



Configuring and Managing Cisco IPICS High Availability

Cisco IPICS provides the option of configuring a secondary hot standby server to provide high availability with no single point of failure. If a primary server fails, the secondary server automatically takes over service without communication interruption. This chapter describes the high availability function in detail It includes these topics:

- Overview, page 10-1
- Configuring Cisco IPICS Servers for HA, page 10-3
- Unconfiguring HA, page 10-6
- HA Affect on the IDC Connections, page 10-6
- HA Affect On Connected Devices, page 10-8
- Synchronizing the Server Time on HA Servers, page 10-8
- Performing a Manual Failover, page 10-9
- Resolving a Split Brain Scenario, page 10-10
- Reestablishing HA Configuration After Prolonged Server Downtime, page 10-14

Overview

The Cisco IPICS high availability (HA) feature allows two Cisco IPICS servers to be configured as a redundant pair: one *active server* to control the system users and connected devices, and a second *standby server* ready to assume control if the active server encounters a problem or goes off-line.

Defining the Active and Standby Servers

During initial configuration, the primary and standby server role is determined by the type of license file installed on the server.

- Install the *IPICS Base server* and *Policy Engine Base* license on the server designated as the primary server. This server will assume the active role.
- Install only the *High Availability* license on the server designated as the secondary server. This server will assume the standby role.

Once the server pair is operating in HA mode, the *active* and *standby* roles can switch between the servers. When a user logs into the Cisco IPICS Administration Console, they are automatically logged in to whichever server is the current active server.

<u>P</u> Tip

See the "Managing Your Licenses" section in *Cisco IP Interoperability and Collaboration System Server Installation Guide* for more information.

Remote Server Locations Using Secure Communication

The redundant Cisco IPICS servers can be located at remote locations, so that a catastrophic event in one location does not affect the other server. Secure communication is established during initial configuration, when the severs perform a one-time exchange of SSH/TLS certificates and public keys. See the "Configuring Cisco IPICS Servers for HA" section on page 10-3 for more information.

Server Failover Due to Local Critical Process Failure

If a local critical process (such as tomcat) fails, the active server instructs the standby server to assume the active role. This process can take several minutes and assumes that the standby server is ready and reachable.

Server Failover Due to Lost Heartbeat Message

The redundant servers also maintain communication by exchanging regular *heartbeat* messages. If the active server does not respond to the messages after a configured time, the standby server assumes the active role. This loss of heartbeat communication can be caused by a loss of network connectivity or a hardware failure.

By default, if the primary and standby servers exchange heartbeat messages every 15 seconds. If 5 heartbeat messages fail (for a total of 75 seconds), the standby server assumes the active role.

See the "Configuring Cisco IPICS Servers for HA" section on page 10-3 for information about modifying the default heartbeat time, if necessary.

Affect of Failover on IDCs and Connected Devices

If a failover occurs, IDC connections may be temporarily lost while the transfer is completed. The consoles continue to operate in off-line mode, however, and are automatically reconnected to the new server with no user interaction. The console remains connected even if the original server comes back online and assumes the standby role. See the "HA Affect on the IDC Connections" section on page 10-6 for more information.

Some connected devices such as iPhones also remain connected during a failover. Other devices, such as Cisco Unified IP Phones or standard telephones, must manually reconnect. See the "HA Affect On Connected Devices" section on page 10-8 for more information.

Manually Failover and Recovering from a Split Brain Scenario

You can manually force the active server to failover to the standby server. See the "Performing a Manual Failover" section on page 10-9 for more information.

If communication is lost between the primary and secondary servers, both servers may temporarily assume the active role. This situation is known as a *split brain scenario*. After communication is reestablished, you must manually transition the secondary server to the standby role. For more instructions, see the "Resolving a Split Brain Scenario" section on page 10-10.

