



Cisco EPN Manager 2.2 High Availability Installation


This chapter provides information and procedures for installing Cisco EPN Manager in a high availability environment:

- [High Availability Overview, on page 1](#)
- [High Availability Deployment Considerations, on page 2](#)
- [Prerequisites for High Availability Installations, on page 5](#)
- [Install Cisco EPN Manager 2.2 in a High Availability Deployment, on page 6](#)

High Availability Overview

The Cisco EPN Manager high availability (HA) system ensures continued system operation in case of failure. HA uses a pair of linked, synchronized Cisco EPN Manager servers to minimize or eliminate the impact of application or hardware failures that may take place on either server.

The following figure shows the main components and process flows for a high availability deployment.



A high availability deployment consists of a primary and a secondary server with Health Monitor (HM) instances (running as application processes) on both servers. When the primary server fails (either automatically

or because it is manually stopped), the secondary server takes over and manages the network while you restore access to the primary server. If the deployment is configured for automatic failover, the secondary server takes over the active role within two to three minutes after the failover.

When issues on the primary server are resolved and the server is in a running state, it remains in standby mode and begins syncing its data with the active secondary server. When failback is triggered, the primary server again takes over the active role. This role switching between the primary and secondary servers generally takes approximately two to three minutes unless the primary server was reinstalled after failure, in which case it would take longer (based on the size of your setup).

For more information about HA, see the High Availability sections in the [Cisco Evolved Programmable Network Manager User and Administrator Guide](#).

High Availability Deployment Considerations

- [High Availability Deployment Models](#)
- [Understand High Availability Limitations](#)
- [Consider Whether You Can Use Virtual Addresses](#)
- [Best Practices if Firewall is Used Between Primary and Secondary Servers](#)

High Availability Deployment Models

Cisco EPN Manager supports the following High Availability (HA) deployment models.

HA Deployment Model	Primary and Secondary Server Location	Example:
Local	On the same subnet (Layer 2 proximity)	Servers located in same data center
Campus	Different subnets connected via LAN	Servers located in same campus, city, state, or province
Remote	Different subnets connected via WAN	Servers are geographically dispersed

Consider the following factors when deciding whether to use the Local, Campus, or Remote HA deployment model:

- **Exposure to disaster**—The more distributed the deployment model, the less risk to the business as a result of a natural disaster. Remote HA deployments are least likely to be affected by natural disaster, allowing for a less complex and costly business continuity model. Local HA deployments are most vulnerable to disaster because of server co-location.
- **Whether you can use a virtual IP address**—Only Local HA deployments can use virtual IP addresses. A virtual IP address is a single IP address that will always point to the active server, even after a failover and failback. It also allows both the primary and secondary servers to share a common management IP address.
- **Bandwidth/latency**—Bandwidth would be highest and latency would be lowest in Local HA deployments because the primary and secondary servers are connected by short network links that have high bandwidth and low latency. Campus HA deployments may have lower bandwidth and higher latency than Local HA deployments. Remote HA deployments have the least bandwidth and the highest latency.
- **Administration**—HA administration is simplest for Local HA deployments, with increasing complexity for Campus and Remote HA deployments. Remote HA deployments will require administrative remedying.

- Configuration of device event forwarding—Configuring event forwarding can be simplest with Local HA deployments because you can use a virtual IP address, and then configure your devices to forward events to that single virtual IP address. Without a virtual IP address, you must configure your devices to forward events to both the primary and secondary servers.

For more details about HA, see the [Cisco Evolved Programmable Network Manager User and Administrator Guide](#).

