Configure High Availability

How High Availability Works

The following figure shows the main components and process flow for a Prime Infrastructure High Availability (HA) setup with the primary server in the active state.

Figure 1: HA Deployment

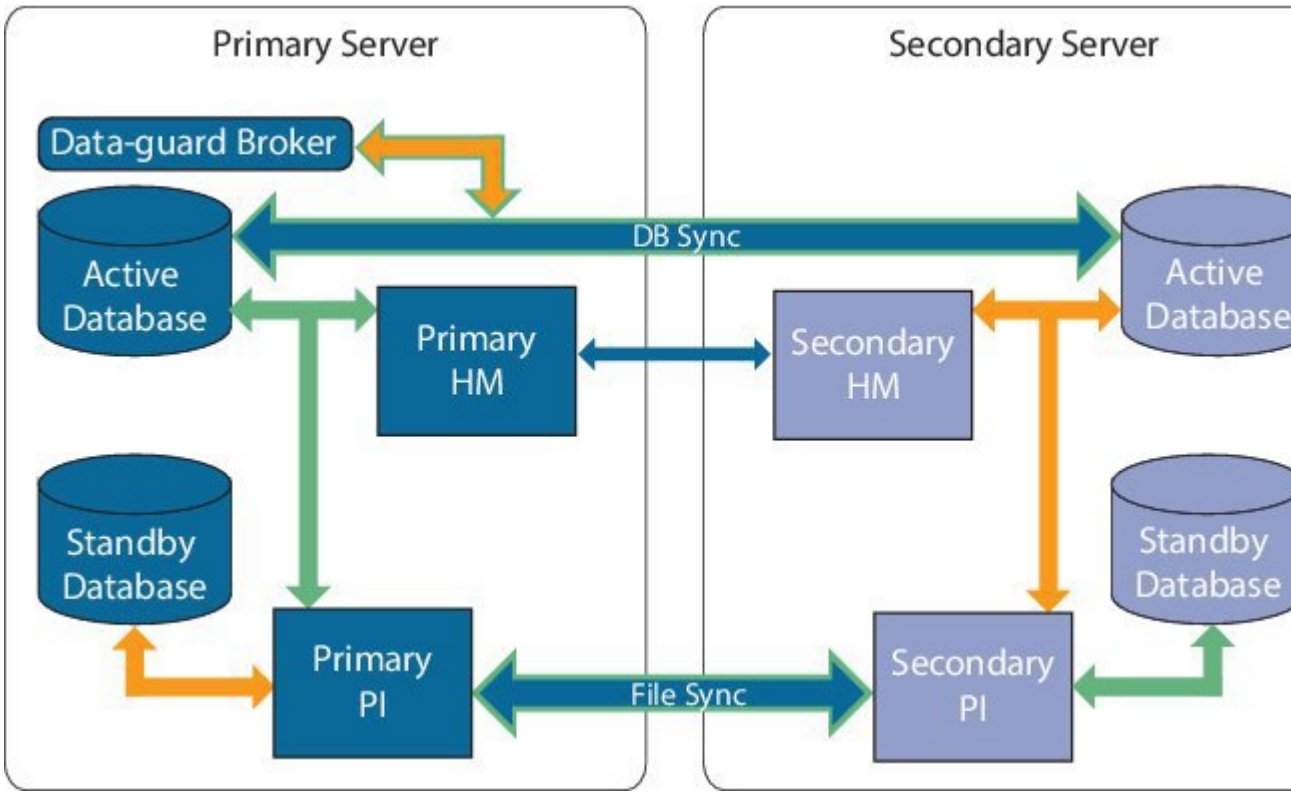

An HA deployment consists of two Prime Infrastructure servers: a primary and a secondary. Each of these servers has an active database and a standby backup copy of the active database. Under normal circumstances, the primary server is active: It is connected to its active database while it manages the network. The secondary server is passive, connected only to its standby database, but in constant communication with the primary server.

The Health Monitor processes running on both servers monitor the status of its opposite server. Oracle Recovery Manager (RMAN) running on both servers creates the active and standby databases and synchronizes the databases when there are changes, with the help of Oracle Data Guard Broker running on the primary server.

When the primary server fails, the secondary takes over, connecting to its active database, which is in sync with the active primary database. You can trigger this switch, called a “failover”, either manually, which is recommended, or have it triggered automatically, You then use the secondary server to manage the network while working to restore access to the primary server. When the primary is available again, you can initiate a switch (called a “failback”) back to the primary server and resume network management using the primary.

If you choose to deploy the primary and secondary servers on the same IP subnet, you can configure your devices to send a notifications to Prime Infrastructure at a single virtual IP address. If you choose to disperse the two servers geographically, such as to facilitate disaster recovery, you will need to configure your devices to send notifications to both servers.

Related Topics

About the Primary and Secondary Servers, on page 3
Sources of Failure, on page 3
File and Database Synchronization, on page 3
HA Server Communications, on page 4

About the Primary and Secondary Servers

In any Prime Infrastructure HA implementation, for a given instance of a primary server, there must be one and only one dedicated secondary server.

Typically, each HA server has its own IP address or host name. If you place the servers on the same subnet, they can share the same IP using virtual IP addressing, which simplifies device configuration. The primary and secondary servers of Prime Infrastructure must be enabled on a network interface ethernet0 (eth0) during HA implementation.

Once HA is set up, you should avoid changing the IP addresses or host names of the HA servers, as this will break the HA setup (see “Reset the Server IP Address or Host Name” in Related Topics).

Related Topics

How High Availability Works, on page 1
Using Virtual IP Addressing With HA, on page 6
Reset the HA Server IP Address or Host Name, on page 49

Sources of Failure

Prime Infrastructure servers can fail due to issues in one or more of the following areas:

- **Application Processes**: Failure of one or more of the Prime Infrastructure server processes, including NMS Server, MATLAB, TFTP, FTP, and so on. You can view the operational status of each of these application processes by running the ncs status command through the admin console.
- **Database Server**: One or more database-related processes could be down. The Database Server runs as a service in Prime Infrastructure.
- **Network**: Problems with network access or reachability issues.
- **System**: Problems related to the server's physical hardware or operating system.
- **Virtual Machine** (VM): Problems with the VM environment on which the primary and secondary servers were installed (if HA is running in a VM environment).

  For more information, see How High Availability Works

File and Database Synchronization

Whenever the HA configuration determines that there is a change on the primary server, it synchronizes this change with the secondary server. These changes are of two types:

1. **Database**: These include database updates related to configuration, performance and monitoring data.
2. **File**: These include changes to configuration files.

Oracle Recovery Manager (RMAN) running on both servers creates the active and standby databases and synchronizes the databases when there are changes, with the help of Oracle Data Guard Broker running on the primary server.

File changes are synchronized using the HTTPS protocol. File synchronization is done either in:

- **Batch**: This category includes files that are not updated frequently (such as license files). These files are synchronized once every 500 seconds.
- **Near Real-Time**: Files that are updated frequently fall under this category. These files are synchronized once every 11 seconds.

By default, the HA framework is configured to copy all the required configuration data, including:

- Report configurations
- Configuration Templates
- TFTP-root
- Administration settings
- Licensing files

Related Topics

How High Availability Works, on page 1

HA Server Communications

The primary and secondary HA servers exchange the following messages in order to maintain the health of the HA system:

- Database Sync: Includes all the information necessary to ensure that the databases on the primary and secondary servers are running and synchronized.
- File Sync: Includes frequently updated configuration files. These are synchronized every 11 seconds, while other infrequently updated configuration files are synchronized every 500 seconds.
- Process Sync: Ensures that application- and database-related processes are running. These messages fall under the Heartbeat category.
- Health Monitor Sync: These messages check for the following failure conditions:
  - Network failures
  - System failures (in the server hardware and operating system)
  - Health Monitor failures

Related Topics

How High Availability Works, on page 1

Health Monitor Process

Health Monitor (HM) is the main component managing HA operations. Separate instances of HM run as an application process on both the primary and the secondary server. HM performs the following functions:

- Synchronizes database and configuration data related to HA (this excludes databases that sync separately using Oracle Data Guard).
- Exchanges heartbeat messages between the primary and secondary servers every five seconds, to ensure communications are maintained between the servers.
- Checks the available disk space on both servers at regular intervals, and generates events when storage space runs low.
- Manages, controls and monitors the overall health of the linked HA servers. If there is a failure on the primary server then it is the Health Monitor's job to activate the secondary server.

Related Topics

How High Availability Works, on page 1

Health Monitor Web Page

You control HA behavior using the Health Monitor web page. Each Health Monitor instance running on the primary server or secondary server has its own web page. The following figure shows an example of the Health Monitor web page for a secondary server in the “Primary Active” and "Secondary Syncing" state.

Figure 2: Health Monitor Web Page (Secondary Server)

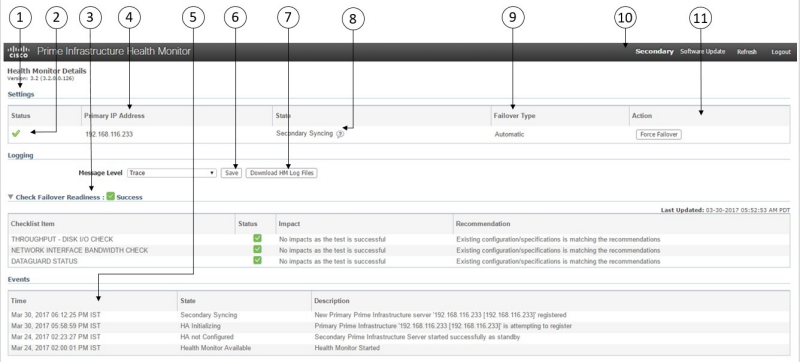

1	Settings area displays Health Monitor state and configuration detail in five separate sections.
2	Status indicates current functional status of the HA setup (green check mark indicates that HA is on and working).
3	Check Failover Readiness field displays the values of system failback and system failover details of the checklist items. For more details, see "Check Failover Readiness" given below the table.
4	Primary IP Address identifies the IP of the peer server for this secondary server (on the primary server, this field is labeled “Secondary IP Address”).
5	Events table displays all current HA-related events, in chronological order, with most recent event at the top.
6	Message Level field lets you change the logging level (your choice of Error, Informational, or Trace). You must press **Save** to change the logging level.
7	Logging **Download** area lets you download Health Monitor log files.
8	State shows current HA state of the server on which this instance of Health Monitor is running.
9	Failover Type shows whether you have Manual or Automatic failover configured.
10	Identifies the HA server whose Health Monitor web page you are viewing.

11	Action shows actions you can perform, such as failover or failback. Action buttons are enabled only when Health Monitor detects HA state changes needing action.

Check Failover Readiness section description:

Checklist Name	Description
SYSTEM - CHECK DISK IOPS	This validates the disk iops in both primary and secondary server. The minimum expected disk iops is 200 MBps.
NETWORK - CHECK NETWORK INTERFACE BANDWIDTH	This checks if the eth0 interface speed matches the recommended speed of 100 Mbps in both primary and secondary sever. This test will not measure network bandwidth by transmitting data between primary and secondary server.
NETWORK - CHECK NETWORK BANDWIDTH SPEED	This checks if the network bandwidth speed matches the recommended speed of 100 Mbps in both primary and secondary sever. This test will measure network bandwidth by transmitting data between primary and secondary server. **Note** In Cisco Prime Infrastructure 3.9, the network bandwidth speed test is calculated only in Mbps. Therefore, GBps, MBps, KBps, and Mbps are changed over to Mbps and given as an input to the speed test.
DATABASE - SYNC STATUS	This ensures the oracle data guard broker configuration which syncs the primary and secondary database.

Trend Graph for Check Failover Readiness :

- Click **Click here** link in the Trend Graph to check the trend graphs for all the check failover readiness test. The trend graphs shows the historical summary of the test and status on the stability of the System/Network.
- Click **Select Date Range** to modify date and time, Click **Apply**. By default, trend graphs displays the latest 6 hours value.

Related Topics

How High Availability Works, on page 1
How to Resolve Database Synchronization Issues, on page 42

Using Virtual IP Addressing With HA

Under normal circumstances, you configure the devices that you manage using Prime Infrastructure to send their syslogs, SNMP traps and other notifications to the Prime Infrastructure server's IP address. When HA is implemented, you will have two separate Prime Infrastructure servers, with two different IP addresses. If

we fail to reconfigure devices to send their notifications to the secondary server as well as the primary server, then when the secondary Prime Infrastructure server goes into Active mode, none of these notifications will be received by the secondary server.

Setting all of your managed devices to send notifications to two separate servers demands extra device configuration work. To avoid this additional overhead, HA supports use of a virtual IP that both servers can share as the Management Address. The two servers will switch IPs as needed during failover and failback processes. At any given time, the virtual IP Address will always point to the correct Prime Infrastructure server.

Note that you cannot use virtual IP addressing unless the addresses for both of the HA servers and the virtual IP are all in the same subnet. This can have an impact on how you choose to deploy your HA servers (see "Planning HA Deployments" and "Using the Local Model" in Related Topics).

Also note that a virtual IP address is in no way intended as a substitute for the two server IP addresses. The virtual IP is intended as a destination for syslogs and traps, and for other device management messages *being sent to the* Prime Infrastructure servers. Polling of devices is always conducted from one of the two Prime Infrastructure server IP addresses. Given these facts, if you are using virtual IP addressing, you must open your firewall to incoming and outgoing TCP/IP communication on all three addresses: the virtual IP address as well as the two actual server IPs.

You can also use virtual IP addressing if you plan to use HA with Operations Center. You can assign a virtual IP as SSO to the Prime Infrastructure instance on which Operations Center is enabled. No virtual IP is needed for any of the instances managed using Operations Center (see "Enable HA for Operations Center").

You can enable virtual IP addressing during HA registration on the primary server, by specifying that you want to use this feature and then supplying the virtual IPv4 (and, optionally, IPv6) address you want the primary and secondary servers to share (see "How to Register HA on the Primary Server").

To remove Virtual IP addressing after it is enabled, you must remove HA completely (see "Remove HA Via the GUI").

Related Topics

What If I Cannot Use Virtual IP Addressing?, on page 12
Planning HA Deployments, on page 9
Using the Local Model, on page 10
Enable HA for Operations Center, on page 13
How to Register HA on the Primary Server, on page 17
How High Availability Works, on page 1
Remove HA Via the GUI, on page 46

How to Use SSL Certificates in an HA Environment?

If you decide to use SSL certification to secure communications between Prime Infrastructure server and users, and also plan to implement HA, you will need to generate separate certificates for both the primary and secondary HA servers.

These certificates must be generated using the FQDN (Fully Qualified Domain Name) for each server. To clarify: You must use the primary server's FQDN to generate the certificate you plan to use for the primary server, and the secondary server's FQDN to generate the certificate you plan to use for the secondary server,

Once you have generated the certificates, import the signed certificates to the respective servers.

Do not generate SSL certificates using a virtual IP address. The virtual IP address feature is used to enable communications between Prime Infrastructure and your network devices.

To set up HTTPS access for Cisco Prime Infrastructure, see Set Up HTTPS Access to Prime Infrastructure

Import Client Certificates Into Web Browsers

Users accessing Prime Infrastructure servers with certificate authentication must import client certificates into their browsers in order to authenticate. Although the process is similar across browsers, the actual details vary with the browser. The following procedure assumes that your users are using a Prime Infrastructure compatible version of Firefox.

You must ensure that the user importing the client certificates has:

- Downloaded a copy of the certificate files to a local storage resource on the client machine
- If the certificate file is encrypted: The password with which the certificate files were encrypted.

Step 1 Launch Firefox and enter the following URL in the location bar: **about:preferences#advanced**.

Firefox displays its **Options** > **Advanced** tab.

Step 2 Select **Certificates** > **View Certificates** > **Your Certificates**, then click **Import...**.