Configuring Cisco IPICS Servers for HA

To configure Cisco IPICS HA, you must install and configure a primary server and a secondary server. These servers can be at different locations because they communicate using a secure, encrypted connection.

Before configuring HA, you must configure both servers with an IP address, an NTP server, and the correct license for each server, as described below.

Before You Begin

Before configuring redundant Cisco IPICS servers, do the following:

- Obtain the Network Time Protocol (NTP) server addresses that will be used to synchronize the Cisco IPICS servers. See the "Synchronizing the Server Time on HA Servers" section on page 10-8 for more information.
- Obtain the IP addresses for the primary and the secondary servers
- Obtain the ipicsadmin user passwords for the secondary server
- Obtain the *IPICS Base server* and *Policy Engine Base* licenses for the Primary server, and the *High Availability* license for the Secondary server. See the "Managing Your Licenses" section in *Cisco IP Interoperability and Collaboration System Server Installation Guide*.
- Verify that the installed HA server pair matches one of the supported configurations that is defined in *Cisco IPICS Compatibility Matrix*.
- Verify that any existing data on either HA server should be merged, or perform a clean install of the IPICS software
- Perform a clean install of the IPICS server software, if necessary. Any data on either server will be
 merged when the servers are configured for HA.
 - Data from a server that was previously configured as either primary or secondary in a different HA pair will be merged to the new HA configuration.
 - Existing data will be merged even if the server was not previously configured for HA.
 - If you do not want the existing data from a server to be merged when creating an HA pair, you
 must perform a clean install of the IPICS server software to remove that data. See the *Cisco IPICS Server Installation and Upgrade Guide* for more information.

To configure redundant Cisco IPICS servers for HA, perform the following procedure:

Procedure

Step 1 Physically install the primary and secondary servers.

See Cisco IP Interoperability and Collaboration System Server Installation Guide for instructions.

Step 2 (Optional) Perform a clean install of the IPICS server software, if necessary, to remove any existing data from the primary or secondary server.



Any data on either server will be merged when the servers are configured for HA. Perform a clean install of the server software to delete any existing data. See the *Cisco IPICS Server Installation and Upgrade Guide* for instructions.

Step 3 Configure an IP address on each server.

See Cisco IP Interoperability and Collaboration System Server Installation Guide for instructions.

Step 4 Configure both the primary and secondary servers with an NTP server to synchronize the system time.

۵. Note

HA servers use the internal time setting to exchange HA heartbeats and data. HA configuration fails if NTP is not configured on both servers.

a. Enter the following command on both servers to enable NTP:

ntpsetup -s enable <ntp-server> <backup-ntp-server>

For example:

ntpsetup -s enable ntp-sj1.cisco.com ntp-sj2.cisco.com

b. Enter the following command on both servers to verify the system settings:

ntpsetup -c

c. Enter the following command on both servers to restart the node manager:

service ipics_nm restart

- **Step 5** Install the following software licenses on the primary and secondary servers.
 - **a.** On the primary server, install the *IPICS Base server* and *Policy Engine Base* licenses. The primary server does not require the HA license (the HA license should never be installed on the primary server).
 - **b.** On the secondary server, install the *High Availability* license only. The secondary server does not require any other licenses.

See the "Managing Your Licenses" section in *Cisco IP Interoperability and Collaboration System Server Installation Guide*.

Step 6 Log in to the Cisco IPICS Administration Console for the primary server.

See the "Accessing the Administration Console" section on page 1-11 for instructions.

Step 7 From the Cisco IPICS Administration Console, navigate to the Configuration > High Availability > Security tab.

If HA partner trust has not yet been configured, the window shows that the Server Status is Not Trusted.

a. In the IP Address, field, enter the IP address of the HA Partner server.

The HA Partner is the secondary Cisco IPICS server. The User Name field displays the user name of the Partner Server linux administrator (ipicsadmin).

