Multi-node Deployment Administration

July 2, 2014

- Multi-node Deployment Models, page 17-1
- Cluster Topology Management, page 17-5
- High Availability Deployments, page 17-12
- High Availability Configuration, page 17-16
- Cluster-wide Routing Information on Cisco Unified Presence, page 17-24
- Static Route Configuration, page 17-24
- Presence Gateway Configuration on Cisco Unified Presence, page 17-29

Multi-node Deployment Models

You need to consider how you are going to deploy the multi-node feature in your network. You configure your desired multi-node deployment model in system topology management GUI in Cisco Unified Presence Administration. Choose **System > Cluster Topology** in Cisco Unified Presence Administration to access system topology management GUI.

This module provides an overview of the deployment model options for the multi-node feature, and provides examples of these deployments on system topology management GUI.

You only use system topology management GUI to configure your *local* Cisco Unified Presence cluster. See the intercluster peer module for information about configuring intercluster peer relationships with remote Cisco Unified Presence clusters.

Note The High Availability deployment models described in this module are only applicable to Cisco Unified Presence Release 8.5.x or later releases.

Balanced User Assignment Redundant High Availability Deployment

You can achieve a balanced mode High Availability deployment by evenly balancing users across all nodes in the subcluster, but only using up to 35% of the CPU of each Cisco Unified Presence server.

The balanced mode High Availability deployment option in a redundant mode supports up to fifteen thousand users per cluster. For example, if you have six Cisco Unified Presence nodes in your deployment, and fifteen thousand users, you assign 2.5 thousand users to each Cisco Unified Presence node.

When you use the balanced mode High Availability deployment option in a redundant mode, as compared to a non-redundant mode, only half the number of users are assigned to each node. However, if one node fails, the other node *will* handle the full load of the additional 50% of users in the subcluster, even at peak traffic. In order to support this failover protection, you must turn on High Availability in each of the subclusters in your deployment.

See Figure 17-1 for an example of this deployment model on system topology management GUI. In this example, there are 15,000 users in total, so 2500 users are evenly balanced across the six nodes.

Figure 17-1 Balanced User Assignment Non- Redundant High Availability Deployment

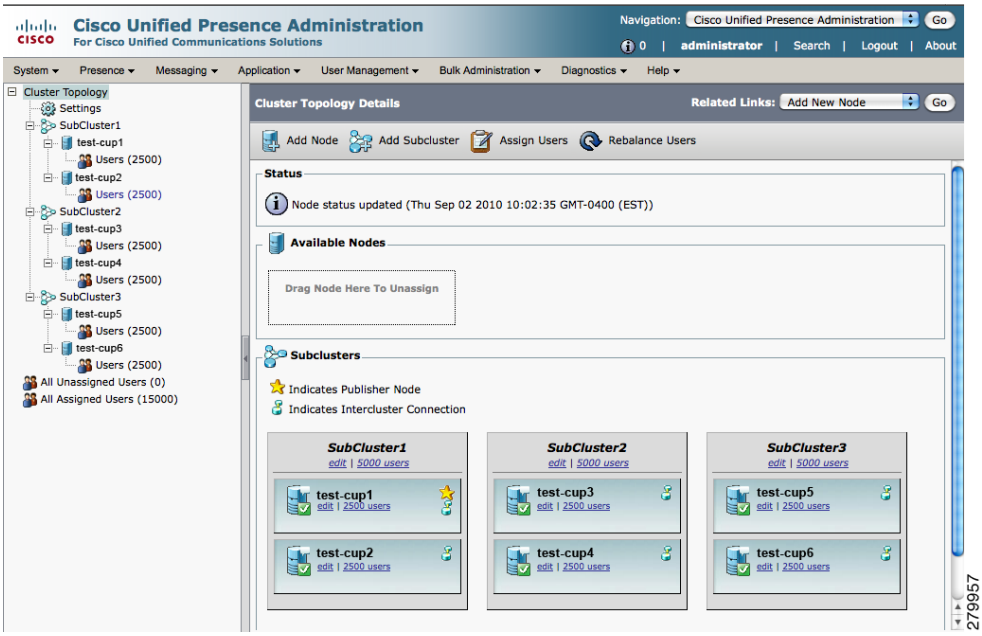

Related Topics

- Configure Routing Communication, page 7-12
- Create Subclusters in System Topology, page 17-9
- High Availability Deployments, page 17-12
- For the hardware user assignment guidelines for the multi-node feature, see the Cisco Unified Presence compatibility matrices at this URL:

  http://www.cisco.com/en/US/products/ps6837/products_device_support_tables_list.html