Step 3 Navigate to the downloaded certificate files, select them, then click **OK** or **Open**.

Step 4 If the certificate files are encrypted: You will be prompted for the password used to encrypt the certificate file. Enter it and click **OK**.

The certificate is now installed in the browser.

Step 5 Press **Ctrl+Shift+Del** to clear the browser cache.

Step 6 Point the browser to the Prime Infrastructure server using certificate authentication.

You will be prompted to select the certificate with which to respond to the server authentication requested. Select the appropriate certificate and click **OK**.

Hot Standby Behavior

When the primary server is active, the secondary server is in constant synchronization with the primary server and runs all Prime Infrastructure processes for fast switch over. When the primary server fails, the secondary server immediately takes over the active role within two to three minutes after the failover.

Once issues in the primary server are resolved and it is returned to a running state, the primary server assumes a standby role. When the primary server is in the standby role, the Health Monitor GUI shows “Primary Syncing” state during which the database and files on the primary start to sync with the active secondary.

When the primary server is available again and a failback is triggered, the primary server again takes over the active role. This role switching between the primary and secondary servers happens within two to three minutes.

Related Topics

How High Availability Works, on page 1

Planning HA Deployments

Prime Infrastructure's HA feature supports the following deployment models:

- **Local**: Both of the HA servers are located on the same subnet (giving them Layer 2 proximity), usually in the same data center.
- **Campus**: Both HA servers are located in different subnets connected via LAN. Typically, they will be deployed on a single campus, but at different locations within the campus.
- **Remote**: Each HA server is located in a separate, remote subnet connected via WAN. Each server is in a different facility. The facilities are geographically dispersed across countries or continents.

The following sections explain the advantages and disadvantage of each model, and discusses underlying restrictions that affect all deployment models.

HA will function using any of the supported deployment models. The main restriction is on HA's performance and reliability, which depends on the bandwidth and latency criteria discussed in "Network Throughput Restrictions on HA". As long as you are able to successfully manage these parameters, it is a business decision (based on business parameters, such as cost, enterprise size, geography, compliance standards, and so on) as to which of the available deployment models you choose to implement.

Related Topics

Network Throughput Restrictions on HA, on page 9
Using the Local Model, on page 10
Using the Campus Model, on page 11
Using the Remote Model, on page 11
What If I Cannot Use Virtual IP Addressing?, on page 12
Automatic Versus Manual Failover, on page 12
Enable HA for Operations Center, on page 13

Network Throughput Restrictions on HA

Prime Infrastructure HA performance is always subject to the following limiting factors:

- The net bandwidth available to Prime Infrastructure for handling all operations. These operations include (but are not restricted to) HA registration, database and file synchronization, and triggering failback.
- The net latency of the network across the links between the primary and secondary servers. Irrespective of the physical proximity of these two servers, high latency on these links can affect how Prime Infrastructure maintains sessions between the primary and secondary servers.
- The net throughput that can be delivered by the network that connects the primary and secondary servers. Net throughput varies with the net bandwidth and latency, and can be considered a function of these two factors.

These limits apply to at least some degree in every possible deployment model, although some models are more prone to problems than others. For example: Because of the high level of geographic dispersal, the Remote deployment model is more likely to have problems with both bandwidth and latency. But both the Local and Campus models, if not properly configured, are also highly susceptible to problems with throughput, as they can be saddled by low bandwidth and high latency on networks with high usage.

You will rarely see throughput problems affecting a failback or failover, as the two HA servers are in more or less constant communication and the database changes are replicated quickly. Most failovers and failbacks take approximately two to three minutes.

The main exception to this rule is the delay for a full database copy operation. This kind of operation is triggered when the primary server has been down for more than the data retention period and you then bring it back up. The data retention period for the express, express-plus and standard configurations server is six hours and for professional and Gen 2 appliance server it is12 hours.

Prime Infrastructure will trigger a full database copy operation from the secondary to the primary. No failback is possible during this period, although the Health Monitor page will display any events encountered while the database copy is going on. As soon as the copy is complete, the primary server will go to the "Primary Synching" state, and you can then trigger failback. Be sure not to restart the primary server or disconnect it from the network while the full database copy is in progress.

Variations in net throughput during a full database copy operation, irrespective of database size or other factors, can mean the difference between a database copy operation that completes successfully in under an hour and one that does not complete at all. Cisco has tested the impact of net throughput on HA deployment in configurations following the Remote model, using typical Prime Infrastructure database sizes of between 105 GB and 156 GB. Based on these tests, Cisco recommends for a typical database of 125 GB (generating a 10 GB backup file):

- For best results: With sub-millisecond latency, and net throughput of 977 Mbps or more, expect a complete database copy time of one hour or less.
- For good results: With latency of 70 milliseconds, and net throughput of 255 Mbps or more, expect a complete database copy time of two hours or less.
- For acceptable results: With latency of 220 milliseconds or less, and net throughput of 86 Mbps or more, expect a complete database copy time of 4.5 hours or less.

With latencies of 330ms or higher, and throughput of 46Mbps or less, you run the risk of the database copy not completing successfully.

Related Topics

Planning HA Deployments, on page 9
Using the Remote Model, on page 11

Using the Local Model

The main advantage of the Local deployment model is that it permits use of a virtual IP address as the single management address for the system. Users can use this virtual IP to connect to Prime Infrastructure, and devices can use it as the destination for their SNMP trap and other notifications.

The only restriction on assigning a virtual IP address is to have that IP address in the same subnet as the IP address assignment for the primary and secondary servers. For example: 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 (/27)
- Primary server IP address: 10.10.101.2
- Secondary server IP address: 10.10.101.3
- Virtual IP address: 10.10.101.[4-30] e.g., 10.10.101.4. Note that the virtual IP address can be any of a range of addresses that are valid and unused for the given subnet mask.

In addition to this main advantage, the Local model also has the following advantages:

- Usually provides the highest bandwidth and lowest latency.
- Simplified administration.
- Device configuration for forwarding syslogs and SNMP notifications is much easier.

The Local model has the following disadvantages:

- Being co-located in the same data center exposes them to site-wide failures, including power outages and natural disasters.
- Increased exposure to catastrophic site impacts will complicate business continuity planning and may increase disaster-recovery insurance costs.

Related Topics

Planning HA Deployments, on page 9
Using the Campus Model, on page 11
Using the Remote Model, on page 11

Using the Campus Model

The Campus model assumes that the deploying organization is located at one or more geographical sites within a city, state or province, so that it has more than one location forming a “campus”. This model has the following advantages:

- Usually provides bandwidth and latency comparable to the Local model, and better than the Remote model.
- Is simpler to administer than the Remote model.

The Campus model has the following disadvantages:

- More complicated to administer than the Local model.
- Does not permit use of a virtual IP address as the single management address for the system, so it requires more device configuration (see “What If I Cannot Use Virtual IP Addressing?” in Related Topics).
- May provide lower bandwidth and higher latency than the Local model. This can affect HA reliability and may require administrative intervention to remedy (see “Network Throughput Restrictions on HA” in Related Topics).
- While not located at the same site, it will still be exposed to city, state, or province-wide disasters. This may complicate business continuity planning and increase disaster-recovery costs.

Related Topics

Planning HA Deployments, on page 9
Network Throughput Restrictions on HA, on page 9
Using the Local Model, on page 10
Using the Remote Model, on page 11
What If I Cannot Use Virtual IP Addressing?, on page 12

Using the Remote Model

The Remote model assumes that the deploying organization has more than one site or campus, and that these locations communicate across geographical boundaries by WAN links. It has the following advantages:

- Least likely to be affected by natural disasters. This is usually the least complex and costly model with respect to business continuity and disaster recovery.
- May reduce business insurance costs.

The Remote model has the following disadvantages:

- More complicated to administer than the Local or Campus models.

- Does not permit use of a virtual IP address as the single management address for the system, so it requires more device configuration (see “What If I Cannot Use Virtual IP Addressing?” in Related Topics).
- Usually provides lower bandwidth and higher latency than the other two models. This can affect HA reliability and may require administrative intervention to remedy (see “Network Throughput Restrictions on HA” in Related Topics).

Related Topics

Planning HA Deployments, on page 9
Network Throughput Restrictions on HA, on page 9
Using the Local Model, on page 10
Using the Campus Model, on page 11
What If I Cannot Use Virtual IP Addressing?, on page 12

What If I Cannot Use Virtual IP Addressing?

Depending on the deployment model you choose, not configuring a virtual IP address may result in the administrator having to perform additional steps in order to ensure that syslogs and SNMP notifications are forwarded to the secondary server in case of a failover. The usual method is to configure the devices to forward all syslogs and traps to both servers, usually via forwarding them to a given subnet or range of IP addresses that includes both the primary and secondary server.

This configuration work should be done at the same time HA is being set up: that is, after the secondary server is installed but before HA registration on the primary server. It must be completed before a failover so that the chance of losing data is eliminated or reduced. Not using a virtual IP address entails no change to the secondary server install procedure. The primary and secondary servers still need to be provisioned with their individual IP addresses, as normal.

This workaround is not available to you if you want to use HA with Operations Center. Enabling virtual IP addressing is a firm requirement in this case (see “Enable HA for Operations Center”).

Related Topics

Using Virtual IP Addressing With HA, on page 6
Planning HA Deployments, on page 9
Network Throughput Restrictions on HA, on page 9
Using the Campus Model, on page 11
Using the Remote Model, on page 11
Enable HA for Operations Center, on page 13

Automatic Versus Manual Failover

Configuring HA for automatic failover reduces the need for network administrators to manage HA. It also reduces the time taken to respond to the conditions that provoked the failover, since it brings up the secondary server automatically.

However, we recommend that the system be configured for Manual failover under most conditions. Following this recommendation ensures that Prime Infrastructure does not go into a state where it keeps failing over to the secondary server due to intermittent network outages. This scenario is most likely when deploying HA using the Remote model. This model is often especially susceptible to extreme variations in bandwidth and latency (see “Planning HA Deployments” and “Network Throughput Restrictions on HA” in Related Topics)

If the failover type is set to Automatic and the network connection goes down or the network link between the primary and secondary servers becomes unreachable, there is also a small possibility that both the primary and secondary servers will become active at the same time. We refer to this as the "split brain scenario".

To prevent this, the primary server always checks to see if the secondary server is Active. As soon as the network connection or link is restored and the primary is able to reach the secondary again, the primary server checks the secondary server's state. If the secondary state is Active, then the primary server goes down on its own. Users can then trigger a normal, manual failback to the primary server.

Note that this scenario *only* occurs when the primary HA server is configured for Automatic failover. Configuring the primary server for Manual failover eliminates the possibility of this scenario. This is another reason why we recommend Manual failover configuration.

Automatic failover is especially ill-advised for larger enterprises. If a particular HA deployment chooses to go with Automatic failover anyway, an administrator may be forced to choose between the data that was newly added to the primary or to the secondary. This means, essentially, that there is a possibility of data loss whenever a split-brain scenario occurs. For help dealing with this issue, see "How to Recover From Split-Brain Scenario" in Related Topics.

To ensure that HA is managed correctly, Cisco recommends that Prime Infrastructure administrators always confirm the overall health of the HA deployment before initiating failover or failback, including:

- The current state of the primary.
- The current state of the secondary.
- The current state of connectivity between the two servers.

Related Topics

Planning HA Deployments, on page 9
Network Throughput Restrictions on HA, on page 9
How to Trigger Failback, on page 31
How to Recover From Split-Brain Scenario, on page 41
Enable HA for Operations Center, on page 13

Enable HA for Operations Center

Operations Center is compatible with Prime Infrastructure's High Availability (HA) framework. You can easily enable HA for Operations Center by setting up primary and secondary Operations Center servers, much as you do when implementing HA for normal Prime Infrastructure server instances that you manage using Operations Center.

No additional Operations Center license is required on the secondary server. HA for Operations Center supports both manual and automatic failover. In the event of a failover, when the secondary Operations Center server becomes active, all managed instances from the primary Operations Center server are automatically carried over to the secondary server. You can enable HA on your primary Operations Center server whether the primary is new or already running Operations Center.

Enabling HA for Operations Center is optional. However, if you choose to enable HA for Operations Center, you may also enable virtual IP addressing while HA registration on Operations Center. Use of virtual IP addressing also requires that the primary and secondary servers be on the same subnet.

To set up HA for Operations Center using virtual IP, follow this workflow:

1. Determine the virtual IP address you will use for both servers. For details, see "Using Virtual IP Addressing With HA" and "Before You Begin Setting Up High Availability", in Related Topics.
2. Install Prime Infrastructure on the server you plan to use as your primary Operations Center HA server.

If you already have a Prime Infrastructure server with Operations Center enabled, and wish to use it as your primary Operations Center server with HA: Remove Single Sign On (SSO) servers from the Operations Center instance and all the Prime Infrastructure instances managed by that Operations Center server. You can easily do this by selecting **Administration > Users > Users, Roles & AAA > SSO Servers** and then using the **Delete SSO Server(s)** command.

3. Install the secondary server and configure it for use with HA. For details, see “How to Install the HA Secondary Server ” in Related Topics.
4. Register the secondary server on the primary, specifying that you want to Enable virtual IP and supplying the virtual IP address you selected. Logout from the Server and login back with the virtual IP. For details, see “How to Register HA on the Primary Server” in Related Topics.
5. If this is a new primary HA server: Apply the Operations Center license file to the primary server to transform it into an Operations Center instance. For details, see “Activate Your Operations Center License”.
6. Setup the virtual IP address as the SSO server on the primary server, specifying the virtual IP address as the IP address for the SSO server.For details, see “Enable SSO for Operations Center” in Related Topics.

Note By default TOFU is enabled in the primary server and if no CA certificate is deployed in primary or secondary, then after failover, delete the Virtual IP TOFU from the PI instances and secondary server. After failback repeat the same from primary server. To remove TOFU for Virtual IP from SSO (primary) client server:

```
ncs certvalidation tofu-certs deletecert host <virtual ip>
```

Note Post the upgrade of Prime Infrastructure Operations center to 3.7, if you have a self-signed certificate then before adding the Prime Infrastructure instance to Operations center, you must remove the VIP from TOFU check.

7. Repeat the virtual IP SSO server setup on all instances of Prime Infrastructure that will be managed by the primary Operations Center server. Make sure you have deleted any old SSO configuration and launch PI server with its own IP.
8. Log out of all Prime Infrastructure instances and log back into the Operations Center instance, using the virtual IP address as the Operations Center server IP.
9. If this is a new primary HA server: Add Prime Infrastructure instances to the Operations Center server, as explained in “Add Cisco Prime Infrastructure Instances to Operations Center” in the Related Topics.

   For more information, see "Activate Your Operations Center License" in Related Topics.

Note It is recommended to use either the host-name or the IP address uniformly for both managed servers and SSO configuration. Including both IP address and host-name may cause unexpected behaviour in SSO when cross launching from OPC to managed PI's.

To set up HA for Operations Center without using virtual IP, follow this workflow:

1. Install Prime Infrastructure on the server you plan to use as your primary Operations Center HA server.

   If you already have a Prime Infrastructure server with Operations Center enabled, and wish to use it as your primary Operations Center server with HA: Remove Single Sign On (SSO) servers from the Operations Center instance and all the Prime Infrastructure instances managed by that Operations Center server. You

can easily do this by selecting **Administration > Users > Users, Roles & AAA > SSO Servers** and then using the **Delete SSO Server(s)** command.

2. Install the secondary server and configure it for use with HA. For details, see “How to Install the HA Secondary Server ” in Related Topics.
3. Register the secondary server on the primary.
4. If this is a new primary HA server: Apply the Operations Center license file to the primary server to transform it into an Operations Center instance. For details, see “Activate Your Operations Center License”.
5. Repeat the primary Server IP address setup on all instances of Prime Infrastructure that will be managed by the primary Operations Center server.
6. Log out of all Prime Infrastructure instances and log back into the Operations Center instance, using the Primary IP address as the Operations Center server IP.
7. If this is a new primary HA server: Add Prime Infrastructure instances to the Operations Center server, as explained in “Add Cisco Prime Infrastructure Instances to Operations Center” in the Related Topics.