- b. In the User Password field, enter the password of the ipicsadmin user.
- **c.** Click **Save** to save the HA Partner IP address and establish trust between the redundant servers. To establish trust, the servers exchange public keys and SSL certificates.

If successful, the Server Status changes to Trusted and the HA Configuration tab appears.

Step 8 Click the HA Configuration tab:



- The HA Configuration window appears only after high availability mode is enabled, as described in Step 7.
- The IP addresses for the primary and secondary server are read-only. The primary server IP address is the Cisco IPICS server to which you are currently logged in. The secondary server IP address is the address you entered to enable high availability mode in Step 7.
- If high availability mode is enabled, but the HA is not yet configured, the Standby Server Status is Not Ready.
- **a.** In the Heartbeat Port field, enter the IP port number for HA heartbeat traffic. The default value is 3444.
- **b.** In the Heartbeat Interval (seconds) field, enter a number between 5 and 600 to define the number of seconds between heartbeats.

Each heartbeat checks to confirm the status and availability of the partner server. The default value is 15 seconds.

c. In the Missed Heartbeat Count field, enter number from 5 to 30 to define the number of missed heartbeats before the active role is transitioned to the secondary server.

The default is five missed heartbeats (75 seconds if the Heartbeat interval is 15 seconds). The transition process takes approximately 150 to 180 seconds.

- d. Click Update to save the changes.
- **Step 9** Click **Configure** to save the changes and activate high availability (server redundancy).

After a brief delay, you are automatically logged out of the IPICS system until HA configuration operations are complete.

- **Step 10** Verify that the **Standby Server Status** changes to Ready.
 - a. Log in to the Cisco IPICS Administration Console.
 - **b.** Navigate to the **Configuration > High Availability** window.
 - c. Click the HA Configuration tab.
 - d. Verify that the Standby Server Status is Ready.

L

Unconfiguring HA

Certain situations require you to temporarily unconfigure HA. These situations include restoring data from a Cisco IPICS database backup or generating SSL certificates for an IPICS server.

To unconfigure Cisco IPICS servers for HA, perform the following procedure:

Procedure

- Step 1From the Cisco IPICS Administration Console on the active server, navigate to the
Configuration > High Availability window.
- **Step 2** Click the **HA Configuration** tab.
- Step 3 Click the Unconfigure button.
- **Step 4** Click **Logout** to log out from the active server.
- **Step 5** Wait a few minutes for the active server to reconfigure.
- **Step 6** (Optional) Delete the security trust certificates used with the partner server:

- **Note** Delete the HA security certificates only if you need to regenerate or replace the SSL certificates. The default IPICS server SSL certificates expire after three years and must be replaced. See Appendix D, "Generating SSL Certificates" for more information.
- **a**. Log in again to the primary server.
- b. Navigate to the HA Security window (expand the Configuration drawer and choose High Availability > HA Security).
- c. Enter the password for the partner server in the password field to enable the Delete button.

A valid password is required to guarantee that the trust certificates are removed from the partner server.

d. Click Delete to delete the trust certificates used with the partner server.

HA Affect on the IDC Connections

When a user logs in to an IDC, the Log in selector displays all active and standby servers. To connect, users must select the active server. If the standby server is selected, the connection is denied and automatically redirected to the active server.



The role of active and standby can switch between the available servers. Users must always log in to the active server. The same user name and password is used for either server.

A failover occurs when the active server goes down or loses connectivity with the standby server. In this case, IDCs are automatically transferred to a standby server and service continues as normal. During the short time that the IDC is establishing a connection with the a standby server, the console operates in offline mode, using the last configuration provided by the previously active server. The console also

displays a message informing the user of the offline status. No action is required by the console user during this offline period. A new message appears when the console automatically reconnects to the new server.

After a failover, consoles remain connected to the new active server even if the original server is brought back online. The restored server becomes the new standby, ready take to control if a failover occurs. If the console user logs out, the user can log back in to the active server (as displayed in the log in server selector).