Active/Standby User Assignment Redundant High Availability Deployment

For this deployment model, assign all your users to the primary Cisco Unified Presence node, and none to the backup node. When you turn on High Availability in the subcluster, the backup node can handle all traffic from the primary node if the primary node fails.

See Figure 17-2 for an example configuration for this deployment model on system topology management GUI. In this example, there are 15,000 users in total, so 5000 users are assigned to the first node of each subcluster.

Figure 17-2 Active/Standby User Assignment High Availability Deployment

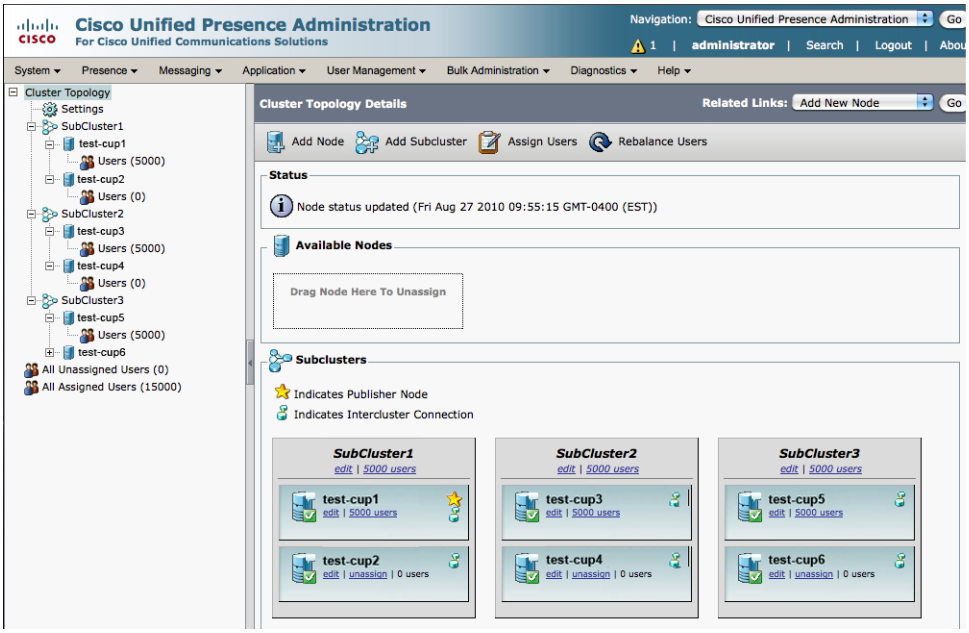

Related Topics

- Configure Routing Communication, page 7-12
- User Redistribution, page 1-8
- Create Subclusters in System Topology, page 17-9
- Cluster Topology Management, page 17-5
- High Availability Deployments, page 17-12
- For the hardware user assignment guidelines for the multi-node feature, see the Cisco Unified Presence compatibility matrices at this URL:

  http://www.cisco.com/en/US/products/ps6837/products_device_support_tables_list.html

High Availability for Cisco Unified Personal Communicator 7.x and 8.x Clients

Cisco Unified Presence provides server-side failback, which uses the same throttle mechanism as server failover. This feature detects when a failed Cisco Unified Presence server in a High Availability deployment comes back in service. It then sends terminating notify messages to Cisco Unified Personal Communicator clients that are failed over to initiate failback to their home node. Also, if a user is moved between nodes in the subcluster, the Cisco Unified Presence server sends terminating notify messages, and the client will sign out and sign in to the new node. To balance the load between two nodes in the subcluster, you can assign the users equally in each node.

