

Main page: [Cisco Unified MeetingPlace, Release 7.0](#)

Up one level: [Configuration](#)

Contents

- [1 Setting Up Failover for Two Newly Installed Application Servers](#)
 - ◆ [1.1 Before You Begin](#)
 - ◆ [1.2 Procedure](#)
 - ◆ [1.3 Verifying](#)
 - ◆ [1.4 Troubleshooting Tips](#)
 - ◆ [1.5 Related Topics](#)
 - ◆ [1.6 What To Do Next](#)
- [2 Setting Up Failover for One Existing Application Server and One Newly Installed Application Server](#)
 - ◆ [2.1 Before You Begin](#)
 - ◆ [2.2 Procedure](#)
 - ◆ [2.3 Verifying](#)
 - ◆ [2.4 Troubleshooting Tips](#)
 - ◆ [2.5 Related Topics](#)
 - ◆ [2.6 What To Do Next](#)
- [3 Configuring the Application Servers in a Failover Deployment](#)
 - ◆ [3.1 Before You Begin](#)
 - ◆ [3.2 Procedure](#)
 - ◆ [3.3 What to Do Next](#)
 - ◆ [3.4 Related Topics](#)

Setting Up Failover for Two Newly Installed Application Servers

Before You Begin

- Complete [Prerequisites for Application Server Failover](#).
- Read [Restrictions for Application Server Failover](#).

In summary, both servers should be ready for failover, but left in standby status. If the primary server is in failover active status, then enter **failoverUtil setServer standby** to bring it to standby status.

Procedure

1. Install the first Application Server (Node 1).
During installation, you configure the hostname and IP address of eth0, which is called "Ethernet Port 1(device eth0)" on the Network Setup page.
2. Install any licenses on the first Application Server (Node 1). See [About Licenses](#) for information about licenses.
3. Log in to the CLI of Node 1.
If you are logging in remotely, use the eth0 IP address or hostname.
4. Enter the following command to set up Node 1 for failover:
failoverUtil setDeployment failover

5. Follow the CLI prompts to configure the virtual network interface (eth0:0) with an IP address, subnet mask, default gateway, and hostname.
Node 1 automatically restarts and enters *standby* mode.
6. Install the second Application Server (Node 2), ensuring the following:
 - ◆ Node 1 and Node 2 use the exact same IP address and hostname for eth0.
 - ◆ Time is synchronized between Node 1 and Node 2.
7. Install any licenses on the second Application Server (Node 2) in active mode and reboot the system.
See [About Licenses](#) for information about licenses.
8. Log in to the CLI of Node 2.
If you are logging in remotely, use the eth0 IP address or hostname.
9. Enter the following command to set up Node 2 for failover:
failoverUtil setDeployment failover
10. Follow the CLI prompts to configure the virtual network interface (eth0:0) with an IP address, subnet mask, default gateway, and hostname.
Node 2 automatically restarts and enters *standby* mode.
11. Enter the following command to initialize database replication:
mp_replication init -s 1 -n 2 -l node2-eth0:0-hostname -r node1-eth0:0-hostname -h
12. Log in to the CLI of Node 1.
If you are logging in remotely, use the Node 1 eth0:0 IP address or hostname.
13. Enter the following commands to initialize and start database replication:
mp_replication init -s 1 -n 1 -l node1-eth0:0-hostname -r node2-eth0:0-hostname -h
mp_replication switchON -r node2-eth0:0-hostname
14. Enter the following command to change Node 1 from standby mode to active mode:
failoverUtil setServer active

Note: Node must be active to add licenses (step 7)

Verifying

Using the hostname or IP address of the virtual eth0:0 interface, log in to the Administration Center of each node, and verify that the correct failover deployment mode (active or standby) appears at the top of the page.

To verify that the replication was successful, log in to both nodes (eth0:0 interface) as root user and enter the **mp_replication status** command. The following should display in your output:

Local Server	Remote Server
<ul style="list-style-type: none"> • Replication Type: Local • Site: 1 • Node: 1 • State: Active • Status: Local 	<ul style="list-style-type: none"> • Replication Type: Intra-site • Site: 1 • Node: 2 • State: Active • Status: Connected

Troubleshooting Tips

- If the Administration Center in *both* nodes displays "Failover deployment. Standby server," something may have interrupted the process initiated in [Step 12](#). To resolve this issue, see the [failoverUtil_setServer](#) command description.
- If you switch on replication with sync via the mp_replication switchON command and replication fails, it is possible that a table is locking due to the processes on an active server. To resolve this issue, run the mp_replication switchON with sync command again.
- During this process, you may see several error messages. These error messages are expected behavior and informational only. The final message of the command explicitly says if the command ran successfully or not.