If a user attempts to log in to the standby server, the console is automatically redirected to the active server.



- Because there is a short pause in Cisco IPICS server operations during a failover, it is possible that not all data becomes synchronized from the active server to the standby. If this situation occurs, some configurations that were entered immediately before the failover could be lost. See the following example for more information.
- Photographs and videos that were uploaded to an incident shortly before a failover may need to be uploaded again because these files can be large and files are replicated at a lower priority than system data.

Example

In this example, a Cisco IPICS system is configured with redundant servers, as described in the "Configuring Cisco IPICS Servers for HA" section on page 10-3. The system has been running for 10 days with no problems. When dispatcher Dan starts his shift, he logs in to his console by selecting the active server from the login dialog. At lunch time, a localized earthquake occurs and the active server falls out of the rack. The console used by Dan detects the broken link to the server and automatically connects to the standby server, a process that takes about 10 seconds. Although Dan notices that the console is disconnected from all servers, the console remains connected to media resources and Dan is still able to notify Adam (the administrator) about the localized earthquake by using a console channel push-to-talk (PTT) feature. 10 seconds later, the console that is used by Dan automatically connects to the standby server. There is no visible change to the console.

45 minutes after the server failover, Adam the administrator remounts the original server involved in the earthquake and restores the server network connection. However, the console that is used by dispatcher Dan continues to connect to the new active server, and the restored server becomes the new standby.

The next day, when Dan returns to work, he logs in to the console and connects to the new active server, which is identified as the active server in the login server selector.

One issue did occur during this process. Before the earthquake, the channel Fire was assigned to user Ursula. Ursula saw the new channel because the console received an update from the active server. But before the channel assignment was propagated to the standby server, the active server was knocked off the network by the earthquake. Because the new Fire channel was not propagated to the standby before the failover occurred, the Fire channel does not appear on the new active server and must be reconfigured.

The next day, Amy attempts to log in to the console. Because she was not online at the time of the failover, her console still shows the original active server in the login server selector. When she attempts to connect to this server, which is now the standby, the application pauses for approximately 5 seconds. When the connection fails, the console is automatically redirected to the other server. The console connects to the new active server and the console is updated to reflect the new server status.

HA Affect On Connected Devices

If a Cisco IPICS server failover occurs, connected devices are affected in the following ways:

- Mobile client—Mobile clients automatically switch to the new active server.
- Cisco Unified IP Phone—The Cisco Unified IP Phone administrator must configure a Cisco Unified Phone Service for both the primary and secondary IPICS server. This configuration allows either the user or the administrator to subscribe to the services for both servers. (For more information, see the "Cisco Unified Phone Services" section in Cisco Unified Communications Manager System Guide.)

If you are using a Cisco Unified IP Phone and the active server goes down, the standby server attempts to redirect the IP phone connection. If the server cannot do so, you must manually reconnect to the active IPICS server. Open the **Services** menu, select the new active Cisco IPICS server, and log in again.

Standard telephone—Standard dial-in phone calls are dropped, but users can call back in to the system after a short delay.



Cisco IPICS does not automatically call back dial-out users participating in a VTG when a failover occurs.

Synchronizing the Server Time on HA Servers

Before servers are configured for HA, you must configure the internal time on each server using the Network Time Protocol (NTP). HA servers use the internal time setting to exchange HA heartbeats and data. HA configuration fails if NTP is not configured on both servers.

In addition, the time setting should not be manually changed on either HA server. If the time settings are more than 30 seconds apart, the servers can lose HA communication and enter a split brain scenario. See the "Resolving a Split Brain Scenario" section on page 10-10 for more information.