Cisco Unified Personal Communicator Sign-in and Redirect

Cisco Unified Presence supports the ability to redirect a Cisco Unified Personal Communicator client application to the Cisco Unified Presence node to which the user is assigned (home node). The redirect feature is supported in intracluster and intercluster deployments. In both types of deployments, redirect occurs automatically when the client application signs in. After the user successfully signs in to the home node, Cisco Unified Personal Communicator caches the server name. As a result, redirect happens only once, unless a user is reassigned.

Using Figure 17-3 as a reference, see the following examples to gain a better understanding of the various redirect scenarios. In Figure 17-3, Cluster1 is assumed to be a Cisco Unified Presence Release 8.6 cluster and Cluster2 is a 7.x or 8.x cluster.

Figure 17-3 Intercluster and Intracluster Redirect Diagram

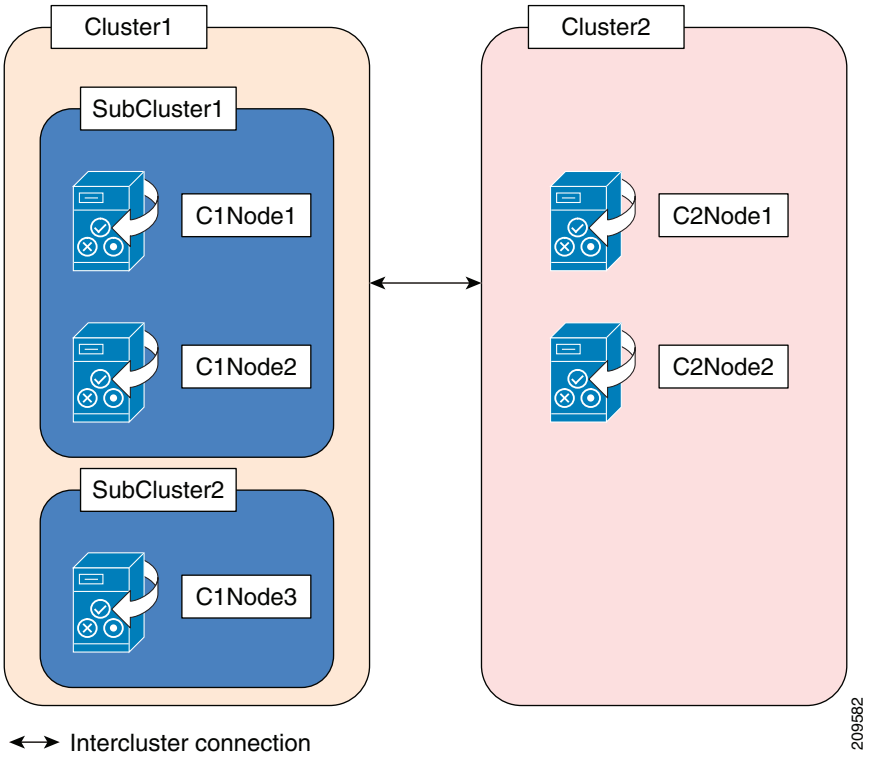

In the preceding figure, Cluster1 has three nodes, a publisher (C1Node1) and two subscribers (C1Node2 and C1Node3) and has an intercluster peer relationship with Cluster2, which contains a publisher (C2Node1) and subscriber (C2Node2). Several different redirect scenarios are possible:

1. A Cisco Unified Personal Communicator user is assigned C1Node1 as a home node and attempts to sign in to C1Node2. C1Node2 automatically redirects the Cisco Unified Personal Communicator client to C1Node1. In this scenario, High Availability is disabled in Subcluster1. If High Availability is enabled in Subcluster1, C1Node2 will process the login request. There is no redirect.

2. A Cisco Unified Personal Communicator user is assigned C1Node3 as a home node and attempts to sign in to C1Node1 or C1Node2. Regardless of whether High Availability is enabled in Subcluster1, C1Node1 or C1Node2 redirects the Cisco Unified Personal Communicator client to C1Node3. High Availability rules do not apply here because C1Node3 is part of Subcluster2.
3. A Cisco Unified Personal Communicator user is assigned C2Node1 or C2Node2 as a home node and attempts to sign in to C1Node1, C1Node2, or C1Node3. C1Node1, C1Node2, or C1Node3 automatically redirects the Cisco Unified Personal Communicator client its home node.