Understand High Availability Limitations

The Cisco EPN Manager HA system is subject to the following limiting factors (this applies to all HA deployment models):

- The HA system requires at least 255 Mbps of network bandwidth to handle HA operations (the ideal is 977 Mbps). These operations include (but are not restricted to) HA registration, database and file synchronization, and triggering failback. Because Cisco EPN Manager uses a single physical port for all its networking needs, there can be occurrences of insufficient bandwidth which in turn will affect HA performance.
- The HA system requires low latency (less than 100 msec) across network links between the primary and secondary servers. Irrespective of the physical proximity of these two servers, high latency on these links can affect how Cisco EPN Manager maintains sessions between the primary and secondary servers. This is because larger databases require more synchronization transactions which require lower latency and higher bandwidth. If you are managing a relatively small network using Cisco EPN Manager, your database would be smaller and therefore, HA might work with a higher network latency and less bandwidth.
- HA performance is always sensitive to the network throughput delivered by the network that connects the primary and secondary servers. This restriction applies (to some degree) to all of the deployment models. For example, in a geographically dispersed deployment, a Remote HA deployment is more likely to have problems due to low bandwidth and high latency. However, if Local and Campus HA deployments are not properly configured, they are highly susceptible to problems with latency that result from bandwidth limitations on high-usage networks.

For assistance in determining whether your network is suitable for any of the HA variations, please contact your Cisco representative.

Consider Whether You Can Use Virtual Addresses

Using virtual IP addresses in a Local HA deployment setup gives your users the ability to connect to the active server using a single IP address or web URL without having to know which server is actually active. Virtual IP addresses also allow both servers to share a common management IP address. During normal operation, the virtual IP address points to the primary server. If a failover occurs, the virtual IP address automatically points to the secondary server. When failback occurs, the virtual IP address automatically switches back to the primary server.

To use a virtual IP addresses, the following IP addresses must be on the same subnet:

- The virtual IP address
- The IP addresses of the primary and secondary servers
- The IP address of the gateway configured on both primary and secondary servers

The following example illustrates how virtual, primary, and secondary IP addresses should be assigned with respect to each other. If the primary and secondary servers have the following IP address assignments within the given subnet, the virtual IP address for both servers can be assigned as follows:

- Subnet mask: 255.255.255.224 (/32)
- Primary server IP address: 10.10.101.1
- Secondary server IP address: 10.10.101.2
- Virtual IP address: 10.10.101.[3-30] e.g., 10.10.101.3. Note that the virtual IP address can be any of a range of addresses that are valid for the given subnet mask.

If you do not use a virtual IP address, you must configure your devices to forward events to both the primary and secondary servers (for example, by forwarding them to a given subnet or range of IP addresses that includes both the primary and secondary server). To reduce (or eliminate) the chance of losing data, you must configure device event forwarding before a failover occurs. You do not need to make any changes to the secondary server during installation; simply provision the primary and secondary servers with their individual IP addresses.

Whether your HA deployment uses a single IP address or not, users should always connect to the Cisco EPN Manager web GUI using the active server IP address/URL.

Best Practices if Firewall is Used Between Primary and Secondary Servers

Firewalls between the primary and secondary servers should be configured to avoid short timeouts for TCP packets to allow enough time for HA registration and other processes.

Following is the procedure for changing the Oracle and OS timeouts, if necessary. Use this procedure if the failback operation fails repeatedly.

Before you Begin

Back up the files specified in the procedure below.



Note This procedure must be performed on both primary and secondary servers.

- Step 1** Open the following file:
`/opt/oracle/base/product/12.1.0/dbhome_1/network/admin/sqlnet.ora`
- Step 2** Add the following parameters:
`SQLNET.EXPIRE_TIME=2`
`DISABLE_OOB=on`
`SQLNET.INBOUND_CONNECT_TIMEOUT=600`
- Step 3** Open the following file:
`/opt/CSColumos/bin/ha_dgmgrl.sh`
- Step 4** Add the following line under the register() function:
`edit database $1 set property NetTimeout=1000`

Following is an example excerpt of the file with the relevant line in bold:

```
connect $4/$5@$3
remove configuration;
create configuration $DGMGRL_CONFIG_NAME as primary database is $1 connect identifier is $1;
edit database $1 set property NetTimeout=1000
;
add database $2 as connect identifier is $2 maintained as physical;
enable configuration;
```

Step 5 Open the following file:

/etc/sysctl.conf

Add the following commands to the end of the file:

```
net.ipv4.tcp_keepalive_time = 80
net.ipv4.tcp_keepalive_probes = 7
net.ipv4.tcp_keepalive_intvl = 10
```