To synchronize the time settings on the primary or secondary servers, perform the following procedure:

Procedure

Step 1

Configure both the primary and secondary servers with an NTP server:

The NTP server sets the system time and ensures the active and standby server times are synchronized.

a. Enter the following command on both servers to enable NTP:

ntpsetup -s enable <ntp-server> <backup-ntp-server>

For example:

ntpsetup -s enable ntp-sj1.cisco.com ntp-sj2.cisco.com

b. Enter the following command on both servers to verify the system settings:

ntpsetup -c

Note

If the same files or records exist on both servers but have different timestamps, only the file or record with the most recent timestamp is retained. The older data is overwritten.

Step 2 Enter the following command on both servers to restart the node manager:

service ipics_nm restart

Step 3 If both servers are in *active* state due to a time misconfiguration, you must force the secondary server back into standby mode. See the "Resolving a Split Brain Scenario" section on page 10-10 for more information.

Performing a Manual Failover

To manually cause the currently active server and standby servers to reverse roles, perform the following procedure. The active role will be transferred to the standby server. This procedure is necessary if you want to take the active server offline, and ensures a stable failover between servers.

Before You Begin

- Either the primary or secondary server can be in the Active state. Use the following procedure to transfer the active role to the server currently in the Standby state.
- To perform a manual failover, HA must be configured and the Standby Server Status must be Ready. See the "Configuring Cisco IPICS Servers for HA" section on page 10-3 for more information.
- Before you begin, review the affect of a failover on consoles and connected devices. See the "HA Affect on the IDC Connections" section on page 10-6 and the "HA Affect On Connected Devices" section on page 10-8.

Procedure

Step 1 Log in to the Cisco IPICS Administration Console. See the "Accessing the Administration Console" section on page 1-11 for instructions.

- **Step 2** Navigate to the **Configuration > High Availability** window.
- **Step 3** Click the **HA Configuration** tab.

- **Note** If the HA Configuration tab is not available, the high availability mode is not enabled. Complete the instructions in the "Configuring Cisco IPICS Servers for HA" section on page 10-3.
- The Standby Server Status is Ready.
- The server in standby mode should assume the active role.

If the Standby Server Status is Not Ready, the high availability mode is enabled, but not configured. Complete Step 8 in the "Configuring Cisco IPICS Servers for HA" section on page 10-3.

Step 4 Click the **Failover Now** button to transfer the active state to the standby server.

The standby server becomes the new active server and the active server becomes the standby server.

L

Resolving a Split Brain Scenario

A *split brain scenario* occurs when the communication between the primary and secondary servers is lost, causing both servers to independently assume the active server role. Although this situation allows consoles and devices to connect with each server for continued operation, the data stored on each server is not synchronized with the other server. Over time, the data differences between the servers becomes greater.

When the communication link between the servers is reestablished, both servers remain in active state. This situation is known as a split brain scenario. To resolve this misconfiguration, one server must be returned to the standby state, and the data on the two servers must be reconciled.

Note

- Although either server can be returned to standby state, Cisco recommends that you keep the primary server in active state and returning the secondary server to standby, as described in the following instructions.
- A *split brain* scenario can occur if the system time on the servers is more than 30 seconds apart. See the "Synchronizing the Server Time on HA Servers" section on page 10-8 for information about synchronizing the server time using NTP.

Overview of Reconciliation Methods

The following methods are used to reconcile server data and restore server redundancy. Review the following descriptions to decide which method to use:

- Method 1: Force the Secondary Server into Standby State, page 10-10—Use this method if the servers have been in split brain mode for a short time (less than 5 days). This process entails forcing the secondary server into standby state, and then using Linux commands to synchronize the databases and file systems of the two servers. This method automatically returns the servers to high availability operation, and users can continue to access the primary server.
- Method 2: Reconfigure High Availability, page 10-12—Use this method if the servers have been in a split brain mode for 5 or more days, or if you are concerned that forcing the secondary into standby state (Method 1) might cause data loss. This method entails bringing down the secondary server, unconfiguring high availability on the primary server, and then manually reconciling the files that exist on both servers. After complete, you must reconfigure HA on both servers to restore server redundancy.
- Method 3: Workaround, page 10-13—Use this method only if Method 1 and Method 2 do not work. This method forces the primary to be the only active server but does not reconcile the data.