Note For more information about establishing intercluster peer relationships and syncing users, see Intercluster Peer Configuration, page 12-1.

Cluster Topology Management

This module is only applicable if you are deploying the multi-node feature. When you configure the multi-node feature, note the following:

- Perform the system topology configuration on the Cisco Unified Presence *publisher* node.
- Before configuring the system topology, read the multi-node planning and deployment information for best practice information about configuring this type of deployment.

Caution Only use the system topology interface to configure your *local* Cisco Unified Presence cluster. See the intercluster peer module for information about configuring intercluster peer relationships with remote Cisco Unified Presence clusters.

- Subcluster, Node and User Management Recommendations, page 17-5
- Add a New Node, page 17-7
- Create, Assign and Move Nodes in System Topology, page 17-10
- High Availability Deployments, page 17-12

Subcluster, Node and User Management Recommendations

- Node Creation and Movement Recommendations, page 17-5
- Node Name Recommendations, page 17-6
- Add a New Node, page 17-7
- Manual User Assignment Recommendations, page 1-7
- User Redistribution, page 1-8

Node Creation and Movement Recommendations

When you create nodes in system topology management GUI you can:

- Assign the nodes to a subcluster in Cisco Unified Presence, or allow the nodes to remain unassigned. These states are interchangeable.

- Assign Cisco Unified Presence users to the nodes, or allow the nodes to remain without any user assignments.
- Turn on or off High Availability on a subcluster. See the section about configuring High Availability deployments later in this chapter.
- Move a node from one subcluster to another if the node is assigned, has no users and high-availability is turned off in the subcluster.
- Move a node from one subcluster to another if the node is assigned and has no users.
- Configure real pingable nodes, or logical nodes which can be installed later and which remain inaccessible until that time.

To move nodes with users assigned, perform one of the following actions:

- Unassign the users, move the node, and then reassign the users to the node. Note that when you unassign the users, they will lose service.
- Create a logical node and move the users to the logical node. Move the node, reassign the users to the node, and remove the logical node.

Note

- Remove all users from a node before you unassign or move it.
- Turn off High Availability in the subcluster before you unassign or move a node in that subcluster.
- We strongly recommend that you perform any node movements that involve unassigning or moving a large numbers of users at off peak times. Such large operations can adversely impact performance.

Related Topics

- Node Name Recommendations, page 17-6
- Add a New Node, page 17-7
- Create, Assign and Move Nodes in System Topology, page 17-10

Node Name Recommendations

Cisco Unified Presence Release 8.6(4) and Earlier

By default, the name for a node is the hostname that you configure during the Cisco Unified Presence installation. For example, if the hostname of your Cisco Unified Presence node is called “cup1”, the node name is “cup1”. You can change the node name to the dotted IP address or the FQDN, for example, "192.168.0.1" or "cup1.acme.com". If you change the default name for the node, note the following:

- You must be able to resolve the hostname or the FQDN from the Cisco Unified Presence server, and Cisco Unified Personal Communicator client computers.
- If either Cisco Unified Presence server or the Cisco Unified Personal Communicator client computer cannot resolve the hostname or the FQDN, configure the IP address for the node name value.
- To test the name resolution from the Cisco Unified Presence server, use the command

  ```
  utils network ping <node_name>
  ```

- To test the name resolution from the Cisco Unified Personal Communicator client computer, use the command

  ```
  ping <node_name>
  ```

- If your network uses DNS that can map to IPv4 addresses, you can enter the Cisco Unified Presence hostname. Otherwise, you must enter the full IPv4 address of the Cisco Unified Presence server.

