodes are configured as duplicates of each other. The LDAP clients (switches) are configured with primary and secondary LDAP IPs, so if the active LDAP fails they try contacting the LDAP running on the standby.

   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" section on page 7-2; subsequent requests to the AMQP Virtual IP address are redirected to OVA-B.

Prerequisites for Cisco DCNM Open Virtual Appliance HA

This section contains the following topics that describe the prerequisites for obtaining a high-availability (HA) environment.

- Oracle Database for DCNM Servers
- Deploying Cisco DCNM OVAs
- Creating an NFS/SCP Repository
- Availability of Virtual IP Addresses
- Installing an NTP Server

Deploying Cisco DCNM OVAs

You must deploy two standalone Open Virtual Appliance (OVAs). When you deploy both OVAs, you must meet the following criteria:

- Both OVAs must have the respective management access (eth0) and enhanced fabric management (eth1) interfaces in the same subnet. The eth0 of the active OVA must be in the same subnet as eth0 of the standby OVA. The eth1 of the active OVA must be in the same subnet as eth1 of the standby OVA.
- Both OVAs must be deployed with the same administrative password. This process ensures that both OVAs are duplicates of each other.

After the DCNM Open Virtual Appliance is powered up, verify that all the applications are up and running by using the **appmgr status all** command.

After all of the applications are up and running, stop the applications by using the **appgmr stop all** command.

Note When the Open Virtual Appliance is started up for the first time, please wait for all the applications to run before you shut down any of the applications or power off the virtual appliance.

Creating an NFS/SCP Repository

The DCNM HA cluster needs a server that has both NFS/SCP capabilities. This server is typically a Linux server.

Note The server has to be in the enhanced fabric management network because the switches will use this server to download images and configurations.

Make sure that the exported directory is writable from both peers. The procedure to export a directory /var/lib/sharedarchive on a CentOS server is listed in the following paragraph. The steps will vary based on your environment.

Note You might need root privileges to execute these commands. If you are a nonroot user, please use them with ‘sudo’.

```
[root@repository ~]# mkdir -p /var/lib/sharedarchive
[root@repository ~]# chmod 777 -R /var/lib/sharedarchive
[root@repository ~]# vi /etc/exports
/var/lib/sharedarchive *(rw,sync)

[root@repository ~]# cd /etc/init.d
[root@repository ~]# service nfs restart
```

The same folder /var/lib/sharedarchive can also be accessed through SCP with SCP credentials.

The **/var/lib/sharedarchive * (rw,sync)** command provides read-write permissions to all servers on /var/lib/sharedarchive. Refer to CentOS documentation for information on restricting write permissions to specific peers.

Availability of Virtual IP Addresses

Two free IPv4 addresses are needed to set up VIP addresses. The first IP address will be used in the management access network; it should be in the same subnet as the management access (eth0) interface of the OVAs. The second IP address should be in the same subnet as enhanced fabric management (eth1) interfaces (switch/POAP management network).

Installing an NTP Server

For most of the HA functionality to work, you must synchronize the time on both OVAs by using an NTP server. The installation would typically be in the management access network (eth0) interfaces.

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 theOVA management network
LDAP	OpenLDAP Mirror-mode replication	No	—
XMPP	Not available in HA	Yes	Two VIPs defined, one on each network
DHCP	ISC DHCPD Failover	No	—
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].

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: **Admin > Federation**. 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 http://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.

OpenLightweight Directory Access Protocol

The DCNM Open Virtual Appliance installs an LDAP server an asset database to the switches.

This section contains the following topics:

- "Using the DCNM Open Virtual Appliance-Packaged (Local) LDAP Server" section on page 7-8
- "Using the Remote LDAP Server" section on page 7-9

Using the DCNM Open Virtual Appliance-Packaged (Local) LDAP Server

LDAP HA is achieved through OpenLDAP mirror mode replication. Each LDAP server that is running on one DCNM Open Virtual Appliance becomes a duplicate of the LDAP server that is running on the other Open Virtual Appliance.

DCNM and LDAP Interaction

Both LDAP IP address show up in the Cisco DCNM Web Client (**Admin->Fabric Settings**) in the following order: LDAP-A, LDAP-B.

Cisco DCNM attempts to write on LDAP-A as follows.