   For more information, see "Activate Your Operations Center License" in Related Topics.

Related Topics

Using Virtual IP Addressing With HA, on page 6
Before You Begin Setting Up High Availability, on page 16
How to Install the HA Secondary Server, on page 17
How to Register HA on the Primary Server, on page 17
Activate Your Operations Center License
Add Cisco Prime Infrastructure Instances to Operations Center

Set Up High Availability

To use the HA capabilities in Prime Infrastructure, you must:

1. Ensure you have the information and settings you need to enable HA. For details, see “ Before You Begin Setting Up High Availability ” in Related Topics.
2. Install a second Prime Infrastructure server, and configure it to act as your secondary HA server. For details, see “ How to Install the HA Secondary Server ”.
3. Configure High Availability mode on the primary server, specifying the installed secondary server as the HA fallback server. For details, see “ How to Register HA on the Primary Server ”.

Related Topics

How High Availability Works, on page 1
Planning HA Deployments, on page 9
Enable HA for Operations Center, on page 13
Before You Begin Setting Up High Availability, on page 16
How to Install the HA Secondary Server, on page 17
How to Register HA on the Primary Server, on page 17
What Happens During HA Registration, on page 22
Monitor High Availability, on page 30
Access the Health Monitor Web Page, on page 30
High Availability Reference Information, on page 42

Before You Begin Setting Up High Availability

Before you begin, you will need:

- The Prime Infrastructure installation software. You will use this software to create the secondary HA server. The version of this software must match the version of Prime Infrastructure installed on your primary server. You can use the CLI **show version**command to verify the current version of the primary server software.
- If you have applied patches to your primary server, you must also patch the secondary server to the same level. Choose **Administration > Licenses and Software Updates > Software Update** to see a list of the patches applied to the primary server. Then, after setting up High Availability, follow the procedure in "How to Patch Paired High Availability Servers" to patch the secondary server to the same level as the primary server.
- A secondary server with hardware and software specifications that match or exceed the requirements for your primary server. For example: If your primary server was installed as a Prime Infrastructure Standard size OVA, your secondary server must also be installed as a Standard server, and must meet or exceed all requirements given for Standard size servers in the Cisco Prime Infastructure Quick Start Guide.
- The IP address or host name of the secondary server. You will need these when configuring HA on the primary server.
- If you plan to use virtual IP addressing: The virtual IPv4 and IPv6 IP address you want to use as the virtual IP for both HA servers. This is required only if you plan to use the virtual IP feature (see "Using Virtual IP Addressing with HA" in Related Topics). Note that virtual IP addressing requires that both HA servers are on the same subnet. You must use virtual IP addressing if you plan to use HA with Operations Center (see "Enable HA for Operations Center" in Related Topics)
- An authentication key of any length. It must contain at least three of the following types of characters: lowercase letters, uppercase letters, digits and special characters. You will enter this authentication key when you install the secondary server. The HA implementation uses this key to authenticate communications between the primary and secondary servers. Administrators also use the key to configure HA in the primary server, and to log on to the secondary server's Health Monitor page to monitor the HA implementation and troubleshoot problems with it.
- A Prime Infrastructure user ID with Administrator privileges on the primary server.
- A valid email address to which HA state-change notifications can be set. Prime Infrastructure will send email notifications for the following changes: HA registration, failure, failover, and failback.
- For acceptable results: Latency of 220 milliseconds or less, and net throughput of 86 Mbps or more, over the link between the primary and secondary servers. Failure to provide at least this link quality will interfere with data replication and may lead to HA failures. For advice on the range of acceptable performance requirements, see "Network Throughput Restrictions on HA".
- If there is a firewall configured between the primary and the secondary servers, ensure that the firewall permits incoming and outgoing TCP/UDP on the following ports:
  - 8082: Used by the Health Monitor porcess to exchange hearbeat messages.
  - 1522: Used by Oracle to synchronize data.
  - 8085: Used by the Health Monitor process to check network bandwidth speed between Primary and Secondary servers when the user executes readiness test under High Availability.
- If you plan on using Operations Center with an HA implementation of Prime Infrastructure: Ensure that all of your HA-enabled Prime Infrastructure servers (both primary and secondary) have fully resolved host names.

  For more information, see Cisco Prime Infastructure Quick Start Guide

Related Topics

Set Up High Availability, on page 15
Using Virtual IP Addressing With HA, on page 6
Enable HA for Operations Center, on page 13
Network Throughput Restrictions on HA, on page 9

How to Install the HA Secondary Server

If your primary server has been patched, be sure to apply the same patches to your secondary server after installation and before registering HA on the primary server.

Make sure you have already decided on an authentication key, as explained in "Before You Begin Setting Up High Availability" in Related Topics.

Step 1 Begin installing the Prime Infrastructure server software on your secondary server just as you would for a primary server. For instructions on installing the server, see the Cisco Prime Infrastructure Quick Start Guide.

Step 2 During the installation, you will be prompted as follows:

Will this server be used as a secondary for HA? (yes/no)

Enter **yes** at the prompt.

Step 3 You will then be prompted for the HA authentication key, as follows:

Enter Authentication Key:

Enter the authentication key at the prompt. Enter it again at the confirmation prompt.

Step 4 When the secondary server is installed:

a) Use the CLI **show version** command on both servers, to verify that they are at the same version and patch level (see "Check Prime Infrastructure Version and Patch Status").
b) Run the ncs status command to verify that all processes are up and running on the secondary server (see "Check Prime Infrastructure Server Status").
c) Register HA on the primary server (see "How to Register HA on the Primary Server").

Related Topics

Set Up High Availability, on page 15
Before You Begin Setting Up High Availability, on page 16
Check Prime Infrastructure Version and Patch Status
Check Prime Infrastructure Server Status
How to Register HA on the Primary Server, on page 17

How to Register HA on the Primary Server

To enable HA, you must register HA on the primary server. The primary server needs no configuration during installation in order to participate in the HA configuration. The primary needs to have only the following information:

- The IP address or host name of the secondary HA server you have already installed and configured (see "How to Install the HA Secondary Server" in Related Topics)

- The authentication key you set during installation of the secondary server.
- One or more email addresses, to which notifications will be sent.
- The Failover Type (see “Automatic Versus Manual Failover”).

If you plan to use virtual IP addressing (see “Using Virtual IP Addressing With HA”), you will also need to:

- Select the **Enable Virtual IP** checkbox.
- Specify the IPv4 virtual IP address to be shared by the primary and secondary HA servers. You may also specify an IPv6 virtual IP address, although this is not required.

The following steps explain how to register HA on the primary server. You follow these same steps when re-registering HA.

Step 1 Log in to Prime Infrastructure with a user ID and password that has administrator privileges.

Step 2 From the menu, select **Administration** > **Settings** > **High Availability**. Prime Infrastructure displays the HA status page.

Step 3 Select **HA Configuration** and then complete the fields as follows:

a. **Secondary Server**: Enter the IP address or the host name of the secondary server.

Note We always recommend to use DNS server for resolving the host name to IP address. If you are using the "/etc/hosts" file instead of DNS server, you should enter the secondary IP address instead of host name.

b. **Authentication Key**: Enter the authentication key password you set during the secondary server installation.

c. **Email Address**: Enter the address (or comma-separated list of addresses) to which notification about HA state changes should be mailed. If you have already configured email notifications using the Mail Server Configuration page (see “Configure Email Server Settings”), the email addresses you enter here will be appended to the list of addresses already configured for the mail server.

d. **Failover Type**: Select either **Manual** or **Automatic**. We recommend that you select **Manual**.

Step 4 If you are using the virtual IP feature: Select the **Enable Virtual IP** checkbox, then complete the additional fields as follows:

a. **IPV4 Virtual IP**: Enter the virtual IPv4 address you want both HA servers to use.

b. **IPV6 Virtual IP**: (Optional) Enter the IPv6 address you want both HA servers to use.

Note that virtual IP addressing will **not** work unless both servers are on the same subnet. You should not use IPV6 address block fe80, it is been reserved for link-local unicast addressing.

Step 5 Click **Check Readiness** to ensure if the HA related environmental parameters are ready for the configuration.

For more details, see "Check Readiness for HA Registration/Configuration".

Step 6 Click **Register** to view the Milestone progress bar, to check the 100% completion of Pre-HA Registration, Database Replication and Post HA Registration as shown below. Prime Infrastructure initiates the HA registration process. When

registration completes successfully, **Configuration Mode** will display the value **Primary Active**.

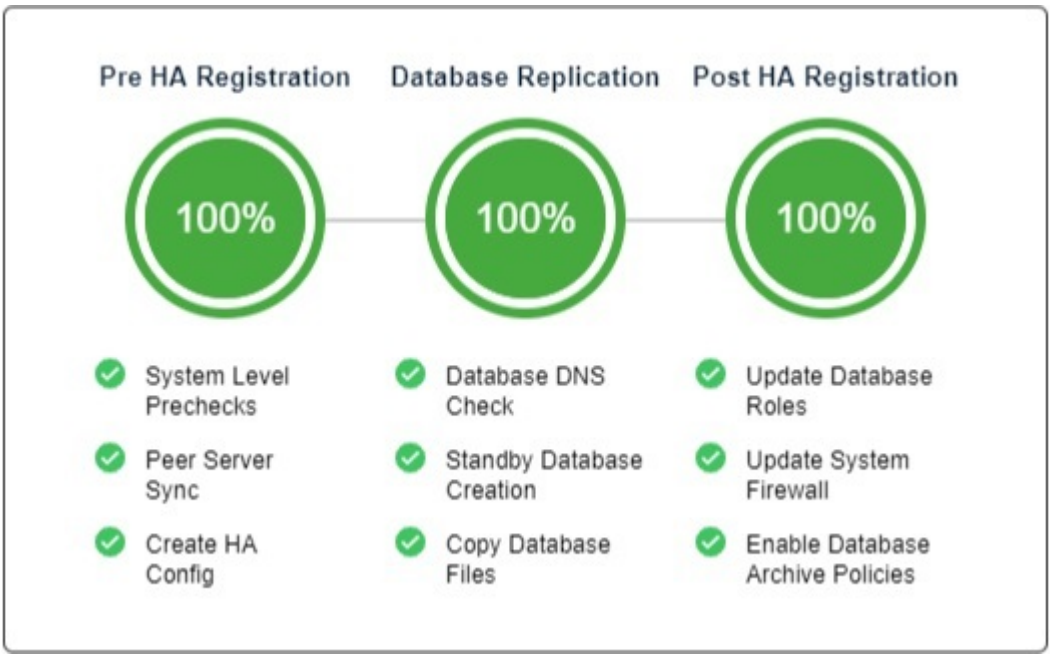

For more information, see Configure Email Server Settings.

Related Topics

How to Install the HA Secondary Server, on page 17
Automatic Versus Manual Failover, on page 12
Using Virtual IP Addressing With HA, on page 6
Before You Begin Setting Up High Availability, on page 16
What Happens During HA Registration, on page 22
Set Up High Availability, on page 15
Check Readiness for HA Registration/Configuration, on page 19

Check Readiness for HA Registration/Configuration

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

An approximate of 15 checks are run in the system to ensure the HA configuration completion 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.

To check readiness for HA configuration, follow these steps:

Step 1 Log in to Prime Infrastructure with a user ID and password that has administrator privileges.

Step 2 From the menu, select **Administration** > **Settings** > **High Availability**. Prime Infrastructure displays the HA status page.

Step 3 Select **HA Configuration**.

Step 4 Provide the secondary server IP address in the **Secondary Server** field and secondary Authentication Key **Authentication Key** field .

Step 5 Click **Check Readiness**.

A pop up window with the system specifications and other parameters will be displayed. The screen will show the Checklist Item name, Status, Impact and Recommendation details.

Below, is the list of checklist test name and the description displayed for Check Readiness:

Table 1: Checklist name and description

Checklist Test Name	Test Description
SYSTEM - Check CPU Count	This validates the CPU count in primary and secondary server. The CPU count in primary server can be less than or equal to the secondary server.
DATABASE - LISTENER STATUS	This checks if the database listeners are up and running in both primary and secondary server. If there is a failure, the test will restart and report the status. This checks if all the wcs instances exist under oracle "listener.ora" file. This is executed in both primary and secondary server.
DATABASE - CHECK MEMORY TARGET	This checks for "/dev/shm" database memory target size for HA setup.
DATABASE - CHECK LISTENER CONFIG CORRUPTION	This checks for all the database instances exist under database listener configuration. This is executed in both primary and secondary server.
SYSTEM - HEALTH MONITOR STATUS	This checks whether the health monitor process is running in both primary and secondary server.
SYSTEM - CHECK DISK IOPS	This validates the disk IOPS in both primary and secondary server. The minimum expected disk IOPS is 200 MBps.
NETWORK - CHECK FIREWALL FOR DATABASE PORT ACCESSIBILITY	This 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 iptables list.
NETWORK - CHECK NETWORK INTERFACE BANDWIDTH	This checks if the eth0 interface speed matches the recommended speed of 100 Mbps in both primary and secondary sever. This test will not measure network bandwidth by transmitting data between primary and secondary server.

NETWORK - CHECK NETWORK BANDWIDTH SPEED	This checks if the network bandwidth speed matches the recommended speed of 100 Mbps in both primary and secondary sever. This test will measure network bandwidth by transmitting data between primary and secondary server.
DATABASE - CHECK ONLINE STATUS	This checks if the database files status is online and accessible in both primary and secondary server.
DATABASE - CHECK TNS CONFIG CORRUPTION	This validates if the tnsping is successful in both primary and secondary server.
DATABASE - TNS REACHABILITY STATUS	This checks if all the wcs instances exist under oracle "listener.ora" file. This is executable in both primary and secondary server.
DATABASE - VALIDATE STANDBY DATABASE INSTANCE	This validates if the standby database instance (stbywcs) is available in both primary and secondary server.
SYSTEM - CHECK RAM SIZE	This checks if the disk size of primary server less than or equal to secondary server.
SYSTEM - CHECK SERVER PING REACHABILITY	This ensures that the primary server can run ping check with the remote (secondary) server.

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

Note The validation failback and failover events during Check Readiness will be sent to the Alarms and Events page; whereas, the registration failure event will not be present in the Alarms and Evens page.

Check High Availability Status

You can check on the status of the High Availability enabled on a Prime Infrastructure server.

Step 1 Open a CLI session with the Prime Infrastructure server (see How to Connect Via CLI).

Step 2 Enter the following command to display the current status of Prime Infrastructure HA processes:

PIServer/admin# **ncs ha status**

Related Topics

Set Up High Availability, on page 15

What Happens During HA Registration

Once you finish entering configuration information and click the Save button on the HA Configuration page, the primary and secondary HA servers will register with each other and begin copying all database and configuration data from the primary to the secondary server.

The time required to complete the copying is a function of the amount of database and configuration data being replicated and the available bandwidth on the network link between the two servers. The bigger the data and the slower the link, the longer the replication will take. For a relatively fresh server (in operation for a few days), with 100 devices and a 1 GB-per-second link, copying will take approximately 25 minutes.

During HA registration, the primary and secondary server state will go through the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: HA Not Configured	From: HA Not Configured
To: HA Initializing	To: HA Initializing
To: Primary Active	To: Secondary Syncing

You can view these state changes on the HA Status page for the primary server, or the Health Monitor web pages for either of the two servers. If you are using the HA Status page, click **Refresh** to view progress. Once the data is fully synchronized, the HA Status page will be updated to show the current state as "Primary Active", as shown in the following figure.