Cisco Unified Presence Release 8.6(5) and Later

If you are using DNS in your deployment, then the name of the publisher node is set to its FQDN by default. This is the Cisco recommended node name value for all subscriber nodes as it ensures that the node is fully resolvable from all clients and servers. The FQDN of a node is a concatenation of the hostname and domain that you configure during the Cisco Unified Presence installation. For example, if the hostname of your Cisco Unified Presence node is called "node1" and the DNS domain is "acme.com", the node name is "node1.acme.com". You can change the node name to the dotted IP address or just simply the hostname, for example, "192.168.0.1" or "node1.com". If you are using the hostname or FQDN as the node name, note the following:

- You must be able to resolve the hostname or the FQDN from all Cisco Unified Presence servers, across all clusters.
- You must be able to resolve the hostname or the FQDN from all Cisco Jabber client computers.
- If either the Cisco Unified Presence server or the Cisco Jabber client computer cannot resolve the hostname or the FQDN, consider the following options:
  - Configure the IP address for the node name value.
  - Make the required DNS configuration to ensure that the hostname or FQDN is resolvable from all client machines and all Cisco Unified Presence servers.
- To test resolution of the node name use the following commands:
  - From the Cisco Unified Presence server, use the command `utils network ping <node_name>`
  - From the Cisco Jabber client computer, use the command `ping <node_name>`

Note

- Cisco Unified Personal Communicator 7 does not support FQDN based node names. If you are planning to deploy this client, you must change the node name to either the IP address or the hostname.
- When using the Cisco Jabber client, certificate warning messages can be encountered if the IP address is configured as the IM and Presence Service node name. To prevent Cisco Jabber from generating certificate warning messages, the FQDN should be used as the node name.

Related Topics

- Node Creation and Movement Recommendations, page 17-5
- Create, Assign and Move Nodes in System Topology, page 17-10
- *Changing the IP Address, Hostname and Domain Name for Cisco Unified Presence* on cisco.com

Add a New Node

Follow this procedure if you need to add new nodes after a multi-node deployment is running.

You must create the new node in your topology before you install the node, specifically before you install the Cisco Unified Presence software on the new node. However, you cannot assign the new node to a subcluster before you install Cisco Unified Presence software on the new node.

Note You no longer have to manually add Cisco Unified Presence as an Application Server on Cisco Unified Communications Manager. When you add or remove a node on the system topology management GUI, the node is automatically added to or removed from the Application Server list on Cisco Unified Communications Manager.

Before You Begin

Check the following:

- From System troubleshooter page, verify that the Cisco UP Replication Watcher service is running on all nodes.
- On the **Network services** screen in Cisco Unified Serviceability (on the subscriber node), verify that all Cisco Unified Presence services are running.
- High Availability is turned off in a subcluster before you move or unassign a node in that subcluste

Restrictions

Your hardware must comply with the multi-node hardware recommendations.

Procedure

Step 1 Create a new subcluster in system topology management GUI (if required).

Step 2 Create a new node in system topology management GUI.

Step 3 Install the Cisco Unified Presence software on the new node.

See the *Installation Guide for Cisco Unified Presence* for the installation procedure.

Step 4 Assign the node to the subcluster (if required).

Cisco Unified Presence assigns the node to the cluster, but the node will not receive traffic until you assign users to it.

Step 5 Turn on High Availability in the subclusters as required.

Step 6 Assign users from other nodes to the new node as required.

Related Topics

- Clustering over WAN for Intracluster and Intercluster Deployments, page 2-4
- Create Subclusters in System Topology, page 17-9
- Create, Assign and Move Nodes in System Topology, page 17-10
- Configure User Assignment in System Topology, page 15-1
- High Availability Deployments, page 17-12

Expand the Cluster

Before You Begin

Check the following:

- From System troubleshooter page, verify that the Cisco UP Replication Watcher service is running on all nodes.
- On the **Network services** screen in Cisco Unified Serviceability (on the subscriber node), verify that all Cisco Unified Presence services are running.
- High Availability is turned off in a subcluster before you move or unassign a node in that subcluste

Restrictions

- Your hardware must comply with the multi-node hardware recommendations.
- We strongly recommend that you perform any node movements that involve you unassigning or moving a large numbers of users at off peak times. Such large operations can adversely impact performance.

Procedure

Step 1 Create the new subcluster(s) in system topology management GUI (if required).

Step 2 Create the new nodes in system topology management GUI.

Step 3 Install each new node.