- If the write operation succeeds, the data gets replicated to LDAP-B.
- If the write operation fails, then Cisco DCNM writes to LDAP-B.

  The data on LDAP-B eventually gets replicated to LDAP-A when it becomes available.

Switch and LDAP Interaction

When you configure the asset databases, every switch is configured with multiple LDAP servers, as shown in the following example.

The first active LDAP server that is configured in the switch becomes the Active LDAP server. The Active LDAP server is queried first for autoconfigurations.

For every read operation that the switch needs to perform, the Active LDAP server is contacted first, followed by the rest of the LDAP servers.

```
Leaf-0 # fabric database type network
Leaf-0 (config-fabric-db)# server protocol ldap host <LDAP-1-IP> vrf management
Leaf-0 (config-fabric-db)# db-table ou=networks,dc=cisco,dc=com key-type 1
Leaf-0 (config-fabric-db)# server protocol ldap host <LDAP-2-IP> vrf management
Leaf-0 (config-fabric-db)# db-table ou=networks,dc=cisco,dc=com key-type 1
```

Use the **show fabric database statistics** command to find the Active LDAP server, which is marked by an asterisk (*) in the output.

```
Leaf-0 # show fabric database statistics
DB-Type         Requests    Dispatched  Not dispatched  Re-dispatched
---------------------------------------------------------------------
network                1             1               0              0
cabling                0             0               0              0
profile                1             1               0              0
---------------------------------------------------------------------
TOTAL                  2             2               0              0

Per Database stats:
 T Prot Server/DB                       Reqs     OK  NoRes    Err  TmOut   Pend
-------------------------------------------------------------------------------
 n ldap 10.77.247.147                      5      2      1      2      0      0
```

```
*n ldap 10.77.247.148                     3      3      0      0      0      0
*p ldap 172.23.244.122                    1      1      0      0      0      0
Legend:
  T-Type (N-Network, C-Cabling, P-Profile)
  *-Active Server
```

In the previous example, during autoconfiguration, a leaf switch first queries 10.77.247.148, which is the active network database (indicated by "*n"). If that is not available, it automatically contacts the second LDAP server configured as an network database (10.77.247.147 in this example).

Using the Remote LDAP Server

This section describes the behavior when you use a remote LDAP server in an HA environment.

Cisco DCNM and LDAP Interaction

Cisco DCNM allows only two external LDAP servers that are assumed to be synchronized with each other.

Switch and LDAP interaction

The switch and LDAP interaction that use the remote LDAP server is the same interaction as when you are using the Open Virtual Appliance-packaged LDAP. The Active LDAP server is contacted first; if it is not reachable, the switch then attempts to read from the next available LDAP server.

DCHP HA

DHCP on both OVAs listen on the interface of the enhanced fabric management network. The native Internet Systems Consortium (ISC) DHCPD failover mechanism is be used for HA. The lease information is automatically synchronized using native code.

DHCP POAP

The switches do a DHCP broadcast and get response from the Active DHCP server.

DHCP Autoconfiguration

When a tenant host or virtual machine (VM) comes up, it sends a broadcast that is relayed by the leaf node. In such a scenario, the VM profiles should be configured with both relay addresses of OVA-A and OVA-B.

```
interface vlan $vlanid
. . .
ip dhcp relay 1.2.3.4 vrf ..# eth1 IP of OVA-A
ip dhcp relay 1.2.3.5 vrf ..# eth1 IP of OVA-B
```

Changing DHCP Scope Configurations

Scope changes through the Cisco DCNM UI ensure proper synchronization of scopes among the peers. We do not recommend that you do a manual configuration of the DHCP scope configuration file.

Note You must update the IP range for the default scope before creating the new scope, otherwise DHCP will be unable to star. See the "Starting DHCP in an HA Setup" section on page 7-15 for information on updating the IP range for the DHCP scope through the Cisco DCNM UI.

Repositories

All repositories must be remote.

Extensible Messaging and Presence Protocol (XMPP)

HA Implementation

XMPP HA is achieved by having two instances of XMPP applications run with a common database and having a Virtual (floating) IP direct the traffic to the active/standby XMPP.

XMPP Virtual IP Usage

An OVA HA setup has two VIP addresses (one for each network) for the Cisco XCP at the default XMPP port. These VIPs can be used for accessing XMPP services on the OVA management network and the enhanced fabric management network. For example, jabber clients can point to the VIP in the OVA 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 group chat interactions.