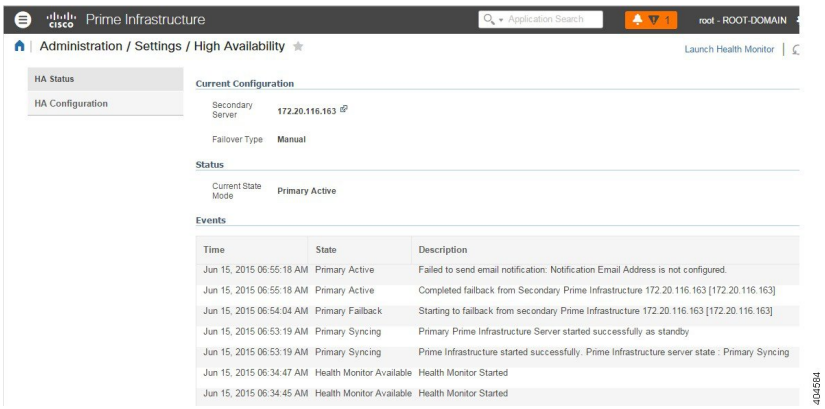

After registration is initiated, Prime Infrastructure initiates synchronization between the primary and the secondary HA servers. The synchronization should not have any impact on user activity, although users may observe slow system response until the synchronization is complete. The length of the synchronization is a function of the total database size and, is handled at the Oracle database level by the Oracle RMAN and Data Guard Broker processes. There is no impact on the execution of user- or system-related activity during the sync.

During registration, Prime Infrastructure performs a full database replication to the secondary server. All processes on the secondary server will be running, but the server itself will be in passive mode. If you execute the Prime Infrastructure CLI command **ncs status** on the secondary server while the secondary server is in the "Secondary Syncing" state, the command output will show all processes as running.

Related Topics

How High Availability Works, on page 1
Planning HA Deployments, on page 9

Set Up High Availability, on page 15

How to Patch HA Servers

You can download and install UBF patches for your HA servers in one of the following ways, depending on your circumstances:

- Install the patch on HA servers that are not currently paired. Cisco recommends this method if you have not already set up HA for Prime Infrastructure.
- Install the patch on existing paired HA servers

For details on each method, see the Related Topics.

Related Topics

How to Patch New HA Servers, on page 23
How to Patch Paired HA Servers, on page 25

How to Patch New HA Servers

If you are setting up a new Prime Infrastructure High Availability (HA) implementation and your new servers are not at the same patch level, follow the steps below to install patches on both servers and bring them to the same patch level.

Step 1 Download the patch and install it on the primary server:

a) Point your browser to the software patches listing for Cisco Prime Infrastructure (see Related Topics) .
b) Click the **Download** button for the patch file you need to install (the file name ends with a UBF file extension), and save the file locally.
c) Log in to the primary server using an ID with administrator privileges and choose **Administration > Licenses and Software Updates > Software Update**.
d) Click the **Upload** link at the top of the page .
e) Use one of the following options to upload the UBF file.

1. Upload from local computer
   - Click the **Upload from local computer** radio button in the **Upload Update** window.
   - Click **Browse**, navigate to the file, and click **OK**. After the successful upload, the software will appear under the **Files** tab.
2. Copy from server's local disk
   - Click the **Copy from server's local disk** radio button in the **Upload Update** window.
   - Click **Select** , select the UBF file from the **Select file from local disk** pop-up and click **Select**. After the successful upload, the software will appear under the **Files** tab.

f) When the upload is complete: On the Software Upload page, verify that the Name, Published Date and Description of the patch file are correct.
g) Select the patch file and click **Install**.

h) Click **Yes** in the warning pop-up. When the installation is complete, the server will restart automatically. The restart typically takes 15 to 20 minutes.
i) After the installation is complete on the primary server, verify that the Status of Updates table on the Software Update page shows “Installed” for the patch.

Step 2 Install the same patch on the secondary server:

a) Access the secondary server’s Health Monitor (HM) web page by pointing your browser to the following URL:

https://*ServerIP***:8082**

where *ServerIP* is the IP address or host name of the secondary server.

Note You will be prompted for the username and authentication key. Enter the username as 'root' and authkey and click **Login**.

Note Verify that the secondary server state displayed on the HM web page is in the *Secondary Syncing* state.

b) You will be prompted for the username and authentication key. Enter the username as 'root' and authkey and click **Login**.
c) Click the HM web page’s **Software Update** link. You will be prompted for the authentication key a second time. Enter it and click **Login** again.
d) Click the **Upload** link at the top of the page.
e) Use one of the following options to upload the UBF file.

1. Upload from local computer
   - Click the **Upload from local computer** radio button in the **Upload Update** window.
   - Click **Browse**, navigate to the file, and click **OK**. After the successful upload, the software will appear under the **Files** tab.
2. Copy from server's local disk
   - Click the **Copy from server's local disk** radio button in the **Upload Update** window.
   - Click **Select** , select the UBF file from the **Select file from local disk** pop-up and click **Select**. After the successful upload, the software will appear under the **Files** tab.

f) When the upload is complete: On the Software Upload page, confirm that the Name, Published Date and Description of the patch file are correct.
g) Select the patch file and click **Install**.
h) Click **Yes** in the warning pop-up. When the installation is complete, the server will restart automatically. The restart typically takes 15 to 20 minutes.
i) After the installation is complete on the secondary server, verify that the Status of Updates table on the Software Update page shows “Installed” for the patch.

Step 3 Verify that the patch status is the same on both servers, as follows:

a) Log in to the primary server and access its Software Update page as you did in step 1, above. The “Status” column should show “Installed” for the installed patch.
b) Access the secondary server’s Health Monitor page as you did in step 2, above. The “Status” column should show “Installed” for the installed patch

Step 4 Register the servers.

For more information, see "Software patches listing for Cisco Prime Infrastructure", "Restart Prime Infrastructure" and "Check Prime Infrastructure Server Status".

Related Topics

Set Up High Availability, on page 15
How to Register HA on the Primary Server, on page 17
How to Patch HA Servers, on page 23

How to Patch Paired HA Servers

If your current Prime Infrastructure implementation has High Availability servers that are not at the same patch level, or you have a new patch you must install on both your HA servers, follow the steps below.

Patching paired HA servers is not supported. You will receive a popup error message indicating that you cannot perform an update on Prime Infrastructure servers while HA is configured. So, you must first disconnect the primary and secondary servers before attempting to apply the patch.

1. Follow the steps in “Remove HA Via the GUI” (see Related Topics) to disconnect the primary and secondary servers.
2. Follow the steps in “How to Patch New HA Servers” to apply the patch.
3. Follow the steps in “Set Up High Availability” to restore your HA configuration.

Related Topics

Set Up High Availability, on page 15
Check High Availability Status, on page 21
Remove HA Via the GUI, on page 46
How to Patch New HA Servers, on page 23

How to Patch Paired HA Servers Set for Manual Failover

If your current Prime Infrastructure implementation has High Availability servers that are not at the same patch level, or you have a new patch you must install on both your HA servers, follow the steps below.

You must start the patch install with the primary server in “Primary Active” state and the secondary server in “Secondary Syncing” state.

Patching of primary and secondary HA servers set for manual failover takes approximately 30 minutes, and does not require failover or failback. Patching of the primary and secondary HA servers takes approximately 30 minutes. Downtime during the primary patch installation restart takes 15 to 20 minutes.

In some cases, you may receive a popup error message indicating that you cannot perform an update on Prime Infrastructure servers while HA is configured. If so, you *must* first disconnect the primary and secondary servers before attempting to apply the patch. In this case, you cannot use the steps in this procedure. Instead, be sure to:

1. Follow the steps in “Remove HA Via the GUI” (see Related Topics) to disconnect the primary and secondary servers.
2. Follow the steps in “How to Patch New HA Servers” to apply the patch.

Note You will be prompted for the username and authentication key Entered when HA was enabled. provie the username as 'root' and authentication key and click Login.

3. Follow the steps in “Set Up High Availability” to restore your HA configuration.

Step 1 Ensure that your HA implementation is enabled and ready for update:

a) Log in to the primary server using an ID with Administrator privileges.
b) Select **Administration > Settings > High Availability**, The primary server state displayed on the HA Status page should be “Primary Active”.
c) Select **HA Configuration**. The current Configuration Mode should show “HA Enabled”. We recommend that you set the Failover Type to “manual” during the patch installation.
d) Access the secondary server’s Health Monitor (HM) web page by pointing your browser to the following URL:

   https://*ServerIP***:8082**

   where *ServerIP* is the IP address or host name of the secondary server.

e) Verify that the secondary server state displayed on the HM web page is in the “Secondary Syncing” state.

Step 2 You will be prompted for the user name and authentication key entered when HA was enabled. Enter username as 'root' with authentication key and click **Login**.

Step 3 Download the UBF patch and install it on the primary server:

a) Point your browser to the software patches listing for Cisco Prime Infrastructure (see Related Topics) .
b) Click the **Download** button for the patch file you need to install (the file name ends with a UBF file extension), and save the file locally.
c) Log in to the primary server using an ID with administrator privileges and choose **Administration > Licenses and Software Updates > Software Update**.
d) Click the **Upload** link at the top of the page .
e) Use one of the following options to upload the UBF file.

   1. Upload from local computer
      - Click the **Upload from local computer** radio button in the **Upload Update** window.
      - Click **Browse**, navigate to the file, and click **OK**. After the successful upload, the software will appear under the **Files** tab.

   2. Copy from server's local disk
      - Click the **Copy from server's local disk** radio button in the **Upload Update** window.
      - Click **Select** , select the UBF file form the **Select file from local disk** pop-up and click **Select**. After the successful upload, the software will appear under the **Files** tab.

f) When the upload is complete: On the Software Upload page, verify that the Name, Published Date and Description of the patch file are correct.
g) Select the patch file and click **Install**.
h) Click **Yes** in the warning pop-up. When the installation is complete, the server will restart automatically. The restart typically takes 15 to 20 minutes.

i) After the server restart is complete on the primary server, select **Administration > Settings > High Availability**, The primary server state displayed on the HA Status page should be “Primary Active”.
j) Verify that the Status of Updates table on the Software Update page shows “Installed” for the patch.

Step 4 Install the same patch on the secondary server once patching is complete on the primary server:

a) Access the secondary server’s HM web page and login if needed.
b) Click the HM web page’s **Software Update** link. You will be prompted for the authentication key a second time. Enter it and click **Login** again.
c) Click the **Upload** link at the top of the page.
d) Use one of the following options to upload the UBF file.

1. Upload from local computer
   - Click the **Upload from local computer** radio button in the **Upload Update** window.
   - Click **Browse**, navigate to the file, and click **OK**. After the successful upload, the software will appear under the **Files** tab.
2. Copy from server's local disk
   - Click the **Copy from server's local disk** radio button in the **Upload Update** window.
   - Click **Select** , select the UBF file form the **Select file from local disk** pop-up and click **Select**. After the successful upload, the software will appear under the **Files** tab.

e) When the upload is complete: On the Software Upload page, confirm that the Name, Published Date and Description of the patch file are correct.
f) Select the patch file and click **Install**.
g) Click **Yes** in the warning pop-up. When the installation is complete, the server will restart automatically. The restart typically takes 15 to 20 minutes.
h) After the server restart is complete on the secondary server, log in to the secondary HM page (https://serverIP:8082) and verify that the secondary server state displayed on the HM web page is “Secondary Syncing”.
i) Verify that the Status of Updates table on the Software Update page shows “Installed” for the patch.

Step 5 Once the server restart is complete, verify the patch installation as follows:

a) Log in to the primary server and access its Software Update page as you did in step 2, above. The “Status” column on the Status of Updates > Update tab should show “Installed” for the patch.
b) Access the secondary server’s Software Update page as you did in step 3, above. The “Status” column on the Status of Updates > Updates tab should show “Installed” for the patch.

For more information, see

- Software patches listing for Cisco Prime Infrastructure.
- Start Prime Infrastructure
- Stop Prime Infrastructure
- Check Prime Infrastructure Server Status

Related Topics

Set Up High Availability, on page 15

Check High Availability Status, on page 21
Remove HA Via the GUI, on page 46
How to Patch New HA Servers, on page 23
How to Patch Paired HA Servers Set for Automatic Failover, on page 28

How to Patch Paired HA Servers Set for Automatic Failover

If your current Prime Infrastructure implementation has High Availability servers that are not at the same patch level, or you have a new patch you must install on both your HA servers, follow the steps below.

You must start the patch install with the primary server in “Primary Active” state and the secondary server in “Secondary Syncing” state.

Patching of primary and secondary HA servers set for automatic failover takes approximately one hour, and requires both failover and failback. Downtime during the failover and failback lasts 10 to 15 minutes.

In some cases, you may receive a popup error message indicating that you cannot perform an update on Prime Infrastructure servers while HA is configured. If so, you *must* first disconnect the primary and secondary servers before attempting to apply the patch. In this case, you cannot use the steps in this procedure. Instead, be sure to:

1. Follow the steps in “Remove HA Via the GUI” (see Related Topics) to disconnect the primary and secondary servers.
2. Follow the steps in “How to Patch New HA Servers” (see Related Topics)to apply the patch.
3. Follow the steps in “Set Up High Availability” (see Related Topics) to restore your HA configuration.

Step 1 Ensure that your HA implementation is enabled and ready for update:

a) Log in to the primary server using an ID with Administrator privileges.
b) Select **Administration > Settings > High Availability**, The primary server state displayed on the HA Status page should be “Primary Active”.
c) Select **HA Configuration**. The current Configuration Mode should show “HA Enabled”.
d) Access the secondary server’s Health Monitor (HM) web page by pointing your browser to the following URL:

   https://*ServerIP***:8082**

   where *ServerIP* is the IP address or host name of the secondary server.

e) You will be prompted for the user name and authentication key entered when HA was enabled. Enter username as 'root' with authentication key and click **Login**.
f) Verify that the secondary server state displayed on the HM web page is in the “Secondary Syncing” state.

Step 2 Download the UBF patch and install it on the primary server:

a) Point your browser to the software patches listing for Cisco Prime Infrastructure (see Related Topics) .
b) Click the **Download** button for the patch file you need to install (the file name ends with a UBF file extension), and save the file locally.
c) Log in to the primary server using an ID with administrator privileges and choose **Administration > Licenses and Software Updates > Software Update**.
d) Click the **upload** link at the top of the page and browse to the location where you saved the patch file.
e) Select the UBF file and then click **OK** to upload the file.
f) When the upload is complete: On the Software Upload page, verify that the Name, Published Date and Description of the patch file are correct.

g) Select the patch file and click **Install**.
h) Click **Yes** in the warning pop-up. Failover will be triggered and the primary server will restart automatically. Failover will take 2 to 4 minutes to complete. After the failover is complete, the secondary server will be in "Secondary Active" state.
i) After the primary server is restarted, run the **ncs status** command (see "Check Prime Infrastructure Server Status") to verify that the primary's processes have re-started. Before continuing: Access the primary server's HM web page and verify that the primary server state displayed is "Primary Synching".

Step 3 Failback to the primary using the secondary server's HM web page:

a) Access the secondary server's HM web page and login if needed.
b) Click **Failback** to initiate a failback from the secondary to the primary server. It will take 2 to 3 minutes for the operation to complete. As soon as failback completes, the secondary server will be automatically restarted in the standby mode. It will take a maximum of 15 minutes for the restart to complete, and it will be synched with the primary server.

   You can verify the restart by logging into the secondary server's HM web page and looking for the message "Prime Infrastructure stopped successfully" followed by "Prime Infrastructure started successfully."

   After failback is complete, the primary server state will change to "Primary Active"

c) Before continuing: Run the **ncs ha status** command on both the primary and secondary servers. Verify that the primary server state changes to "Primary Active" and the secondary server state is "Secondary Synching".