Step 6 Run the following command from the root shell:

```
sysctl -system
```

Prerequisites for High Availability Installations

The following prerequisites must be met before installing Cisco EPN Manager in a high availability deployment:

- Make sure that your hardware and software meet the requirements listed in the relevant prerequisites topic:
 - [OVA/VM Requirements](#).
 - [Bare Metal Requirements for Remote High Availability Deployments](#).
- Make sure the secondary server is configured as follows:
 - The secondary server's hardware and software specifications must be the same as those of the primary server. For example, if you installed Cisco EPN Manager on the primary server and specified the Professional system size, your secondary server must also be installed using the Professional system size, and must meet all requirements for Professional-size servers in [System Requirements](#).
 - The secondary server must be running the same software level as the primary server (including the patch level).
 - If you plan to use a virtual IP address for a Local HA deployment, the virtual IP address, primary, and secondary servers must be on the same subnet. The gateway on the primary and secondary servers must also reside on the same subnet.
- If there is a firewall between the primary and secondary servers, there must be permission from the firewall for the ports used by HA. The ports are listed in [Ports Used by Cisco EPN Manager](#).
- Prepare the following information which you will need to enter during the installation:
 - The IP address or host name of the secondary server (if you are not using a virtual IP address). You will need it when configuring HA on the primary server.
 - The virtual IPv4 and IPv6 (if used) IP addresses you want to use for both servers (if you plan to use a virtual IP address).

- The password you want to use for the HA authentication key. It will be used to authenticate communications between the primary and secondary servers. You will need to enter it when you configure HA—that is, when you register the secondary server on the primary server (also called *pairing* the servers). Finally, you will need it to log in to the secondary server's Health Monitor page.
- A Cisco EPN Manager web GUI user ID with Administration privileges on the primary server. You will also need the user's password.
- A valid email address to which HA notifications can be sent.

Install Cisco EPN Manager 2.2 in a High Availability Deployment

The procedure in this section is for a fresh installation of the product in a high availability environment. If you are upgrading to Cisco EPN Manager 2.2 from a previous version, see [Upgrade to Cisco EPN Manager 2.2 \(High Availability\)](#).

Before You Begin

Make sure your servers meet the requirements listed in [Prerequisites for High Availability Installations](#).

-
- Step 1** Install Cisco EPN Manager on the primary server as described in [Install Cisco EPN Manager 2.2 \(No HA\)](#).
- Step 2** Install Cisco EPN Manager on the secondary server as described in [Install Cisco EPN Manager 2.2 \(No HA\)](#).
- Step 3** When you are prompted to choose whether you want this newly-installed server to act as a secondary failback server in an HA implementation, enter **yes**.
- Step 4** Enter a password which will be used as the *HA authentication key* for communication between the primary and secondary servers. You will need this key to configure HA. (During normal operation, you will need to enter the HA authentication key to log in to the secondary server's Health Monitor page.)
- Step 5** Enter the password again to confirm.
- Step 6** Enter **Y** to confirm that you want to install this server as a secondary server. When the installation is complete, the VM (OVA/VM) or Cisco UCS server (ISO/bare metal) will reboot.
- Step 7** Log in using the Cisco EPN Manager CLI admin username and password you specified during the installation.
- Step 8** Verify that all processes are running on the secondary server using the **ncs status** command. Ensure that, at a minimum, the following services are up and running: Health Monitor, Database, NMS, SAM Daemon, DA Daemon, Compliance Engine. Note that for optimal Cisco EPN Manager functionality, all services should be up and running.
- Step 9** Make sure all devices are configured to forward events (syslogs, traps, and TL1 messages) to both servers (or the virtual IP address, if you are using one).
- Note** If you do not perform this step *before* registering the secondary server on the primary server and a failover occurs, you may lose some data.
- Step 10** Configure HA by registering the secondary server on the primary server. The registration process must be performed from the primary server. For more information, see the section on registering the secondary server on the primary server in the [Cisco Evolved Programmable Network Manager User and Administrator Guide](#).
-