Application Failovers

If the XMPP on OVA-A fails, VIP address still resides on OVA-A but the traffic gets redirected to the XMPP that is running on OVA-B.

Application Failbacks

When the XMPP on OVA-A comes up, the VIP address automatically redirects the successive requests to the XMPP running on OVA-A

Virtual-IP Failovers

The VIP address that is configured by default on OVA-A can fail due to two reasons:

- The load-balancing software running on OVA-A fails.
- OVA-A fails.
  - If a load-balancing software failure occurs, the VIP address floats over to OVA-B, and from there it directs the requests to the XMPP on OVA-A.
  - If an OVA-A failure occurs, the VIP address floats over to OVA-B, and directs the requests to XMPP-B. In both cases, the VIP address of Cisco DCNM automatically migrates to OVA-B. The only difference is which XMPP will be used after the failover.

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.

When OVA-A is brought up and XMPP is starts 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. XMPP starts on OVA-A.
3. OVA-B goes down or the load-balancing software fails on OVA-B

Configuring DCNM HA

Because both of the OVAs in an HA environment are deployed identically, either one of them can be the Active peer. The other Open Virtual Appliance would be the Standby peer. All of the configuration CLI commands in the following sections are executed from the secure shell (SSH) terminal.

Configuring the Active Peer

Note Before starting the High Availability setup in active and standby, you need to stop all applications using the **appmgr stop all** command

Step 1 Log in to the SSH terminal of the Open Virtual Appliance that you want to become the Active peer and enter the **appmgr setup ha active** command.

```
Active-peer# appmgr setup ha active
*********************************************************
You are about to enable High Availability in this DCNM virtual appliance.
Please make sure that you the following
1.      An Oracle Database with one user defined for DCNM and one for XMPP
2.      A repository with NFS/SCP capabilities
3.      An NTP server for time synchronization
4.      A couple of free IP addresses to be used as Virtual IPs (one on each port group)
5.      A peer DCNM deployed with the same user profile (same username/password)
6.      Shut down all applications in this server using 'appmgr stop all'

*********************************************************
Do you want to continue? [y/n] [y]
```

Step 2 Make sure that each prerequisite is in place and press **y**; if not all of the prerequisites are in place, press **n** to exit.

A prompt for the root password appears.

```
. . .
Enter the root password of this DCNM : <root-password-of-active-peer>
Enter it again for verification: <root-password-of-active-peer>
. . .
```

Step 3 Enter the administrative password created during Open Virtual Appliance installation.

You will now be prompted for the management access interface (eth0 IP address) of the Standby peer.

Step 4 Enter the management IP address of the peer DCNM.

The active Open Virtual Appliance generates a pair of authentication keys and transfers it to the peer's authorized keys.

a. Enter the root password of the Standby peer when prompted.

All of the other network information needed from the Standby peer is automatically picked up by the Active peer and displayed for confirmation.

b. Ensure that it is the correct peer and press **y** to continue.

```
. . .
Enter the mgmt IP of the peer DCNM (eth0 IP)  : <peer eth0 IP>
Generating ssh keys..
Enter the root password of the peer
root@10.77.247.148's password: <standby-peer root password>
Retrieving information...
Peer Details :
=============
Hostname:  abc.xyz.com
Eth0 IP :  1.2.3.4
Eth1 IP :  192.168.57.148
Do you want to continue? [y/n] [y]
```

Step 5 Enter the VIP addresses for both the management access (eth0) and enhanced fabric management networks (eth1).

Make sure that the VIP addresses are currently not used by any other interfaces in their respective networks.

```
Setting the Virtual IP addresses
============================
The Virtual IP in the eth0 network.
It serves as a single point of access for the following applications: DCNM REST API, AMQP
Enter the VIP  : <a free IP from eth0 subnet>

The Virtual IP in the eth1 network.
It serves as a single point of access for the following applications: DCNM REST API from
the switch network
Enter the VIP  : <a free IP from eth1 subnet>
```

Step 6 Enter the database URL to set the database. The script uses a JDBC thin driver, so you should enter the URL in the same format.

a. Enter the database password.

b. Enter the database password again for verification.