Step 4 Once failback completes, verify the patch installation by logging in to the primary server and accessing its Software Update page (as you did in step 2, above). The "Status" column on the Status of Updates > Update tab should show "Installed" for the patch.

Step 5 Install the same patch on the secondary server once patching is complete on the primary server:

a) Access the secondary server's HM web page and login if needed.
b) Click the HM web page's **Software Update** link. You will be prompted for the authentication key a second time. Enter it and click **Login** again.
c) Click **Upload Update File** and browse to the location where you saved the patch file.
d) Select the UBF file and then click **OK** to upload the file.
e) When the upload is complete: On the Software Upload page, confirm that the Name, Published Date and Description of the patch file are correct.
f) Select the patch file and click **Install**.
g) Click **Yes** in the warning pop-up. The server will restart automatically. The restart typically takes 15 to 20 minutes.
h) After the installation is complete on the secondary server, verify that the Status of Updates table on the Software Update page shows "Installed" for the patch.
i) After the server restart is complete on the secondary server, log in to the secondary HM page and verify that the secondary server state displayed on the HM web page is "Secondary Syncing".

Step 6 Once server restart is complete, verify the patch installation as follows:

a) Log in to the primary server and access its Software Update page as you did in step 2, above. The "Status" column on the Status of Updates > Update tab should show "Installed" for the patch.
b) Access the secondary server's Software Update page as you did in step 5, above. The "Status" column on the Status of Updates > Updates tab should show "Installed" for the patch.

   For more information, see Software patches listing for Cisco Prime Infrastructure, Stop Prime Infrastructure, Start Prime Infrastructure and Check Prime Infrastructure Server Status.

Related Topics

Set Up High Availability, on page 15
Check High Availability Status, on page 21
Remove HA Via the GUI, on page 46
How to Patch New HA Servers, on page 23
How to Patch Paired HA Servers Set for Manual Failover, on page 25

Monitor High Availability

Once you have configured HA and registered it on the primary server, most of your interactions with HA will involve accessing the server Health Monitor web page and responding to email notifications by triggering a failover or failback. These processes, as well as special situations requiring more complicated responses, are covered in the following related topics.

Related Topics

Access the Health Monitor Web Page, on page 30
How to Trigger Failover, on page 31
How to Trigger Failback, on page 31
Force Failover, on page 32
Respond to Other HA Events, on page 32

Access the Health Monitor Web Page

You can access the Health Monitor web page for the primary or secondary server at any time by pointing your browser to the following URL:

https://**Server**:8082

where **Server** is the IP address or host name of the primary or secondary server whose Health Monitor web page you want to see.

Note You will be prompted for the username and authentication key .Enter the username as 'root' and authentication key and click **Login**.

You can also access the Health Monitor web page for the currently active server by logging in to Prime Infrastructure, selecting **Administration > Settings > High Availability**, and then clicking the **Launch Health Monitor** link at the top right of the HA Status page.

Related Topics

Monitor High Availability, on page 30
How to Trigger Failover, on page 31
How to Trigger Failback, on page 31
Force Failover, on page 32

How to Trigger Failover

Failover is the process of activating the secondary server in response to a detected failure on the primary.

Health Monitor (HM) detects failure conditions using the heartbeat messages that the two servers exchange. If the primary server is not responsive to three consecutive heartbeat messages from the secondary, it is considered to have failed. During the health check, HM also checks the application process status and database health; if there is no proper response to these checks, these are also treated as having failed.

The HA system takes approximately 10 to 15 seconds to detect a process failure on the primary server and initiate a failover. If the secondary server is unable to reach the primary server due to a network issue, it might take more time to initiate a failover. In addition, it may take additional time for the application processes on the secondary server to be fully operational.

As soon as HM detects the failure, it sends an email notification. The email includes the failure status along with a link to the secondary server's Health Monitor web page.

If HA is currently configured for automatic failover, the secondary server will activate automatically and there is no action you need to perform.

If HA is currently configured for manual failover, you must trigger the failover as mentioned in the below procedure.

Failover should be considered temporary. The failed primary Prime Infrastructure instance should be restored to normal as soon as possible, and failback should be re-initiated.

Step 1 Access the secondary server's Health Monitor web page using the web link given in the email notification, or using the steps in “Accessing the Health Monitor Web Page”.

Step 2 Trigger the failover by clicking the **Failover** button.

Related Topics

How High Availability Works, on page 1
How to Trigger Failback, on page 31
Monitor High Availability, on page 30
How to Register HA on the Primary Server, on page 17
Access the Health Monitor Web Page, on page 30

How to Trigger Failback

Failback is the process of re-activating the primary server once it is back online. It also transfers Active status from the secondary server to the primary, and stops active network monitoring processes on the secondary.

During failback, the secondary server is available except during the period when processes are re-started on the secondary. Both servers’ Health Monitor web pages are accessible for monitoring the progress of the failback. Additionally, users can also connect to the secondary server to access all normal functionality, except for these caveats:

- Do not initiate configuration or provisioning activity while the failback is in progress.
- Be aware that, after a successful failback, the secondary server will go into passive (“Secondary Syncing”) mode and control will switch over to the primary server. During this process, Prime Infrastructure will be inaccessible to the users for a few moments.

You must always trigger failback manually, as follows:

Step 1 Access the secondary server's Health Monitor web page using the link given in the email notification, or using the steps in “Accessing the Health Monitor Web Page”.

Step 2 Trigger the failback by clicking the **Failback** button.

The secondary server is automatically restarted in the standby mode after the failback and is automatically synced with the primary server. The primary server will now be the available Prime Infrastructure server.

Related Topics

How High Availability Works, on page 1
How to Trigger Failover, on page 31
Force Failover, on page 32
Monitor High Availability, on page 30
Access the Health Monitor Web Page, on page 30

Force Failover

A forced failover is the process of making the secondary server active while the primary server is still up. You will want to use this option when, for example, you want to test that your HA setup is fully functional.

Forced failover is available to you only when the primary is active, the secondary is in the “Secondary syncing” state, and all processes are running on both servers. Forced failover is disabled when the primary server is down. In this case, only the normal Failover is enabled.

Once the forced failover completes, the secondary server will be active and the primary will restart in standby automatically. You can return to an active primary server and standby secondary server by triggering a normal failback.

Step 1 Access the secondary server's Health Monitor web page using the steps in “Accessing the Health Monitor Web Page”.

Step 2 Trigger the forced failover by clicking the **Force Failover** button. The forced failover will complete in 2 to 3 minutes.

Related Topics

How High Availability Works, on page 1
How to Trigger Failover, on page 31
How to Trigger Failback, on page 31
Monitor High Availability, on page 30
How to Register HA on the Primary Server, on page 17
Access the Health Monitor Web Page, on page 30

Respond to Other HA Events

All the HA related events are displayed on the HA Status page, the Health Monitor web pages, and under the Prime Infrastructure Alarms and Events page. Most events require no response from you other than triggering failover and failback. A few events are more complex, as explained in the related topics.

Related Topics

HA Registration Fails, on page 33
Network is Down (Automatic Failover), on page 33
Network is Down (Manual Failover), on page 34
Process Restart Fails (Manual Failover), on page 37
Primary Server Restarts During Sync (Manual Failover), on page 38
Secondary Server Restarts During Sync, on page 38
Both HA Servers Are Down, on page 38
Both HA Servers Are Down and the Secondary Will Not Restart, on page 40
Replace Primary MSEs, on page 66
How to Recover From Split-Brain Scenario, on page 41

HA Registration Fails

If HA registration fails, you will see the following HA state-change transitions for each server (instead of those detailed in "What Happens During HA Registration":

Primary HA State Transitions...	Secondary HA State Transitions...
From: HA Initializing	From: HA Initializing
To: HA Not Configured	To: HA Not Configured

To recover from failed HA registration, follow the steps below.

Step 1 Use ping and other tools to check the network connection between the two Prime Infrastructure servers. Confirm that the secondary server is reachable from the primary, and vice versa.

Step 2 Check that the gateway, subnet mask, virtual IP address (if configured), server hostname, DNS, NTP settings are all correct.

Step 3 Check that the configured DNS and NTP servers are reachable from the primary and secondary servers, and that both are responding without latency or other network-specific issues.

Step 4 Check that all Prime Infrastructure licenses are correctly configured.

Step 5 Once you have remedied any connectivity or setting issues, try the steps in "How to Register High Availability on the Primary Server" again in related topics.

Related Topics

Respond to Other HA Events, on page 32
What Happens During HA Registration, on page 22
How to Register HA on the Primary Server, on page 17

Network is Down (Automatic Failover)

If there is a loss of network connectivity between the two Prime Infrastructure servers, you will see the following HA state-change transitions for each server, assuming that the Failover Type is set to "Automatic":

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Lost Secondary	To: Secondary Lost Primary
To: Primary Lost Secondary	To: Secondary Failover
To: Primary Lost Secondary	To: Secondary Active

You will get an email notification that the secondary is active.

Step 1 Check on and restore network connectivity between the two servers. Once network connectivity is restored and the primary server can detect that the secondary is active, all services on the primary will be restarted and made passive automatically. You will see the following state changes:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Lost Secondary	From: Secondary Active
To: Primary Failover	To: Secondary Active
To: Primary Syncing	To: Secondary Active

Step 2 Trigger a failback from the secondary to the primary. You will then see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Syncing	From: Secondary Active
To: Primary Failback	To: Secondary Failback
To: Primary Failback	To: Secondary Post Failback
To: Primary Active	To: Secondary Syncing

Related Topics

Respond to Other HA Events, on page 32
How to Trigger Failback, on page 31

Network is Down (Manual Failover)

If there is a loss of network connectivity between the two Prime Infrastructure servers, you will see the following HA state-change transitions for each server, assuming that the Failover Type is set to “Manual”:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Lost Secondary	To: Secondary Lost Primary

You will get email notifications that each server has lost the other.

Step 1 Check on and, if needed, restore the network connectivity between the two servers.

You will see the following state changes once network connectivity is restored.:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Lost Secondary	From: Secondary Lost Primary
To: Primary Active	To: Secondary Syncing

No administrator response is required.

Step 2 If network connection cannot be restored for any reason, use the HM web page for the secondary server to trigger a failover from the primary to the secondary server. You will see the following state changes:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Lost Secondary	From: Secondary Lost Primary
To: Primary Lost Secondary	To: Secondary Failover
To: Primary Failover	To: Secondary Active

You will get an email notification that the secondary server is now active.

Step 3 Check and restore network connectivity between the two servers. Once network connectivity is restored and the primary server detects that the secondary server is active, all services on the primary server will be restarted and made passive. You will see the following state changes:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Lost Secondary	From: Secondary Active
To: Primary Failover	To: Secondary Active
To: Primary Syncing	To: Secondary Active

Step 4 Trigger a failback from the secondary to the primary.

You will then see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Syncing	From: Secondary Active
To: Primary Failback	To: Secondary Failback
To: Primary Failback	To: Secondary Post Failback
To: Primary Active	To: Secondary Syncing

Related Topics

Respond to Other HA Events, on page 32
How to Trigger Failback, on page 31

Process Restart Fails (Automatic Failover)

The Prime Infrastructure Health Monitor process is responsible for attempting to restart any Prime Infrastructure server processes that have failed. Generally speaking, the current state of the primary and secondary servers should be “Primary Active” and “Secondary Syncing” at the time any such failures occur.

If HM cannot restart a critical process on the primary server, then the primary server is considered to have failed. If your currently configured Failover Type is “automatic”, you will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Uncertain	To: Secondary Lost Primary
To: Primary Failover	To: Secondary Failover
To: Primary Failover	To: Secondary Active

When this process is complete, you will get an email notification that the secondary server is now active.

Step 1 Restart the primary server and ensure that it is running. Once the primary is restarted, it will be in the state “Primary Syncing”. You will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Failover	From: Secondary Active
To: Primary Preparing for Failback	To: Secondary Active
To: Primary Syncing	To: Secondary Active

Step 2 Trigger a failback from the secondary to the primary. You will then see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Syncing	From: Secondary Active
To: Primary Failback	To: Secondary Failback
To: Primary Failback	To: Secondary Post Failback
To: Primary Active	To: Secondary Syncing

Related Topics

Respond to Other HA Events, on page 32
How to Trigger Failback, on page 31

Process Restart Fails (Manual Failover)

The Prime Infrastructure Health Monitor process is responsible for attempting to restart any Prime Infrastructure server processes that have failed. Generally speaking, the current state of the primary and secondary servers should be “Primary Active” and “Secondary Syncing” at the time any such failures occur. If HM cannot restart a critical process on the primary server, then the primary server is considered to have failed. You will receive an email notification of this failure. If your currently configured Failover Type is “Manual”, you will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Uncertain	To: Secondary Lost Primary

Step 1 Trigger on the secondary server a failover from the primary to the secondary. You will then see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Uncertain	From: Secondary Syncing
To: Primary Failover	To: Secondary Failover
To: Primary Failover	To: Secondary Active

Step 2 Restart the primary server and ensure that it is running. Once the primary server is restarted, the primary’s HA state will be “Primary Syncing”. You will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Failover	From: Secondary Active
To: Primary Preparing for Failback	To: Secondary Active
To: Primary Syncing	To: Secondary Active

Step 3 Trigger a failback from the secondary to the primary. You will then see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Syncing	From: Secondary Active
To: Primary Failback	To: Secondary Failback
To: Primary Failback	To: Secondary Post Failback
To: Primary Active	To: Secondary Syncing

Related Topics

Respond to Other HA Events, on page 32

How to Trigger Failover, on page 31
How to Trigger Failback, on page 31

Primary Server Restarts During Sync (Manual Failover)

If the primary Prime Infrastructure server is restarted while the secondary server is syncing, you will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Alone	To: Secondary Lost Primary
To: Primary Active	To: Secondary Syncing

The “Primary Alone” and “Primary Active” states occur immediately after the primary comes back online. No administrator response should be required.

Related Topics

Respond to Other HA Events, on page 32

Secondary Server Restarts During Sync

If the secondary Prime Infrastructure server is restarted while syncing with the primary server, you will see the following state transitions:

Primary HA State Transitions...	Secondary HA State Transitions...
From: Primary Active	From: Secondary Syncing
To: Primary Lost Secondary	From: Secondary Lost Primary
To: Primary Active	To: Secondary Syncing

No administrator response should be required.

Related Topics

Respond to Other HA Events, on page 32

Both HA Servers Are Down

If both the primary and secondary servers are down at the same time, you can recover by bringing them back up in the correct order, as explained in the steps below.

Step 1 Restart the secondary server and the instance of Prime Infrastructure running on it. If for some reason you cannot restart the secondary server, see “Both HA Servers Are Down and Secondary Will Not Restart” in Related Topics.

Step 2 When Prime Infrastructure is running on the secondary, access the secondary server’s Health Monitor web page. You will see the secondary server transition to the state “Secondary Lost Primary”.

Step 3 Restart the primary server and the instance of Prime Infrastructure running on it. When Prime Infrastructure is running on the primary, the primary will automatically sync with the secondary. To verify this, access the primary server's Health Monitor web page. You will see the two servers transition through the following series of HA states:

Primary HA State Transitions...	Secondary HA State Transitions...
To: Primary Lost Secondary	To: Secondary Lost Primary
To: Primary Active	To: Secondary Syncing

Related Topics

Both HA Servers Are Down and the Secondary Will Not Restart, on page 40
Access the Health Monitor Web Page, on page 30
Respond to Other HA Events, on page 32

Both HA Servers Are Powered Down

If both the primary and secondary servers are powered down at the same time, you can recover by bringing them back up in the correct order, as explained in the steps below.

Step 1 Power on the secondary server and the instance of Prime Infrastructure running on it.

The secondary HA restart will fail at this stage because the primary is not reachable. However, the secondary Health Monitor process will be running with an error.