An example of an error message that you can ignore is:

```
* Error:
'/opt/cisco/meetingplace/database/versions/IIF.10.00.UC5XL/etc/sqlhosts not initialized. xxxxxx i
* Forcing the fix
```

The final message, which indicates if the process was successful, is:

```
Database Replication changes for this machine (Host = <hostname>; IP Addr = <ip addr>) is SUCCESS
```

Related Topics

- [Installing Cisco Unified MeetingPlace, Release 7.0](#)
- [Using the Command-Line Interface \(CLI\) in Cisco Unified MeetingPlace](#)

What To Do Next

Proceed to [Configuring the Application Servers in a Failover Deployment](#).

Setting Up Failover for One Existing Application Server and One Newly Installed Application Server

In this task:

- Node 1 is the existing Application Server.
- Node 2 is the new Application Server.

Before You Begin

- Complete [Prerequisites for Application Server Failover](#).
- Read the [Restrictions for Application Server Failover](#).

- Back up and archive the data on the existing Application Server (Node 1). See [Backing Up, Archiving, and Restoring Data on the Cisco Unified MeetingPlace Application Server](#).

Caution! Performing this task temporarily brings down your Cisco Unified MeetingPlace system. Proceed only during a scheduled maintenance period.

Procedure

1. Install the second Application Server (Node 2), ensuring the following:
 - ◆ To reduce system downtime during this procedure, enter the Node 2 eth0:0 (or any valid and currently unused) IP address and hostname when you configure "Ethernet Port 1(device eth0)" on the Network Setup page.

This configuration is temporary; you will change the eth0 IP address and hostname to the correct value later in this procedure.
 - ◆ Node 1 and Node 2 server times must be synchronized with each other.
2. Log in to the CLI of Node 1.

If you are logging in remotely, use the Node 1 eth0 IP address or hostname.
3. Enter the following command to set up Node 1 for failover:
failoverUtil setDeployment failover
4. Follow the CLI prompts to configure the virtual network interface (eth0:0) with an IP address, subnet mask, default gateway, and hostname.

Node 1 automatically restarts and enters *standby* mode.
5. Log in to the CLI of Node 2.

If you are logging in remotely, use the temporary Node 2 eth0 IP address or hostname that you entered during installation in [Step 1](#).
6. Enter the **net** command to change the Node 2 eth0 IP address and hostname to match the Node 1 eth0 IP address and hostname.

Because Node 1 is still in standby mode, Node 2 is now the active server.
7. Enter **mpx_sys restart** to restart Node 2.
8. Log in to the CLI of Node 2, this time using the shared eth0 IP address or hostname.
9. Enter the following command to set up Node 2 for failover:
failoverUtil setDeployment failover
10. Follow the CLI prompts to configure the virtual network interface (eth0:0) with an IP address, subnet mask, default gateway, and hostname.

Node 2 automatically restarts and enters *standby* mode.
11. Enter the following command to initialize database replication:
mp_replication init -s 1 -n 2 -l node2-eth0:0-hostname -r node1-eth0:0-hostname -h
12. Log in to the CLI of Node 1.

If you are logging in remotely, use the Node 1 eth0:0 IP address or hostname.
13. Enter the following commands to initialize and start database replication and synchronize existing data:
mp_replication init -s 1 -n 1 -l node1-eth0:0-hostname -r node2-eth0:0-hostname -h
mp_replication switchON -r node2-eth0:0-hostname -S -F node1-eth0:0-hostname

For more information about synchronization, see [About Synchronization](#).
14. Enter the following command to change Node 1 from standby mode to active mode:
failoverUtil setServer active

Verifying

Using the hostname or IP address of the virtual eth0:0 interface, log in to the Administration Center of each node, and verify that the correct failover deployment mode (active or standby) appears at the top of the page.

To verify that the replication was successful, log in to both nodes (eth0:0 interface) as root user and enter the **mp_replication status** command. The following should display in your output:

Local Server	Remote Server
<ul style="list-style-type: none"> • Replication Type: Local • Site: 1 • Node: 1 • State: Active • Status: Local 	<ul style="list-style-type: none"> • Replication Type: Intra-site • Site: 1 • Node: 2 • State: Active • Status: Connected

Troubleshooting Tips

- If the Administration Center in *both* nodes displays "Failover deployment. Standby server," something may have interrupted the process initiated in [Step 14](#). To resolve this issue, see the [failoverUtil setServer](#) command description.
- If you switch on replication with sync via the mp_replication switchON command and replication fails, it is possible that a table is locking due to the processes on an active server. To resolve this issue, run the mp_replication switchON with sync command again.
- During this process, you may see several error messages. These error messages are expected behavior and informational only. The final message of the command explicitly says if the command ran successfully or not.

