



Cisco EPN Manager High Availability Installation


This chapter provides information about Cisco EPN Manager in a high availability environment.

- [High Availability Overview, on page 1](#)
- [High Availability Deployment Models, on page 2](#)
- [Understand High Availability Limitations, on page 3](#)
- [Evaluate the Use of Virtual IP Addresses, on page 3](#)
- [Prerequisites for High Availability Installations, on page 4](#)
- [Install Cisco EPN Manager in a High Availability Deployment, on page 5](#)
- [Check Readiness for HA Configuration, on page 5](#)

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 server and a secondary server with Health Monitor (HM) instances (running as application processes) on both servers. When the primary server fails (due to a problem

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 primary server failure.

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 *Configure and Manage High Availability* section in the [Cisco Evolved Programmable Network Manager 8.0 User and Administrator Guide](#).

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 the 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, enabling a less complex and less expensive 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 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 require administrative remedying.
- **Configuration of device event forwarding**—Configuring event forwarding can be the 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 information about HA, see the *Configure and Manage High Availability* section in the [Cisco Evolved Programmable Network Manager 8.0 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 a minimum of 500 Mbps (Mega bit per second) or higher of network bandwidth to handle HA operations. 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 affects HA performance.
- The HA system requires low latency (maximum 100 ms, preferably under 70 ms) 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 will be smaller and therefore, HA will 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 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 configured correctly, 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.

Evaluate the Use of Virtual IP 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 the need to determine which server is currently 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 a failback occurs, the virtual IP address automatically switches back to the primary server.

To use a virtual IP address, 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 that is configured on both primary and secondary servers

The following example demonstrates the proper assignment of virtual, primary, and secondary IP addresses within a given subnet. Consider the scenario where the primary and secondary servers are assigned specific IP addresses:

- Subnet mask: 255.255.255.224 (/27)

- Primary server IP address: 10.10.101.1
- Secondary server IP address: 10.10.101.2

In this case, the virtual IP address, which represents the active server in a high availability setup, can be chosen from a range of valid addresses within the subnet. For instance, you could assign the virtual IP address as 10.10.101.3. The range of valid virtual IP addresses extends from 10.10.101.3 to 10.10.101.30, as defined by the 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.

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 that are listed in [OVA/VM Requirements](#).
- Make sure that 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 that are used by HA. The ports are listed in [Ports Used by Cisco EPN Manager](#).
- The following information are required during the installation process:
 - The IPv4 IP address or host name of the secondary server (if you are not using a virtual IP address). You need it when configuring HA on the primary server.
 - The virtual IPv4 and IPv6 (if used) IP addresses that you want to use for both servers (if you plan to use a virtual IP address).
 - The password that you want to use for the HA authentication key. This password was provided by the user during the installation of the secondary server. It will be used to authenticate communications between the primary and secondary servers. You must enter it when you configure HA—That is,

when you register the secondary server on the primary server (also called *pairing* the servers). Finally, it is necessary 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 along with its password.
- A valid email address to which HA notifications can be sent.

Install Cisco EPN Manager 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 8.0 from a previous version, see .

Before You Begin

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

-
- Step 1** Install Cisco EPN Manager on the primary server as described in .
- Step 2** Install Cisco EPN Manager on the secondary server as described in .
- Step 3** When you are prompted to choose whether you want this newly installed server to act as a secondary fallback server in an HA implementation, enter **yes**.
- Step 4** Enter a password that will be used as the *HA authentication key* for communication between the primary and secondary servers. You need this key to configure HA. (During normal operation, you must 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) reboots.
- Step 7** Log in using the Cisco EPN Manager CLI admin username and password you specified during the installation.
- Step 8** 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 9** 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 *Configure and Manage High Availability* section in the [Cisco Evolved Programmable Network Manager 8.0 User and Administrator Guide](#).
-

Check Readiness for HA Configuration

During the HA configuration, other environmental parameters that are related to HA like system specification, network configuration, and bandwidth between the servers determine the completion of HA configuration.