Method 1: Force the Secondary Server into Standby State

Use Method 1 to manually force the secondary server into standby state. After the secondary server is in standby state, use an SSH client to resynchronize the server databases and file systems.

Resynchronizing the server databases causes the following to occur:

- Records that exist on the primary server but not the secondary server are replicated from the primary server to the secondary server
- Records that exist on the secondary server but not on the primary server are replicated from the secondary server to the primary server

• If a record exists on both servers, the record on the primary server is considered the *master*, and the corresponding record on the secondary server is replaced

Resynchronizing the file system files causes the following to occur:

- Files that exist on the primary server but not on the secondary server are copied from the primary server to the secondary server.
- Files that exist on the secondary server but not on the primary server are deleted.
- If the same file exists on both servers, the file with the more recent timestamp (per coordinated universal time, or UTC) is retained and copied to the other HA server. The old version of the file is overwritten. If you prefer to merge the files instead of deleting the older file, ask your Linux systems administrator for information about reconciling file differences using scp or sftp.

This process automatically returns the servers to high availability operation, and users can continue to access the primary server.

Procedure

Step 1 (Optional) View the files and folders that are replicated during synchronization.

View the file at the following location:

/opt/cisco/ipics/conf/fileDirectory

This file contains the local source directories and the remote destination directories:

```
/idspri/backup/ /idspri/backup
/opt/cisco/ipics/tomcat/current/webapps/ipics_files/
/opt/cisco/ipics/tomcat/current/webapps/ipics_server/pmclogs/
/opt/cisco/ipics/tomcat/current/webapps/ipics_server/pmclogs
/idspri/archive/ /idspri/db_table_archive/
/idspri/db_table_archive
/opt/cisco/ipics/tomcat/current/webapps/documents/
/opt/cisco/ipics/tomcat/current/webapps/documents/
```



This file is used to determine what will be replicated. Do not alter the file in any way.

- **Step 2** Force the secondary server into standby state:
 - **a**. Determine which server was configured as the secondary server.

See the "Configuring Cisco IPICS Servers for HA" section on page 10-3.

b. Log in to the Cisco IPICS Administration Console for the secondary server.

See the "Accessing the Administration Console" section on page 1-11 for instructions.

- c. Navigate to the Configuration > High Availabily window.
- d. Click the HA Configuration tab:
- e. Verify that the Standby Server Status is Not Ready.
- f. Click the Go Standby button at the bottom of the window.

The **Go Standby** button is enabled only when both servers are in active mode, and communication between them has been reestablished.

g. Click **OK** when the confirmation message appears.

The secondary server logs out all current user sessions.

h. Wait for the split brain remediation process to complete.

The remediation process reestablishes the HA server pair with an active and standby server configuration.

When the process is complete, any login attempts to the secondary server are redirected to the primary server. See the "HA Affect on the IDC Connections" section on page 10-6 for more information.

- **Step 3** Resynchronize the database between the primary and secondary server:
 - **a.** Use an SSH client to log in to either the primary or secondary HA server with the user name **Informix**. The command will work the same on either server.
 - **b.** Start the database synchronization process by entering the following command:

server> /opt/cisco/ipics/database/bin/ipicsedr_control_repl REPAIR

c. Monitor the database synchronization process by entering the following command:

server> /sbin/service ipics ha-status

Note The database replication repair processes are concurrent and run in background. The status of the different processes are listed under "Pending Database Replication Synchronization Processes". The repair process is completed when the section has no outstanding entries.