Step 4 Assign the nodes to the (new) subclusters.

Step 5 Turn on High Availability in the subclusters as required.

Step 6 Once all the nodes are online, assign users to the new nodes using the following user assignment options:

- Using the Find User Assignment feature, unassign chosen users from each node, and use the User Assignment Mode parameter to reassign new users to new subcluster(s) and nodes.
- Using the Find User Assignment feature, manually move users to new nodes.
- Unassign all users, and then reassign the users to the cluster using the appropriate User Assignment Mode parameter setting for the whole cluster.

Related Topics

- Clustering over WAN for Intracluster and Intercluster Deployments, page 2-4
- Create Subclusters in System Topology, page 17-9
- Create, Assign and Move Nodes in System Topology, page 17-10
- Configure User Assignment in System Topology, page 15-1
- High Availability Configuration, page 17-16

Create Subclusters in System Topology

The system automatically assigns the first Cisco Unified Presence node that you install as the publisher node. After you install the publisher node, create the required subclusters and subsequent nodes in your Cisco Unified Presence cluster in system topology management GUI. After a subcluster has been created, you can update or view the status of the subcluster by clicking **Edit**.

Repeat this procedure for each subcluster that you require for your deployment.

Note Perform this procedure on the publisher Cisco Unified Presence node.

Before You Begin

Plan your multi-node deployment model.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Cluster Topology**.

Step 2 Click **Add New Subcluster**.

Step 3 Define a unique name for the subcluster.

Step 4 Click **Save**.

Related Topics

- Cisco Unified Presence Planning Requirements, page 3-1
- High Availability Deployments, page 17-12

Create, Assign and Move Nodes in System Topology

Create the required subsequent nodes for your deployment. By creating the subsequent nodes in the topology view of the publisher node, Cisco Unified Presence associates the subsequent nodes with the publisher node.

Note
- Perform this procedure on the publisher Cisco Unified Presence node.
- Perform this procedure *before* you install any of the subsequent Cisco Unified Presence nodes. If you assign a subsequent Cisco Unified Presence node to a subcluster prior to installing it, users in remote clusters will not receive availability information. An availability outage will occur until the node is installed.

Before You Begin

- Create the required subclusters for your deployment.
- Depending on how you plan to configure your node name, obtain the required value for your nodes (for example hostname, dotted IP address, FQDN or DNS-SRV).

Restrictions

- If you wish to change the default node name, there are certain node name restrictions. Read the node name recommendations topic.
- You can only move a node from one subcluster to another if the node is assigned and has no users.
- You must turn off High Availability in a subcluster before you move or unassign a node in that subcluster.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Cluster Topology**.

Step 2 Create the required subsequent nodes for your deployment:

a. Click **Add New Node**.

b. Define a unique name for the node.

c. Click **Save**.

Step 3 Perform one of these actions:

If you want to:	Action	Notes
Assign a node to a subcluster	Drag the node into the empty slot in the subcluster	• Do not assign the subsequent node to a subcluster until *after* you install it, and you have checked the status of the node. • Before you assign a node to a subcluster, check the following – From System troubleshooter page, verify that the Cisco UP Replication Watcher service is running on all nodes. – On the Network services screen in Cisco Unified Serviceability (on the subscriber node), verify that all Cisco Unified Presence services are running on the assigned node.
To move a previously assigned node.	Drag the node from the subcluster and drop it into the empty slot of the peer subcluster.	• Turn off high -availability in the subcluster before you move the node. • Unassign all users from the node before you move it.
To update or view the status of a node.	Click the **edit** link on the node to view the **Node Detail** screen.	• View the total users assigned to the node. • Verify the status of the node. • If you turn on High Availability in the subcluster, the critical services that Cisco Unified Presence monitors on the node for failover are marked in the 'Monitored' column. • If you turn on High Availability, you can also view the High Availability state of the node, and the reason for this state.