The script tries to do a sample query from the database to check the details entered. The Cisco DCNM schema and related data are loaded after you confirm that all the data are valid.

```
Setting the Database for DCNM and XMPP
======================================
Enter the DB URL {ex. jdbc:oracle:thin:@ipaddr:1521:<SID or Servicename>}
Enter the DB username for DCNM tables: <dcnm-dbuser>
Enter the DB password for DCNM tables :
Enter it again for verification:
Enter the DB username for XMPP tables: <xmpp-dbuser>
Enter the DB password for XMPP tables :
Enter it again for verification:
```

Step 7 Enter an FQDN that will be used as a common XMPP domain name.

```
Common FQDN for Virtual IPs on both DCNM mgmt and EFM networks
==============================================================
Enter the common FQDN for VIP on both DCNM mgmt and EFM networks:
```

Step 8 Enter repository settings:

a. Enter an SCP/NFS repository IP address for the enhanced fabric management network.

b. Enter the IP/exported-directory location.

The script does a test mount and unmounts it shortly after. It is permanently mounted after user confirmation. Similar checks are done for SCP repository users.

c. You will have to enter the SCP password three times (twice for the script and the third time when the script does a test write on the repository).

d. Enter an NTP server IP address. This step is very important for all the applications that run on a cluster.

```
Repository/NTP Details

note: A repository server in the DFA network that has both NFS and SSH/SCP capability.
=======================
Enter the SCP/NFS repository IP  : <repository IP>
NFS Exported location {ex. /var/shared/dcnm/}  : /var/lib/dcnmuser
Performing a test mount to ensure that the server is reachable..
Performing a test-write to ensure the exported directory is writable
test-write successful. Proceeding..
Enter the SCP username for <repository IP>  : <repository user>
Enter the SCP password :
Enter it again for verification:
Performing a test-write to ensure the directory is writable through SCP..
root@<repository-ip>'s password:
test-write successful. Proceeding..
Enter an NTP server for time synchronization  : 10.56.14.161
```

Step 9 A summary of the details entered will be displayed. If you want to reenter the details, press **n**.

Once the HA setup is complete, you can check the role of the ha as follows:

```
OVA-A # appmgr show ha-role
Active
```

Configuring the Standby peer

Step 1 Log in to the SSH terminal of OVA-B and enter the **appmgr setup ha standby** command.

```
OVA-B # appmgr setup ha standby
*********************************************************
You are about to enable High Availability in this DCNM virtual appliance.
Please make sure that you the following
1.      A peer DCNM virtual appliance deployed with the same user and configured as Active
peer
2.      Shut down all applications in this server using 'appmgr stop all'
*********************************************************
Do you want to continue? [y/n] [y]
```

Step 2 Press **y** to continue.

The standby Open Virtual Appliance generates a pair of authentication keys and transfers it to the peer's authorized keys.

a. Enter the root password of the Active peer when prompted.

All the other network information entered during active the Open Virtual Appliance setup is automatically picked up by the Standby peer and displayed for confirmation.

b. Carefully check if it is the correct peer and press **y** to continue.

```
Retrieving information from details entered on Active...
Generating ssh keys..
Enter the root password of the peer
Warning: Permanently added '10.77.247.147' (RSA) to the list of known hosts.
Peer Details :
=============
Hostname        :  somehost.cisco.com
Eth0 IP :  10.77.247.147
Eth1 IP :  192.168.57.147

****************************************
Summary of details entered
****************************************

Virtual IP
=========================
Virtual IP in eth0 n/w  : 10.77.247.143
Virtual IP in eth1 n/w  : 192.168.57.143

Database for DCNM and XMPP
==========================
DB URL : jdbc:oracle:thin:@10.77.247.11:1521:XE
DB username for DCNM    : dcnmuser
DB Username for XMPP    : xmppuser


Archives/Repositories
=========================
SCP/NFS repository IP   : 10.77.247.11
NFS Exported location   : /var/lib/dcnmuser
SCP username            : root
NTP server              : 10.56.14.161

***************************************
Do you want to continue? [y/n] [y]
```

Once confirmed, OVA-B is configured to be a Standby peer, and the following message is displayed.

```
…
******************************************************************************
This node has been configured as standby
Please run 'appmgr start all' first on the active peer (10.77.247.147), and then on the
standby peer(10.77.247.148) to start using applications.
** note ** : dhcpd will not be up until the default poap scopes are updated with free IP
addresses from DCNM GUI
******************************************************************************
```

Note For information about updating default POAP scopes and starting DHCP using HA, please see, Starting DHCP in an HA Setup, page 7-15.