Fifteen checks are run in the system to ensure the completion of HA configuration without any error or failure. The checklist name and the corresponding status with recommendations if any, will be displayed when you run the Check Readiness feature.



Note The **Check Readiness** does not block the HA configuration. You can configure HA even if some of the checks do not pass.

If the primary and secondary authentication keys are different, then Check Readiness will not proceed. You can proceed with HA Registration.

To check readiness for HA configuration, follow these steps:

- Step 1** Log in to Cisco EPN Manager with a user ID and password that has administrator privileges.
- Step 2** From the menu, select **Administration > Settings > High Availability**. Cisco EPN Manager displays the HA status page.
- Step 3** Select **HA Configuration**.
- Step 4** Provide the secondary server IP address in the **Secondary Server** field and the secondary authentication key in the **Authentication Key** field.
- Step 5** Click **Check Readiness**.

A pop-up window with the system specifications and other parameters will be displayed. The screen shows the checklist item name, status, impact, and recommendation details.

Below is the list of checklist test names and the description that is displayed for Check Readiness:

Table 1: Checklist name and description

Checklist Test Name	Test Description
SYSTEM - CHECK CPU COUNT	Checks the CPU count in the primary and secondary servers. The CPU count in both servers must satisfy the requirements.
SYSTEM - CHECK DISK IOPS	Checks the disk speed in the primary and secondary servers. The minimum expected disk speed is 200 Mbps.
SYSTEM - CHECK RAM SIZE	Checks the RAM size of the primary and secondary servers. The RAM size of both servers must satisfy the requirements.
SYSTEM - CHECK DISK SIZE	Checks the disk size of the primary and secondary servers. The disk size of both servers must satisfy the requirements.
SYSTEM - CHECK SERVER PING REACHABILITY	Checks that the primary server can reach the secondary server through a ping.
SYSTEM - CHECK OS COMPATABILITY	Checks that the primary server and secondary servers have the same OS version.
SYSTEM - HEALTH MONITOR STATUS	Checks whether the health monitor process is running in the primary and secondary servers.

NETWORK - CHECK NETWORK INTERFACE BANDWIDTH	Checks if the speed of interface eth0 matches the recommended 100 Mbps in primary and secondary servers. This test will not measure the network bandwidth by transmitting the data between the primary and secondary servers.
NETWORK - CHECK FIREWALL FOR DATABASE PORT ACCESSIBILITY	Checks if the database port 1522 is open in the system firewall. If the port is disabled, the test will grant permission for 1522 in the IP tables list.
DATABASE - CHECK ONLINE STATUS	Checks if the database files status is online and accessible in primary and secondary servers.
DATABASE - CHECK MEMORY TARGET	Checks for "/dev/shm" database memory target size for HA setup.
DATABASE - LISTENER STATUS	Checks if the database listeners are up and running in primary and secondary servers. If there is a failure, the test will attempt to start the listener and report the status.
DATABASE - CHECK LISTENER CONFIG CORRUPTION	Checks if all the database instances exist under the database listener configuration file "listener.ora"
DATABASE - CHECK TNS CONFIG CORRUPTION	Checks if all the "WCS" instances exist under the database TNS listener configuration file "tnsnames.ora"
DATABASE - TNS REACHABILITY STATUS	Checks if TNSPING is successful in primary and secondary server.

Step 6 Once the check is completed for all the parameters, check their status and click **Clear** to close the window.

Note Failback and failover events during **Check Readiness** are forwarded to the Alarms and Events page. Configuration failure events are not present in the Alarms and Events list.

Note After changing the HA authentication key using the `ncs ha authkey` command, the HA readiness check may fail with the error *Authentication key does not match between primary and secondary servers*, even though the key is correct. This occurs because the readiness check might be referencing outdated or cached information.

To resolve this issue,

- if the servers are unpaired, run the `ncs ha authkey` command on both the primary and secondary servers with the same password, and restart both servers.
- If the servers are paired, run the command only on the primary server.