Related Topics

- DNS Domain Configuration, page 7-3
- Add a New Node, page 17-7
- Node Name Recommendations, page 17-6

- Node Creation and Movement Recommendations, page 17-5
- High Availability Deployments, page 17-12
- Intercluster Peer Configuration, page 12-1

High Availability Deployments

Requirements for High Availability

The requirements for High Availability are:

- You must be running Cisco Unified Presence release 8.5 (x), or a later 8.x release. Any earlier Cisco Unified Presence 8.0(x) releases do not support High Availability.
- Cisco Unified Presence supports High Availability at a subcluster level. Both nodes in the subcluster must be running the same version of Cisco Unified Presence 8.x software for High Availability to work.

High Availability in a Subcluster

Cisco Unified Presence supports High Availability in a subcluster meaning if a node in the subcluster fails, the Instant Message and Availability services from that node can failover to the second node in the subcluster.

You must *manually* turn on High Availability in a subcluster on the Cluster Topology interface on Cisco Unified Presence Administration interface. On the main Cluster Topology interface, the subcluster icon indicates that you have turned on High Availability on the subcluster.

A green tick beside the High Availability icon indicates that High Availability in the subcluster is running normally. A red ‘x’ beside the High Availability icon indicates that the subcluster is in a failed state.

Cisco Unified Presence automatically detects failover in a subcluster by monitoring the heartbeat and monitoring the critical services on the peer node. When Cisco Unified Presence detects failover, it automatically moves all users to the backup node. From the Cisco Unified Presence Administration interface, you can initiate a manual fallback to the primary node. Cisco Unified Presence Release 8.6(4) and later supports automatic fallback to the primary node after failover.

Caution Cisco Unified Presence Release 8.6(3) and earlier does not perform an automatic fallback to the primary node after failover. You must manually perform the fallback from the Cluster Topology interface, otherwise the users that were moved will remain on the backup node.

Note Cisco Unified Presence performs an automatic fallback when the backup activated node fails due to a critical service failure and the peer node is in the “Failed Over” state and supports the automatic recovery fallback.

To monitor and troubleshoot the status of the High Availability functionality on a subcluster, view the High Availability states that Cisco Unified Presence assigns to each node. See Node State Definitions, page 17-21 and Node States, Causes and Recommended Actions, page 21-1 for descriptions of these states and recommended actions if the subcluster is in a failed state. If a failover occurs, on the node detail screen, Cisco Unified Presence marks the users that have failed over to the backup node.

Related Topics

- Automatic Failover Detection, page 17-13
- Automatic Fallback, page 17-15
- Manual Failover and Fallback, page 17-16
- High Availability Configuration, page 17-16

Impact of Failover to Cisco Unified Presence Clients and Services

Cisco Unified Presence supports High Availability for Cisco Unified Personal Communicator Release 7.x and Cisco Unified Personal Communicator Release 8.5(x) and later.

During failover to the backup node, availability and instant messaging services are temporarily unavailable on client applications. After failover is complete, the availability and instant messaging services become available on the client again when the client signs back in. Similarly, if fallback occurs, availability and instant messaging services are temporarily unavailable on client applications until fallback completes and the client signs back in. Cisco Unified Personal Communicator signs users back in automatically.

The impact of failover on temporary adhoc chat messages depends on the particular client application. On Cisco Unified Personal Communicator, any adhoc chat windows that were open before failover should display again after the failover is complete. However, if all of the users in a chat room automatically exit the chat room as part of a failover or fallback process, or if the adhoc chat room is hosted on a failed node, the adhoc chat windows will not display again after failover and a message is displayed explaining that the chat room was deleted. On all clients, any persistent chat rooms that users create on the failed node cannot be accessed again until recovery.

If Cisco Unified Personal Communicator is operating in softphone mode (the user is on a voice call) during failover, the voice call is not disconnected.