- **Step 4** Resynchronize the file system files:
 - a. Use an SSH client to log in to the primary server with the user name ipicsadmin.



te The following commands must be run on the primary server.

b. Start the file system synchronization process by entering the following command:

server> /opt/cisco/ipics/database/bin/ipicsrsync run ipicsadmin

c. Monitor the synchronization process by entering the following command:

server> /opt/cisco/ipics/database/logs/rsync.log

Step 5 Log in to the Cisco IPICS Administration Console for the primary server and verify that the primary server is in active state and the secondary server is in standby state.

Method 2: Reconfigure High Availability

Use Method 2 to manually inspect the data integrity between the two servers and reconcile the data, if necessary. For more information, see the "Overview of Reconciliation Methods" section on page 10-10.

Procedure

Step 1	Log in to the secondary server using SSH.
Step 2	Enter the following commands to stop access to the secondary server:
	• server> /sbin/service ipics stop
	 server> /sbin/service ipics nm stop

Step 3 Take these actions to unconfigure HA on the primary server:

- a. Log in to the Cisco IPICS Administration Console for the primary server.
 See the "Accessing the Administration Console" section on page 1-11 for instructions.
- b. Expand the Configuration drawer and click High Availability.
- c. Click the HA Configuration tab.
- d. Click the Unconfigure button.
- e. Click Logout to logout from the active server.
- f. Wait a few minutes for the primary server to reconfigure.
- g. Log in again to the primary server.
- h. Navigate again to the HA Security screen (expand the Configuration drawer, click High Availability, and then click the HA Security tab.
- i. Click **Delete** in the **HA Security** screen to delete the HA security certificates and disable HA mode.
- **Step 4** Inspect the following directories in both servers to verify that files in both servers are the same:
 - /idspri/backup
 - /opt/cisco/ipics/tomcat/current/webapps/ipics_files
 - /opt/cisco/ipics/tomcat/current/webapps/ipics_server/pmclogs
 - /idspri/archive
 - /idspri/db_table_archive
 - /opt/cisco/ipics/tomcat/current/webapps/documents

Note The /documents directory contains all uploaded iPhone and IDC content, and is considered critical data.

Step 5 If the files are not the same, move the correct files to the primary server directories

Step 6 Reconfigure high availability on the primary and secondary servers.

See the "Configuring Cisco IPICS Servers for HA" section on page 10-3.



If the role of the original primary server is changed to the secondary role, new HA licenses are required.

Method 3: Workaround

If neither Method 1 nor Method 2 corrects the split brain scenario, use the following workaround to force the primary server to be the only server in active state.

Procedure

- **Step 1** Shut down the secondary server:
 - a. Log in to the secondary server using SSH.

b. Shut down the server using the following command:

shutdown -h now

- c. Power off the machine
- **Step 2** Stop and restart all services on the primary server:
 - a. Log in to the primary server using SSH
 - **b.** Restart the IPICS services using the following commands:

service ipics stop-all service ipics start-all

Step 3 (Optional) If the stop-all and start-all commands are not available on the primary server, enter the following alternative commands:

service ipics_nm stop
service ipics stop
service ipics start

(wait 10 seconds)

service ipics_nm start

Step 4 Power up the secondary server.

Reestablishing HA Configuration After Prolonged Server Downtime

When one of the servers in an HA server pair goes down for an extended length of time, the remaining active server saves database updates in a transaction log. Under normal operation, the database updates are restored to the down HA server when it comes back online.

If the second server remains offline for an extended time, however, and the system experiences heavy activity, the database log may run out of space. If the database log reaches 100 percent capacity, a database replication block state (DDRBLOCK state) can occur, blocking all database updates until cleared.

To prevent this situation, the Cisco IPICS system automatically unconfigures HA if usage exceeds 90 percent of capacity. An error message is also issued when capacity reaches 75 percent of capacity and again at 90 percent of capacity (when HA is also unconfigured.

If HA is unconfigured following extended downtime for one of the servers, you must reconfigure HA on the active server, as described in the "Configuring Cisco IPICS Servers for HA" section on page 10-3.