Step 2 When Prime Infrastructure is running on the secondary, access the secondary server's Health Monitor web page. You will see the secondary server transition to the state "Secondary Lost Primary".

Step 3 Power on the primary server and the instance of Prime Infrastructure running on it.

Step 4 When Prime Infrastructure is running on the primary, the primary will automatically sync with the secondary. To verify this, access the primary server's Health Monitor web page. You will see the two servers transition through the following series of HA states:

Primary HA State Transitions...	Secondary HA State Transitions...
To: Primary Lost Secondary	To: Secondary Lost Primary
To: Primary Active	To: Secondary Syncing

Step 5 Restart the secondary server and the instance of Prime Infrastructure running on it. This is required because not all processes will be running on the secondary at this point.

If for some reason you cannot restart the secondary server, see "Both HA Servers Are Down and Secondary Will Not Restart" in Related Topics.

Step 6 When Prime Infrastructure finishes restarting on the secondary server, all processes should be running. Verify this by running the ncs status command (see "Check Prime Infrastructure Server Status" in Related Topics).

Related Topics

Both HA Servers Are Down and the Secondary Will Not Restart, on page 40

Access the Health Monitor Web Page, on page 30
Respond to Other HA Events, on page 32
Check Prime Infrastructure Server Status

Both HA Servers Are Down and the Secondary Will Not Restart

If both HA servers are down at the same time and the secondary will not restart, you will need to remove the HA configuration from the primary server in order to use it as a standalone until you can replace or restore the secondary server.

The following steps assume that you have already tried and failed to restart the secondary server.

Step 1 Attempt to restart the primary instance of Prime Infrastructure. If the primary is able to restart at all, the restart will abort with an error message indicating that you must remove the HA configuration.

Step 2 Open a CLI session with the primary Prime Infrastructure server (see How to Connect Via CLI).

Step 3 Enter the following command to remove the HA configuration on the primary server:

PIServer/admin# **ncs ha remove**

Step 4 You will be prompted to confirm that you want to remove the HA configuration. Answer **Y** to the prompt.

You should now be able to restart the primary instance of Prime Infrastructure without the error message and use it as a standalone.

When you are able to restore or replace the secondary server, proceed as explained in "How to Register High Availability on the Primary Server" in Related Topics.

Related Topics

Access the Health Monitor Web Page, on page 30
How to Register HA on the Primary Server, on page 17
Remove HA Via the CLI, on page 47
Respond to Other HA Events, on page 32

How to Replace the Primary Server

Under normal circumstances, the state of your primary and secondary servers will be "Primary Active" and "Secondary Syncing", respectively. If the primary server fails for any reason, a failover to the secondary will take place, either automatically or manually.

You may find that restoring full HA access requires you to reinstall the primary server using new hardware. If this happens, you can follow the steps below to bring up the new primary server without data loss.

Step 1 Ensure that the secondary server is currently in "Secondary Active" state. If you have set the Failover Type on the primary server to "manual", you will need to trigger the failover to the secondary manually.

Step 2 Ensure that the old primary server you are replacing has been disconnected from the network.

Step 3 Ensure that the new primary server is ready for use. This will include connecting it to the network and assigning it the same server IP, subnet mask, gateway as the old primary server. You will also need to enter the same authentication key that you entered when installing the secondary server.

Step 4 Ensure that both the primary and secondary servers are at the same patch level and if you want to replace the primary server, then you must :

a) Ensure the primary and secondary server are in TOFU Mode.
b) Login to Secondary server admin CLI.
c) Execute the following command in the secondary server CLI:
d) PIServer/admin# ncs certvalidation tofu-certs deletecert host <primaryserver's-hostname>

This is required to re-establish the communication between the Primary and Secondary servers.

Step 5 Trigger a failback from the secondary to the newly installed primary. During failback to the new primary HA server, a full database copy will be performed, so this operation will take time to complete depending on the available bandwidth and network latency (see "Network Throughput Restrictions on HA" in Related Topics). You will see the two servers transition through the following series of HA states:

Primary HA State Transitions...	Secondary HA State Transitions...
From: HA not configured	From: Secondary Active
To: Primary Failback	To: Secondary Failback
To: Primary Failback	To: Secondary Post Failback
To: Primary Active	To: Secondary Syncing

Related Topics

How to Trigger Failover, on page 31
How to Trigger Failback, on page 31
Respond to Other HA Events, on page 32
Network Throughput Restrictions on HA, on page 9

How to Recover From Split-Brain Scenario

As explained in "Automatic Versus Manual Failover" (see Related Topics), the possibility of data loss always exists on the rare occasions when a "split-brain scenario" occurs. In this case, you can choose to save the newly added data on the secondary and forget the data that was added on the primary, as explained in the following steps.

Step 1 Once the network is up, and the secondary server is up, the primary will restart itself automatically, using its standby database. The HA status of the primary server will be, first, "Primary Failover" transitioning to "Primary Synching". You can verify this by logging on to the primary server's Health Monitor web page.

Step 2 Once the primary server's status is "Primary Syncing, confirm that a user can log into the secondary server's Prime Infrastructure page using the web browser (for example, https://x.x.x.x:443). Do not proceed until you have verified this.

Step 3 Once access to the secondary is verified, initiate a failback from the secondary server's Health Monitor web page (see How to Trigger Failback, on page 31). You can continue to perform monitoring activities on the secondary server until the switchover to the primary is completed.

For more information, see Restart Prime Infrastructure Using CLI.

Related Topics

Automatic Versus Manual Failover, on page 12
Remove HA Via the CLI, on page 47
How to Register HA on the Primary Server, on page 17

How to Resolve Database Synchronization Issues

To resolve the database synchronization issue, when the primary server is in "Primary Active" state and the secondary server is in "Secondary Syncing" state, do the following:

Step 1 Remove HA, see Remove HA Via the CLI, on page 47 and Remove HA Via the GUI, on page 46.

Step 2 After both the primary and secondary servers reaches "HA not configured" state, perform the HA registration. See Set Up High Availability, on page 15

High Availability Reference Information

The following sections supply reference information on HA.

Related Topics

HA Configuration Mode Reference, on page 42
HA State Reference, on page 43
HA State Transition Reference, on page 44
High Availability CLI Command Reference, on page 46
Reset the HA Authentication Key, on page 46
Remove HA Via the GUI, on page 46
Remove HA Via the CLI, on page 47
Remove HA During Restore, on page 47
Remove HA During Upgrade, on page 48
Using HA Error Logging, on page 48
Reset the HA Server IP Address or Host Name, on page 49

HA Configuration Mode Reference

The following table lists all possible HA configuration modes.

Table 2: High Availability Modes

Mode	Description
HA not configured	HA is not configured on this Prime Infrastructure server
HA initializing	The HA registration process between the primary and secondary server has started.

Mode	Description
HA enabled	HA is enabled between the primary and secondary server.
HA alone	Primary server is now running alone. HA is enabled, but the primary server is out of sync with the secondary, or the secondary is down or otherwise unreachable.

Related Topics

High Availability Reference Information, on page 42

HA State Reference

The following table lists all possible HA states, including those that require no response from you.

Table 3: High Availability States

State	Server	Description
Stand Alone	Both	HA is not configured on this Prime Infrastructure server
Primary Alone	Primary	Primary restarted after it lost secondary. Only Health Monitor is running in this state.
HA Initializing	Both	HA Registration process between the primary and secondary server has started.
Primary Active	Primary	Primary server is now active and is synchronizing with secondary server.
Primary Database Copy Failed	Primary	Primary servers being restarted will always check to see if a data gap has occurred due to the primary being down for 24 hours or more. If it detects such a gap, it will automatically trigger a data copy from the active secondary server. In rare cases, this database copy can fail, in which case this transition state is set on the primary. All attempts to failback to the primary are blocked until the database copy completes successfully. As soon as it does, the primary state is set to “Primary Syncing”.
Primary Failover	Primary	Primary server detected a failure.
Primary Failback	Primary	Failback triggered by the User is currently in progress.
Primary Lost Secondary	Primary	Primary server is unable to communicate with the secondary server.
Primary Preparing for Failback	Primary	This state will be set on primary server startup after a failover to the secondary. This state signifies that the primary server has started up in standby mode (because the secondary server is still active) and is ready for failback. Once the primary server is ready for failback, its state will be set to “Primary Syncing”.
Primary Syncing	Primary	Primary server is synchronizing the database and configuration files from the active secondary. Primary gets into this state when primary processes are brought up after failover to secondary and secondary is playing the active role.
Primary Uncertain	Primary	Primary server's application processes are not able to connect to its database.
Secondary Alone	Secondary	Primary server is not reachable from secondary after primary server restart.
Secondary Syncing	Secondary	Secondary server is synchronizing the database and configuration files from the primary.

State	Server	Description
Secondary Active	Secondary	Failover from the primary server to the secondary server has completed successfully.
Secondary Lost Primary	Secondary	Secondary server is not able to connect to the primary server (occurs when the primary fails or network connectivity is lost). In case of automatic failover from this state, the secondary will automatically move to Active state. In case of a manual failover, the user can trigger a failover to make the secondary active.
Secondary Failover	Secondary	Failover triggered and in progress.
Secondary Failback	Secondary	Failback triggered and in progress (database and file replication is in progress).
Secondary Post Failback	Secondary	This state occurs after failback is triggered, replication of database and configuration files from the secondary to the primary is complete, and Health Monitor has initiated changes of the secondary server's status to Secondary Syncing and the primary server's status to Primary Active. These status changes and associated process starts and stops are in progress.
Secondary Uncertain	Secondary	Secondary server's application processes are not able to connect to secondary server's database.

Related Topics

High Availability Reference Information, on page 42

HA State Transition Reference

The following figure details all possible state transitions for the primary server.

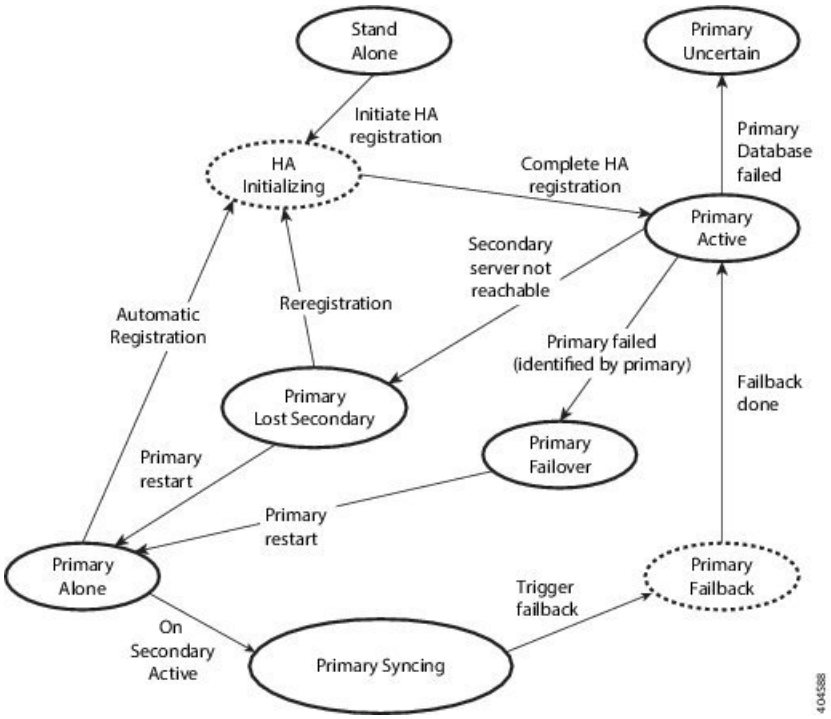

The following figure details all possible state transitions for the secondary server.

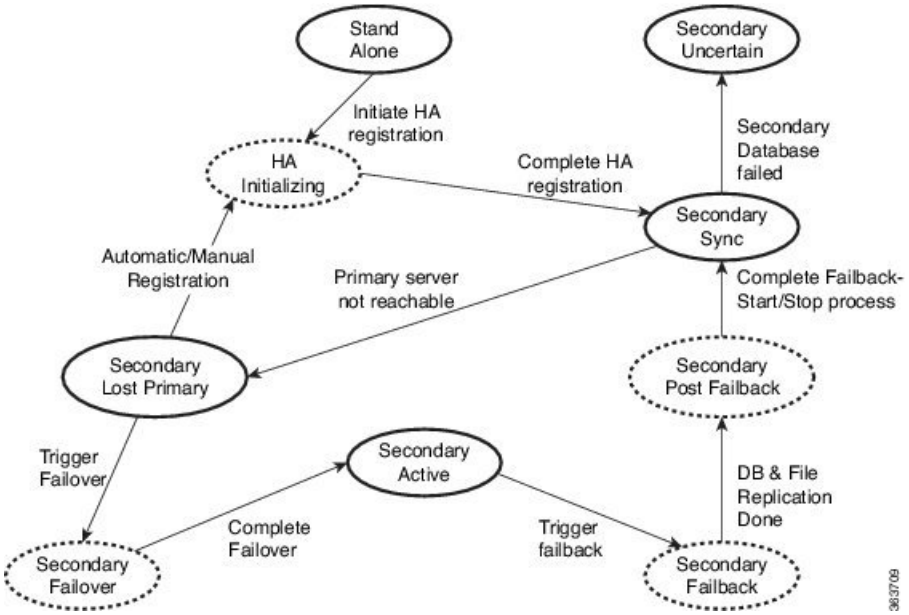

Related Topics

High Availability Reference Information, on page 42

High Availability CLI Command Reference

The following table lists the CLI commands available for HA management. Log in as admin to run these commands on the primary server (see How to Connect Via CLI):

Table 4: High Availability Commands

Command	Description
ncs ha ?	Get help with high availability CLI commands
ncs ha authkey authkey	Update the authentication key for high availability
ncs ha remove	Remove the High Availability configuration
ncs ha status	Get the current status for High Availability

Related Topics

High Availability Reference Information, on page 42

Reset the HA Authentication Key

Prime Infrastructure administrators can change the HA authentication key using the **ncs ha authkey** command. You will need to ensure that the new authorization key meets the password standards.

Step 1 Connect to the primary server via CLI. Do not enter “configure terminal” mode.

Step 2 Enter the following at the command line:

admin# **ncs ha authkey** *MyNewAuthKey*

Where *MyNewAuthKey* is the new authorization key. For more information, see How to Connect Via CLI.

Related Topics

Before You Begin Setting Up High Availability, on page 16
High Availability Reference Information, on page 42

Remove HA Via the GUI

The simplest method for removing an existing HA implementation is via the GUI, as shown in the following steps. You can also remove the HA setup via the command line.

Note that, to use this method, you must ensure that the primary Prime Infrastructure server is currently in the “Primary Active” state. If for any reason the secondary server is currently active, perform a failback and then try to remove the HA configuration after the failback is complete and the secondary’s automatic restart has finished.

Step 1 Log in to the primary Prime Infrastructure server with a user ID that has administrator privileges.

Step 2 Select **Administration > Settings > High Availability > HA Configuration**.

Step 3 Select **Remove**. Removing the HA configuration takes from 3 to 4 minutes.

Once the removal is complete, ensure that the HA configuration mode displayed on the page now reads “HA Not Configured”.

Related Topics

Remove HA Via the CLI, on page 47
How to Trigger Failback, on page 31
High Availability Reference Information, on page 42

Remove HA Via the CLI

If for any reason you cannot access the Prime Infrastructure GUI on the primary server, administrators can remove the HA setup via the command line, using the steps below.

Note that, to use this method, you must ensure that the primary Prime Infrastructure server is currently in the “Primary Active” state. If for any reason the secondary server is currently active, perform a failback and then try to remove the HA configuration after the failback is complete and the secondary’s automatic restart has finished.

Step 1 Connect to the primary server via CLI. Do not enter “configure terminal” mode.