Step 3 Check the HA role of the node by entering the **appmgr show ha-role** command.

```
OVA-A # appmgr show ha-role
```

```
    Standby
```

Starting Applications in the Active Peer

Step 1 Log in to the SSH terminal of the Active peer (OVA-A) and start all applications by entering the **appmgr start all** command.

Step 2 Wait for all the applications to start. Once all applications (except dhcpd) are up and running, go to the next procedure.

Note To start DHCP using HA, see the "Starting DHCP in an HA Setup" section on page 7-15.

Starting Applications in the Standby Peer

Step 1 Login to the SSH terminal of the Standby peer and start all applications using the **appmgr start all** command. Wait for all the applications to start.

Step 2 Once all applications (except dhcpd) are up/running, proceed to the next step.

Note For starting DHCP using HA, please see, Starting DHCP in an HA Setup, page 7-15

Starting DHCP in an HA Setup

In an HA setup, DHCPD will be initially down. In this procedure, you will update the IP range address for the POAP DHCP scope. Use the following procedure to bring up DHCP.

Note You must update the IP range for the default scope before creating the new scope, otherwise DHCP will be unable to start.

Step 1 Log in to Cisco DCNM web UI.

Step 2 On the menu bar, choose **Config> POAP > DHCP Scope** and enter the free IP range address for the default DHCP scope named enhanced_fabric_mgmt_scope.

Step 3 Click **Apply**.

DHCP is automatically started on both the OVAs.

Step 4 Verify all applications are running by opening an SSH terminal session and using the **appmgr status all** command.

Troubleshooting DCNM HA

Switchover in DCNM HA Pair

In Cisco DCNM HA, switchover occurs every time the old Active becomes functional.

When an Active node goes down in an HA setup, the Standby node takes the role of the Active node. When the Active node is functional again, it takes the role of Active node.

When active node (A) goes down, backup node (B) takes role of active. But when (A) comes up again, it takes role of active again. This is not expected behaviour. Old active should not become active again unless a fail-over is triggered or heartbeat instances cannot talk to each other

This happens when shut/no shut is done on the switch interfaces connected to DCNM eth1 interfaces. Heartbeat detects that there has been a split-brain. Since both nodes detect this condition, they both shut down and restart. This is why GUI moves back from node B to node A when node A is reachable again.

'HA Ping' is a feature which can help with this by shutting down heartbeat instances that cannot reach (ping) a specified machine.

Details of HA ping:

Since HA has already been setup, the following needs to be manually enabled:

1) Choose the machine to ping (e.g. the Nexus switch)

2) Run these commands on both instances, setting the appropriate HA ping IP address and peer IP address:

```
HA_PING_ADDRESS=
PEER_ETH1_IP=
echo "* * * * * root /sbin/ha-ping.sh" > /etc/cron.d/ha-ping
echo "IP=$HA_PING_ADDRESS" > $DCNM_HA_HOME/ha-ping.conf
echo "PEER_IP=$PEER_ETH1_IP" >> $DCNM_HA_HOME/ha-ping.conf
chkconfig heartbeat off
sed -i "s/APP_STATUS_HEARTBEAT=.*/APP_STATUS_HEARTBEAT=ha-ping/g"
/root/.DO_NOT_DELETE
```

A probable solution could be to increase heartbeat's deadtime to 60 or even 90 seconds to avoid running into the issue again if shut/no shuts are 30 to 60 seconds apart. Note: This will make takeovers slower.

Commands to increase the timers -

you need to edit /etc/ha.d/ha.cfg, deadtime is specified there.

Please follow this procedure:

1) run appmgr stop ha-apps on standby node

2) run appmgr stop ha-apps on active node

3) edit /etc/ha.d/ha.cfg on both nodes

4) run appmgr start ha-apps on old active node (which was shut down in step 2)

5) wait for old active node to become active again (appmgr show ha-role to check)

6) run appmgr start ha-apps on old standby node (which was shut down in step 1)