An example of an error message that you can ignore is:

```
* Error:
'/opt/cisco/meetingplace/database/versions/IIF.10.00.UC5XL/etc/sqlhosts not initialized. xxxxxx i
* Forcing the fix
```

The final message, which indicates if the process was successful, is:

```
Database Replication changes for this machine (Host = <hostname>; IP Addr = <ip addr>) is SUCCESSF
```

Related Topics

- [Installing Cisco Unified MeetingPlace, Release 7.0](#)
- [Using the Command-Line Interface \(CLI\) in Cisco Unified MeetingPlace](#)

What To Do Next

Proceed to [Configuring the Application Servers in a Failover Deployment](#).

Configuring the Application Servers in a Failover Deployment

Whenever you need to configure the Application Server in a failover deployment, use this procedure to make sure that both the active and standby servers have the same configuration.

Before You Begin

- Complete one of the following tasks:
 - ◆ [Setting Up Failover for Two Newly Installed Application Servers](#)
 - ◆ [Setting Up Failover for One Existing Application Server and One Newly Installed Application Server](#)
- Determine the hostname or IP address of both of the following virtual network interfaces:
 - ◆ eth0:0 of the active server
 - ◆ eth0:0 of the standby server

Procedure

1. Go to **http://<application-server>/admin/**.
Use the shared eth0 hostname or IP address, which always goes to the *active* Application Server.
The "copyConfigFiles" commands will include recorded name files as well.
2. Log in as a **System administrator**.
3. Configure the active server.
Take note of any changes you make on pages or fields that are labeled as:
 - Copied
 - Server-specific
 - Partially replicated
4. If you made changes to any *copied* pages or parameters, perform the following steps:
 1. Log in to the CLI of the *active* server.
If you are logging in remotely, use the eth0 IP address or hostname.
 2. Enter the following command to compress and transfer the files from the active server to the standby server:
failoverUtil copyConfigFiles
 3. Log in to the CLI of the *standby* server.
If you are logging in remotely, use eth0:0 IP address or hostname.
 4. Enter the following command to decompress the transferred files and put them in the correct directories on the standby server:
failoverUtil restoreConfigFiles
5. If you made changes to any *server-specific* pages or parameters, perform the following steps.
Tip: If your workstation screen is large enough to accommodate two full web browser windows without overlapping, you can simultaneously view the Administration Center for both the active and standby servers. This may help you configure the server-specific

parameters to match between the active and standby servers.

1. Go to **http://standby-eth0:0/admin/**.
Use either the hostname or IP address of the virtual eth0:0 interface of the *standby* Application Server.
2. Log in as a System administrator.
3. Verify that "Failover deployment. Standby server." appears at the top of the page.
4. Configure the server-specific parameters on the standby server.

What to Do Next

Perform the following from both Application Servers to replicate custom prompts and some configurations from the Primary application server to the Standby application server during a maintenance window or off-peak hours.

1. Log in to the CLI of the active server.
If you are logging in remotely, use the eth0 IP address or hostname.
2. Enter the following command to compress and transfer the files from the active server to the standby server:
failoverUtil copyConfigFiles
This will compress and copy over configuration files and user prompts to the remote server.
3. Log in to the CLI of the standby server.
If you are logging in remotely, use eth0:0 IP address or hostname.
4. Enter the following command to decompress the transferred files and put them in the correct directories on the standby server:
failoverUtil restoreConfigFiles
This action copies over custom prompts including customized system prompts and user name recordings. Without this action .wav files are not replicated to the standby server.

Related Topics

- [Logging In to the Cisco Unified MeetingPlace Administration Center](#)
- [Using the Command-Line Interface \(CLI\) in Cisco Unified MeetingPlace](#)
- [Integrating with Microsoft Outlook](#)
- [Integrating Cisco Unified MeetingPlace with Cisco WebEx](#)
- [Customizing Music and Voice Prompts for Cisco Unified MeetingPlace](#)
- [How to Perform Application Server Failover](#)
- [Installing Cisco Unified MeetingPlace, Release 7.0](#)