Step 2 Enter the following at the command line:

admin# **ncs ha remove**. For more information, see How to Connect Via CLI.

Related Topics

Remove HA Via the GUI, on page 46
How to Trigger Failback, on page 31
High Availability Reference Information, on page 42

Remove HA During Restore

Prime Infrastructure does not back up configuration settings related to High Availability.

In order to restore a Prime Infrastructure implementation that is using HA, be sure to restore the backed up data to the primary server only. The restored primary will automatically replicate its data to the secondary server. Running a restore on the secondary server is not needed and will generate an error message if you attempt it.

To restore a Prime Infrastructure implementation that uses HA, follow the steps below.

Step 1 Use the GUI to remove the HA settings from the primary server (see “Remove HA Via the GUI” in Related Topics).

Step 2 Restore the primary server as needed.

Step 3 Once the restore is complete, perform the HA registration process again.

For more information, see Restore Prime Infrastructure Data and How to Connect Via CLI.

Related Topics

Remove HA Via the GUI, on page 46
How to Register HA on the Primary Server, on page 17
High Availability Reference Information, on page 42

Remove HA During Upgrade

To upgrade a Prime Infrastructure implementation that uses HA, follow the steps below.

Step 1 Use the GUI to remove the HA settings from the primary server (see "Remove HA Via the GUI" in Related Topics, below).

Step 2 Upgrade the primary server as needed.

Step 3 Re-install the secondary server using the current image.

Note that upgrading the secondary server from the previous version or a beta version is not supported. The secondary server must always be a fresh installation.

Step 4 Once the upgrade is complete, perform the HA registration process again.

Note After upgrade, health montior page will display the below health monitor event message:

Primary Authentication Key was changed by Admin

For more information, see How to Connect Via CLI.

Related Topics

Remove HA Via the GUI, on page 46
How to Register HA on the Primary Server, on page 17
High Availability Reference Information, on page 42

Using HA Error Logging

Error logging for the High Availability feature is disabled by default, to save disk space and maximize performance. If you are having trouble with HA, the best place to begin is by enabling error logging and to examine the log files.

Step 1 View the Health Monitor page for the server having trouble.

Step 2 In the **Logging** area, in the **Message Level** dropdown, select the error-logging level you want.

Step 3 Click **Save**.

Step 4 When you want to download the log files: In the **Logs** area, click **Download**. You can open the downloaded log files using any ASCII text editor.

Related Topics

Access the Health Monitor Web Page, on page 30
High Availability Reference Information, on page 42

Reset the HA Server IP Address or Host Name

Avoid changing the IP address or hostname of the primary or secondary HA server, if possible. If you must change the IP address or hostname, remove the HA configuration from the primary server before making the change. When finished, re-register HA.

Related Topics

Remove HA Via the GUI, on page 46
How to Register HA on the Primary Server, on page 17
High Availability Reference Information, on page 42

Resolve TOFU Failure at Any State

When the primary and secondary servers communicate, there is a possibility of the below TOFU error occurrence.

You must correct the following error(s) before proceeding. 'A Trust-on-first-use (TOFU) based Certificate is configured for this connection. The current certificate on the remote host is different than what was used earlier.'

To rectify this issue, you must perform the following.

- Clear the existing certificate using the NCS CLI command on both the primary and secondary servers.

```
ncs certvalidation tofu-certs deletecert host <server-hostname>
```

Configure MSE High Availability

The Cisco Mobility Services Engine (MSE) is a platform for hosting multiple mobility applications. Under an MSE high availability (HA) configuration, an active MSE is backed up by another inactive instance of MSE. The active MSE is called the primary MSE and the inactive MSE is called the secondary MSE.

Related Topics

Overview of the MSE High Availability Architecture, on page 49
Set Up MSE High Availability: Workflow, on page 52

Overview of the MSE High Availability Architecture

The main component of MSE high availability is the health monitor. The health monitor configures, manages, and monitors the HA setup on each MSE. Heartbeat is maintained between the primary and secondary MSE. Health monitor is responsible for setting up the database, file replication, and monitoring the application. When the primary MSE fails and the secondary MSE takes over, the virtual address of the primary MSE is switched transparently to the secondary MSE. Note that:

- Every active primary MSE is backed up by another inactive instance. The purpose of the secondary MSE is to monitor the availability and state of the primary MSE. The secondary MSE becomes active only after the failover procedure is initiated.
- One secondary MSE can support one primary MSE.

The MSEs, Synchronize Services, Synchronization History, High Availability, Context Aware Notifications, and Mobile Concierge pages on the Services tab are available only in the virtual domain in Release 7.3.

The following related topics provide additional details on the MSE high availability architecture.

Related Topics

MSE High Availability Pairing Matrix, on page 50
Guidelines and Limitations for MSE High Availability, on page 50
Failover Scenario for MSE High Availability, on page 51
Failback Scenario for MSE High Availability, on page 51
Licensing Requirements for MSE High Availability, on page 52
Configure MSE High Availability , on page 49

MSE High Availability Pairing Matrix

The following table lists the types of MSE servers that can be paired in a high-availability configuration.

Table 5: MSE High Availability Server Pairing Matrix

Primary Server Type	Secondary Server Type				
3355	VA-2	VA-3	VA-4	VA-5	
3355	Y	N	N	N	N
VA-2	N	Y	Y	Y	Y
VA-3	N	N	Y	Y	Y
VA-4	N	N	N	Y	Y
VA-5	N	N	N	N	Y

Related Topics

Using the Remote Model, on page 11
Guidelines and Limitations for MSE High Availability, on page 50

Guidelines and Limitations for MSE High Availability

Administrators implementing MSE High Availability and planning to manage it via Prime Infrastructure should observe the following guidelines and limitations:

- Both the health monitor IP and virtual IP should be accessible from Prime Infrastructure.
- The health monitor IP and virtual IP should always be different. The health monitor and virtual interface can be on the same network interface or different interfaces.
- You can use either manual or automatic failover. Failover should be considered temporary. The failed MSE should be restored to normal as soon as possible, and failback should be re-initiated. The longer it

takes to restore the failed MSE, the longer you are running with a single MSE without high availability support.
- You can use either manual or automatic failback.
- Both the primary and secondary MSE should be running the same software version.
- High Availability over WAN is not supported.
- High Availability over LAN is supported only when both the primary and secondary MSEs are in the same subnet.
- The ports over which the primary and secondary MSEs communicate must be open (not blocked with network firewalls, application fireways, gateways, and so on). The following input/output ports should be opened: 80, 443, 8080, 8081, 22, 8001, 1521, 1411, 1522, 1523, 1524, 1525, 9006, 15080, 61617, 59000, 12091, 1621, 1622, 1623, 1624, 1625, 8083, 8084, and 8402.

Related Topics

Overview of the MSE High Availability Architecture, on page 49
MSE High Availability Pairing Matrix, on page 50
Failover Scenario for MSE High Availability, on page 51

Failover Scenario for MSE High Availability

When a primary MSE failure is detected, the following events occur:

- The primary MSE is confirmed as non-functioning (hardware fail, network fail, and so on) by the health monitor on the secondary MSE.
- If automatic failover isn enabled, the secondary MSE starts immediately.
- If manual failover is enabled, an e-mail is sent to the administrator asking if they want to manually start failover. This e-mail is sent only if the e-mail is configured for MSE alarms.
- The result of the failover operation is indicated as an event in the Health Monitor UI, and a critical alarm is sent to Prime Infrastructure.

Related Topics

Overview of the MSE High Availability Architecture, on page 49
Guidelines and Limitations for MSE High Availability, on page 50
Failback Scenario for MSE High Availability, on page 51

Failback Scenario for MSE High Availability

When the primary MSE is restored to its normal state, if the secondary MSE is already in failover state for the primary, then failback can be invoked.

Failback can occur only if the secondary MSE is in one of the following states for the primary instance:

- The secondary MSE is actually failing over for the primary MSE.
- Manual failover is configured but the administrator did not invoke it.
- The primary MSE failed but the secondary MSE cannot take over because it has encountered errors.
- Failback can occur only if the administrator starts up the failed primary MSE.

Related Topics

Overview of the MSE High Availability Architecture, on page 49
Failover Scenario for MSE High Availability, on page 51
Licensing Requirements for MSE High Availability, on page 52

Licensing Requirements for MSE High Availability

For high availability, an activation license is required on the primary and secondary virtual appliances. No other service license is required on the secondary MSE. It is required only on the primary MSE.

Related Topics

Overview of the MSE High Availability Architecture, on page 49
Failback Scenario for MSE High Availability, on page 51

Set Up MSE High Availability: Workflow

During the installation of the MSE software (or using the MSE setup script), configure some critical elements. Pair up the primary and secondary MSE from the Prime Infrastructure UI.

By default, all MSEs are configured as primary. If you do not want high availability support and are upgrading from an earlier release, you can continue to use the IP address for the MSE. If you want to set up high availability, then you must configure the health monitor IP address. The health monitor then becomes a virtual IP address.

Configuring MSE high availability consists of the following steps:

1. Prepare the MSEs for High Availability
2. Configure the Primary MSE
3. Configure the Secondary MSE

You may also need to reconfigure MSE high availability if you must replace the primary MSE server.

For details, see the corresponding Related Topics, below.

Related Topics

Prepare the MSEs for High Availability, on page 52
Configure MSE High Availability on Primary MSEs, on page 53
Configure MSE High Availability on Secondary MSEs, on page 61
Replace Primary MSEs, on page 66
Configure MSE High Availability

Prepare the MSEs for High Availability

To prepare your primary and secondary MSEs for high availabililty, follow these steps:

Step 1 Ensure that the network connectivity between the primary and secondary MSEs is functioning and that all the necessary ports are open.

Step 2 Install the correct version of MSE on the primary MSE.

Step 3 Make sure that the same MSE version is installed on the secondary MSE.

Related Topics

Replace Primary MSEs, on page 66
Configure MSE High Availability , on page 49

Configure MSE High Availability on Primary MSEs

To configure a primary MSE for high availabililty, follow these steps:

Step 1 On the intended primary MSE, enter the following command:

/opt/mse/setup/setup.sh

The setup script displays the following prompts, which you can answer using the suggested responses given in bold (in this and later steps):

Welcome to the Cisco Mobility Services Engine Appliance Setup.

You may exit the setup at any time by typing <Ctrl+c>.

Would you like to configure MSE using:

1. Menu mode

2. Wizard mode

Choose 1 or 2: **1**

Mobility Services Engine Setup

Please select a configuration option below and enter the requested information. You may exit setup at any time by typing <Ctrl +C>.

You will be prompted to choose whether you wish to configure a parameter, skip it, or reset it to its initial default value. Skipping a parameter will leave it unchanged from its current value.

Please note that the following parameters are mandatory and must be configured at lease once.

-> Hostname

-> Network interface eth0

-> Timezone settings

-> Root password

-> NTP settings

-> Prime Infrastructure password

You must select option 24 to verify and apply any changes made during this session.

PRESS <ENTER> TO CONTINUE:

Configure MSE:

1) Hostname * 13) Remote syslog settings

2) Network interface eth0 settings* 14) Host access control settings

3) Timezone settings* 15) Audit Rules

4) Root password * 16) Login banner

5) NTP settings * 17) System console restrictions

6) Prime Infrastructure password * 18) SSH root access

7) Display current configuration 19) Single user password check

8) Domain 20) Login and password settings

9) High availability role 21) GRUB password

10) Network interface eth1 settings 22) Root access control

11) DNS settings 23) Auto start MSE on system boot up

12) Future restart time 24) ## Verify and apply changes ##

Please enter your choice [1 - 24]:

Step 2 Configure the primary MSE hostname:

Please enter your choice [1 - 24]: **1**

Current Hostname=[mse]

Configure Hostname? (Y)es/(S)kip/(U)se default [Skip]: **y**

The host name should be a unique name that can identify the device on the network. The hostname should start with a letter, end with a letter or number, and contain only letters, numbers, and dashes.

Enter a Host name [mse]:**mse1**

Step 3 Configure the primary MSE domain:

Please enter your choice [1-24]: **8**

Current domain=[ ]

Configure domain name? (Y)es/(S)kip/(U)se default [Skip]: **S**

Step 4 Configure the primary MSE network interface eth0 settings.

Please enter your choice [1 - 24]: **2**

Current eth0 interface IP address=[10.0.0.1]

Current eth0 interface netmask=[255.0.0.0]

Current IPv4 gateway address=[172.20.104.123]

Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]: **y**

Enter an IP address for first Ethernet interface of this machine.

Enter eth0 IP address [10.0.0.2]:

Enter the network mask for IP address 172.21.105.126

Enter network mask [255.255.255.224]:

Enter the default gateway address for this machine.

Note that the default gateway must be reachable from the first Ethernet interface.

Enter default gateway address [172.20.104.123]:

Step 5 Configure the primary MSE root password:

Please enter your choice [1 - 24]: **4**

Root password has not been configured

Configure root password? (Y)es/(S)kip/(U)se default [Skip]: **Y**

Changing password for user root.

You can now choose the new password.

A valid password should be a mix of upper and lower case letters, digits, and other characters. You can use an 8 character long password with characters from all of these classes. An upper case letter that begins the password and a digit that ends it do not count towards the number of character classes used.

Enter new password: **password**

Step 6 Configure the primary MSE's high availability role:

Current role=[Primary]

Configure High Availability? (Y)es/(S)kip/(U)se default [Skip]: **y**

High availability role for this MSE (Primary/Secondary)

Select role [1 for Primary, 2 for Secondary] [1]: 1

Health monitor interface holds physical IP address of this MSE server.

This IP address is used by Secondary, Primary MSE servers and Prime Infrastructure to communicate among themselves

Select Health Monitor Interface [eth0/eth1] [eth0]: eth0

Direct connect configuration facilitates use of a direct cable connection between the primary and secondary MSE servers.

This can help reduce latencies in heartbeat response times, data replication and failure detection times.

Please choose a network interface that you wish to use for direct connect. You should appropriately configure the respective interfaces.

"none" implies you do not wish to use direct connect configuration.

Select direct connect interface [eth0/eth1/none] [none]:

Enter a Virtual IP address for the Primary MSE server

Enter Virtual IP address [1.1.1.1]: **10.10.10.11**

Enter network mask for IP address 10.10.10.1

Enter network mask [1.1.1.1]: **255.255.255.0**

Select to start the server in recovery mode.

You should choose yes only if this primary MSE was paired earlier and you have now lost the configuration from this box.

And, now you want to restore the configuration from Secondary via Cisco Prime Infrastructure

Do you wish to start this MSE in HA receovery mode?: (yes/no} [no]:no

Current IP address = [1.1.1.10]

Current eth0 netmask=[255.255.255.0]

Current gateway address=[1.1.1.1]

Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Skip]:

Enter an IP address for first Ethernet interface of this machine.

Enter eth0 IP address [1.1.1.10]: **10.10.10.12**

Enter the network mask for IP address 10.10.10.12

Enter network mask [255.255.255.0]: **255.255.255.0**

Enter an default gateway address for this machine.

Note that the default gateway must be reachable from the first Ethernet interface. Enter default gateway address [1.1.1.1]:**10.10.10.1**

The second Ethernet interface is currently disabled for this machine.

Configure eth1 interface parameters? (Y)es/(S)kip/(U)se default [Yes]: S

Step 7 Configure the primary MSE timezone settings:

Please enter your choice [1 - 24]: 3

Current Timezone=[America/New_York]

Configure Timezone? (Y)es/(S)kip/(U)se default [Skip]: y

Enter the current date and time.

Please identify a location so that time zone rules can be set correctly. Please select a continent or ocean.

1) Africa

2) Americas

3) Antarctica

4) Arctic Ocean

5) Asia

6) Atlantic Ocean

7) Australia

8) Europe

9) Indian Ocean

10) Pacific Ocean

11) UTC - I want to use Coordinated Universal Time.

#? 2

Please select a country.

1) Anguilla 27) Honduras

2) Antigua & Barbuda

5) Bahamas 31) Montserrat

6) Barbados 32) Netherlands Antilles

7) Belize 33) Nicaragua

8) Bolivia 34) Panama

9) Brazil 35) Paraguay

10) Canada 36) Peru

11) Cayman Islands 37) Puerto Rico

12) Chile 38) St Barthelemy

13) Colombia 39) St Kitts & Nevis

14) Costa Rica 40) St Lucia

41) St Martin (French part)

16) Dominica 42) St Pierre & Miquelon

17) Dominican Republic 43) St Vincent

18) Ecuador 44) Suriname

19) El Salvador 45) Trinidad & Tobago

20) French Guiana 46) Turks & Caicos Is

21) Greenland 47) United States

22) Grenada 48) Uruguay

23) Guadeloupe 49) Venezuela

24) Guatemala 50) Virgin Islands (UK)

25) Guyana 51) Virgin Islands (US)

26) Haiti

#? 47

Please select one of the following time zone regions.

1) Eastern Time

2) Eastern Time - Michigan - most locations

3) Eastern Time - Kentucky - Louisville area

4) Eastern Time - Kentucky - Wayne County

5) Eastern Time - Indiana - most locations

6) Eastern Time - Indiana - Daviess, Dubois, Knox & Martin Counties

7) Eastern Time - Indiana - Pulaski County

8) Eastern Time - Indiana - Crawford County

9) Eastern Time - Indiana - Pike County

10) Eastern Time - Indiana - Switzerland County

12) Central Time - Indiana - Perry County

13) Central Time - Indiana - Starke County

14) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties

15) Central Time - North Dakota - Oliver County

16) Central Time - North Dakota - Morton County (except Mandan area)

18) Mountain Time - south Idaho & east Oregon

20) Mountain Standard Time - Arizona

21) Pacific Time

22) Alaska Time

23) Alaska Time - Alaska panhandle

24) Alaska Time - Alaska panhandle neck

25) Alaska Time - west Alaska

26) Aleutian Islands

27) Hawaii

#? 21

The following information has been given:

United States

Pacific Time

Therefore TZ='America/Los_Angeles' will be used.

Local time is now: Sun Apr 6 18:45:27 PDT 2020. Universal Time is now: Mon Apr 7 01:45:27 UTC 2020. Is the above information OK?

1) Yes

2) No

#? 1

Step 8 Configure the primary MSE DNS settings:

Please enter your choice [1 - 24]: 11

Domain Name Service (DNS) Setup

Enable DNS (yes/no) [no]: y

Default DNS server 1=[8.8.8.8]

Enter primary DNS server IP address:

DNS server address must be in the form #.#.#.#, where # is 0 to 255 or hexadecimal :

separated v6 address

Enter primary DNS server IP address [8.8.8.8]:

Enter backup DNS server IP address (or none) [none]:

Step 9 Configure the primary MSE NTP settings:

Please enter your choice [1 - 24]: 5

Network Time Protocol (NTP) Setup.

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

NTP is currently disabled.

Configure NTP related parameters? (Y)es/(S)kip/(U)se default [Skip]: y

Enter whether or not you would like to set up the

Network Time Protocol (NTP) for this machine.

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

Enable NTP (yes/no) [no]: y

Default NTP server 1=[time.nist.gov] Enter NTP server name or address:

NTP server address must be in the form #.#.#.3, where # is 0 to 255 hexadecimal :

separated v6 address.

Enter NTP server name or [time.nist.gov]:

Enter another NTP server IP address (or none) [none]:

Configure NTP Authentication? (Y)es/(S)kip/(U)se default [Skip]: y

Enter NTP Auth key Number [1]:

Enter NTP Auth key Value (String) [Secret]: Do you want to continue (yes/no) [no]: y

Step 10 Configure the Prime Infrastructure password:

Please enter your choice [1 - 24]: 6

Cisco Prime Infrastructure communication password has not been configured. Configure Prime Infrastructure password? (Y)es/(S)kip/(U)se default [Yes]:

Enter a password for the admin user.

The admin user is used by the Prime Infrastructure and other northbound systems to authenticate their SOAP/XML session with the server. Once this password is updated, it must correspondingly be updated on the NCS page for MSE General Parameters so that the Prime Infrastructure can communicate with the MSE.

Step 11 Verify and apply your changes:

Please enter your choice: 24

Please verify the following setup information.

----------BEGIN---------- Hostname=mse1

Role= 1, Health Monitor Intercace=eth0, Direct connect interface=none

Virtual IP Address=10.10.10.11, Virtual IP Netmask=255.255.255.0

Eth0 IP address=10.10.10.12, Eth0 network mask=255.0.0.0

Default Gateway=10.10.10.1

Time zone=America/Los_Angeles

Enable DNS=yes, DNS servers=8.8.8.8

Enable NTP=yes, NTP servers=time.nist.gov

Time zone=America/Los_Angeles

Root password is changed.

Cisco Prime Infrastructure password is changed.

----------END----------

You may enter "yes" to proceed with configuration, "no" to make

more changes.

Configuration Changed

Is the above information correct (yes or no): yes

Checking mandatory configuration information...

Root password: Not configured

WARNING

The above parameters are mandatory and need to be configured.

Ignore and proceed (yes/no): yes

Setup will now attempt to apply the configuration. Restarting network services with new settings. Shutting down interface eth0:

The system is minimally configured right now. It is strongly recommended that you run the setup script under /opt/mse/setup/setup.sh command to configure all appliance related parameters immediately after installation is complete.

PRESS <ENTER> TO EXIT THE INSTALLER:

Step 12 Reboot the system:

[root@mse1]# reboot Stopping MSE Platform

Flushing firewall rules: [OK]

Setting chains to policy ACCEPT: nat filter [OK] Unloading iptables modules: [ok]

Broadcast message from root (pts/0) (Tue Apr29 14:15:27:2014):

The system is going down for reboot NOW:

Step 13 Start the MSE services:

[root@mse1]# /etc/init.d/msed start

Starting MSE Platform.

Starting Health Monitor, Waiting to check the status. Starting Health Monitor, Waiting to check the status. Health Monitor successfully started

Starting Admin process... Started Admin process. Starting database .....

Database started successfully. STarting framework and services..... Framework and services successfully started

Step 14 After all services have started, confirm MSE services are working properly by entering the following command:

[root@mse1]# getserverinfo

Related Topics

Prepare the MSEs for High Availability, on page 52
Configure MSE High Availability on Secondary MSEs, on page 61
Configure MSE High Availability , on page 49

Configure MSE High Availability on Secondary MSEs

To prepare your secondary MSE for high availabililty, follow these steps:

Step 1 On the intended secondary MSE, enter the following command:

/opt/mse/setup/setup.sh

The setup script displays the same prompts as for the primary MSE:

Step 2 Configure the secondary MSE hostname:

Please enter your choice [1 - 24]: **1**

Current hostanme=[mse1]

Configure hostname? (Y)es/(S)kip/(U)se default [Yes]: yes

The host name should be a unique name that can identify the device on the network. The hostname should start with a letter, end with a letter or number, and contain only letters, numbers, and dashes.

Enter a hostname [mse]: **mse2**

Step 3 Configure the secondary MSE domain:

Please enter your choice [1-24]: 8

Current domain=[ ]

Configure domain name? (Y)es/(S)kip/(U)se default [Skip]: S

Step 4 Configure the secondary MSE high availability role:

Current role=[Primary]

Configure High Availability? (Y)es/(S)kip/(U)se default [Skip]: High availability role for this MSE (Primary/Secondary)

Select role [1 for Primary, 2 for Secondary] [1]: 2

Health monitor interface holds physical IP address of this MSE server.

This IP address is used by Secondary, Primary MSE servers and Prime Infrastructure to communicate among themselves

Select Health Monitor Interface [eth0/eth1] [eth0]: eth0

Direct connect configuration facilitates use of a direct cable connection between the primary and secondary MSE servers.This can help reduce latencies in heartbeat response times, data replication and failure detection times.Please choose a network interface that you wish to use for direct connect. You should appropriately configure the respective interfaces.

"none" implies you do not wish to use direct connect configuration.

Select direct connect interface [eth0/eth1/none] [none]:

Current IP address=[1.1.1.10]

Current eth0 netmask=[255.255.255.0] Current gateway address=[1.1.1.1]

Configure eth0 interface parameters? (Y)es/(S)kip/(U)se default [Yes]:

Enter an IP address for first Ethernet interface of this machine. Enter eth0 IP address [1.1.1.10]: 10.10.10.13

Enter the network mask for IP address 10.10.10.13

Enter network mask [255.255.255.0]:

Enter an default gateway address for this machine.

Note that the default gateway must be reachable from the first Ethernet interface. Enter default gateway address [1.1.1.1]:10.10.10.1

The second Ethernet interface is currently disabled for this machine. Configure eth1 interface parameters? (Y)es/(S)kip/(U)se default [Yes]: S

Step 5 Configure the secondary MSE timezone settings:

Please enter your choice [1 - 24]: 3

Current Timezone=[America/New_York]

Configure Timezone? (Y)es/(S)kip/(U)se default [Skip]: y

Enter the current date and time.

Please identify a location so that time zone rules can be set correctly. Please select a continent or ocean.

1) Africa

2) Americas

3) Antarctica

4) Arctic Ocean

5) Asia

6) Atlantic Ocean

7) Australia

8) Europe

9) Indian Ocean

10) Pacific Ocean

11) UTC - I want to use Coordinated Universal Time.

#? 2

Please select a country.

1) Anguilla 27) Honduras

2) Antigua & Barbuda

5) Bahamas 31) Montserrat

6) Barbados 32) Netherlands Antilles

7) Belize 33) Nicaragua

8) Bolivia 34) Panama

9) Brazil 35) Paraguay

10) Canada 36) Peru

11) Cayman Islands 37) Puerto Rico

12) Chile 38) St Barthelemy

13) Colombia 39) St Kitts & Nevis

14) Costa Rica 40) St Lucia

41) St Martin (French part)

16) Dominica 42) St Pierre & Miquelon

17) Dominican Republic 43) St Vincent

18) Ecuador 44) Suriname

19) El Salvador 45) Trinidad & Tobago

20) French Guiana 46) Turks & Caicos Is

21) Greenland 47) United States

22) Grenada 48) Uruguay

23) Guadeloupe 49) Venezuela

24) Guatemala 50) Virgin Islands (UK)

25) Guyana 51) Virgin Islands (US)

26) Haiti

#? 47

Please select one of the following time zone regions.

1) Eastern Time

2) Eastern Time - Michigan - most locations

3) Eastern Time - Kentucky - Louisville area

4) Eastern Time - Kentucky - Wayne County

5) Eastern Time - Indiana - most locations

6) Eastern Time - Indiana - Daviess, Dubois, Knox & Martin Counties

7) Eastern Time - Indiana - Pulaski County

8) Eastern Time - Indiana - Crawford County

9) Eastern Time - Indiana - Pike County

10) Eastern Time - Indiana - Switzerland County

11) Central Time

12) Central Time - Indiana - Perry County

13) Central Time - Indiana - Starke County

14) Central Time - Michigan - Dickinson, Gogebic, Iron & Menominee Counties

15) Central Time - North Dakota - Oliver County

16) Central Time - North Dakota - Morton County (except Mandan area)

17) Mountain Time

18) Mountain Time - south Idaho & east Oregon

19) Mountain Time - Navajo

20) Mountain Standard Time - Arizona

21) Pacific Time

22) Alaska Time

23) Alaska Time - Alaska panhandle

24) Alaska Time - Alaska panhandle neck

25) Alaska Time - west Alaska

26) Aleutian Islands

27) Hawaii

#? 21

The following information has been given: United States

Pacific Time

Therefore TZ='America/Los_Angeles' will be used.

Local time is now: Sun Apr 6 18:45:27 PDT 2014. Universal Time is now: Mon Apr 7 01:45:27 UTC 2014. Is the above information OK?

1) Yes

2) No

#? 1

Step 6 Configure the secondary MSE NTP settings:

Please enter your choice [1 - 24]: 5

Network Time Protocol (NTP) Setup.

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

NTP is currently disabled.

Configure NTP related parameters? (Y)es/(S)kip/(U)se default [Skip]: y

Enter whether or not you would like to set up the Network Time Protocol (NTP) for this machine.

If you choose to enable NTP, the system time will be configured from NTP servers that you select. Otherwise, you will be prompted to enter the current date and time.

Enable NTP (yes/no) [no]: y

Default NTP server 1=[time.nist.gov] Enter NTP server name or address:

NTP server address must be in the form #.#.#.3, where # is 0 to 255 hexadecimal :

separated v6 address.

Enter NTP server name or [time.nist.gov]:

Enter another NTP server IP address (or none) [none]:

Configure NTP Authentication? (Y)es/(S)kip/(U)se default [Skip]: y

Enter NTP Auth key Number [1]:

Enter NTP Auth key Value (String) [Secret]: Do you want to continue (yes/no) [no]: y

Step 7 Verify and apply your changes:

Please enter your choice: 24

Please verify the following setup information.

----------BEGIN---------- Hostname=mse2

Role= 2, Health Monitor Intercace=eth0, Direct connect interface=none

Eth0 IP address=10.10.10.13, Eth0 network mask=255.255.255.0

Default Gateway=10.10.10.1

Time zone=America/Los_Angeles

Enable NTP=yes, NTP servers=time.nist.gov

Time zone=America/Los_Angeles

----------END----------

You may enter "yes" to proceed with configuration, "no" to make more changes.

Configuration Changed

Is the above information correct (yes or no): yes

Checking mandatory configuration information...

Root password: Not configured

WARNING

The above parameters are mandatory and need to be configured.

Ignore and proceed (yes/no): yes

Setup will now attempt to apply the configuration.

Restarting network services with new settings. Shutting down interface eth0:

The system is minimally configured right now. It is strongly recommended that you run the setup script under /opt/mse/setup/setup.sh command to configure all appliance related parameters immediately after installation is complete.

PRESS <ENTER> TO EXIT THE INSTALLER:

Step 8 Reboot the system:

[root@mse2 installers]# reboot

Stopping MSE Platform

Flushing firewall rules: [OK]

Setting chains to policy ACCEPT: nat filter [OK] Unloading iptables modules: [ok]

Broadcast message from root (pts/0) (Tue Apr29 14:15:27:2014):

The system is going down for reboot NOW:

Step 9 Start the MSE services:

[root@mse2]# /etc/init.d/msed start

Starting MSE Platform.

Starting Health Monitor, Waiting to check the status. Starting Health Monitor, Waiting to check the status. Health Monitor successfully started

Starting Admin process... Started Admin process. Starting database .....

Database started successfully. STarting framework and services..... Framework and services successfully started

Related Topics

Prepare the MSEs for High Availability, on page 52
Configure MSE High Availability on Primary MSEs, on page 53
Configure MSE High Availability , on page 49

Replace Primary MSEs

If for any reason you need to replace a primary MSE, you will want to recover the current pairing information to a newly configured primary MSE, as explained in the followng steps.

Step 1 Configure the MSE as a primary using the setup script.

Step 2 Set up a pairing between the primary and secondary MSE using Prime Infrastructure.

Step 3 Initiate failover from the primary MSE to the secondary MSE.

Step 4 Configure the replacement MSE as a primary using the setup script. The new primary MSE must have the same version of the software as the secondary, and the same settings as the old primary MSE.

Step 5 Choose the recovery mode and follow the instructions.

Step 6 Initiate the failback to the new primary using Prime Infrastructure.

A new license is required on the this new primary MSE, as the original license will not match the UDI of the primary, and will not work.

Related Topics

Configure MSE High Availability on Primary MSEs, on page 53
Configure MSE High Availability , on page 49