Automatic Failover Detection

Cisco Unified Presence uses these methods to automatically detect if a node fails:

- **Peer Heartbeat** - In a subcluster, each node sends heartbeat intervals to the other node to check if the node is up and running. If a node detects a loss of heartbeat in the peer node, the node initiates a failover. You can configure the heartbeat interval and the heartbeat timeout from the Service Parameters page on Cisco Unified Presence Administration interface.
- **Monitor Critical Services** - Each node monitors a list of critical services. If the node detects that any critical service is not running for a configurable outage period (ninety seconds is the default value), it instructs the peer node to initiate a failover. You can configure this critical service delay from the Service Parameters page on Cisco Unified Presence Administration interface. These are the list of critical services that the node monitors:
  - Cisco DB (internal IDS database)
  - Cisco UP Presence Engine (if you activate this service)

- Cisco UP XCP Router
- Cisco UP Message Archiver (if you integrate Cisco Unified Presence with a third-party off-board database, and you activate this service)
- Cisco UP SIP Proxy (if you configure SIP federation, enable Partitioned Intradomain Federation, or you have an intercluster connection with a Cisco Unified Presence Release 7.x cluster, and you activate this service)
- Cisco UP XCP SIP Federation Connection Manager (if you configure SIP federation, enable Partitioned Intradomain Federation, or you have an intercluster connection with a Cisco Unified Presence Release 7.x cluster, and you activate this service)
- Cisco UP Presence Datastore—Cisco Unified Presence Release 8.6(4) or later only
- Cisco UP Route Datastore (if you configure SIP federation, enable Partitioned Intradomain Federation, or you have an intercluster connection with a Cisco Unified Presence Release 7.x cluster, and you activate this service)—Cisco Unified Presence Release 8.6(4) or later only

You can view the critical services that Cisco Unified Presence monitors for failover on the node details screen on the Cluster Topology interface. The critical services that Cisco Unified Presence monitors are marked in the 'Monitored' column in the services list.

Note

- Cisco Unified Presence only detects a failover if a critical service is not running for the duration of the outage period. It does not detect a failover in the case where one or more critical services are not running during the outage period, but not for the duration of the outage period, for example, a rolling outage. In this case, Cisco Unified Presence generates alarms indicating that services are starting and stopping, and you can perform a manual failover on Cisco Unified Presence.
- If you manually stop a critical service, and the service is stopped for longer than the permitted outage period, failover will occur.

Prior to Cisco Unified Presence Release 8.6, if Cisco Unified Presence detects the situation where both nodes in the subcluster think they own the same user, both nodes go into a failed state, and you need to perform a manual recovery from the Cluster Topology interface. In Cisco Unified Presence Release 8.6, manual recovery is not required. When the network issue is resolved, auto-recovery occurs without administrator intervention.

If manual recovery is required for another reason, you may experience IDS replication delays.

To check the status of the IDS replication on a node either:

- Use this CLI command:

  ```
  utils dbreplication runtimestate
  ```

- Use the Cisco Unified Reporting Tool (CURT). The 'Unified CUP Database Status' report displays a detailed status of the cluster.

Related Topics

- Perform a Manual Failover to Backup Node, page 17-21
- Configure the Advanced Service Parameters for the Server Recovery Manager, page 17-18
- Cisco UP Replication Watcher Service, page 22-3

Automatic Fallback

Cisco Unified Presence Release 8.6(4) and later supports automatic fallback to the primary node after a failover. Automatic fallback is the process of moving users back to the primary node after a failover without manual intervention. You can enable automatic fallback with the Enable Automatic Fallback service parameter on the Cisco Unified Presence Administration interface.

Automatic fallback occurs in the following scenarios:

- A critical service on Node A fails—A critical service (for example, the Presence Engine) fails on Node A. Automatic failover occurs and all users are moved to Node B. Node A is in a state called "Failed Over with Critical Services Not Running". When the critical service recovers, the node state changes to "Failed Over." When this occurs Node B tracks the health of Node A for 30 minutes. If no heartbeat is missed in this time frame and the state of each node remains unchanged, automatic fallback occurs.
- Node A is rebooted—Automatic failover occurs and all users are moved to Node B. When Node A returns to a healthy state and remains in that state for 30 minutes automatic fallback occurs.
- Node A loses communications with Node B—Automatic failover occurs and all users are moved to Node B. When communications are re-established and remain unchanged for 30 minutes automatic fallback occurs.

If failover occurs for a reason other than one of the three scenarios listed here, you must recover the node manually. If you do not want to wait 30 minutes before the automatic fallback, you can perform a manual fallback to the primary node.

Related Topics

- High Availability Deployments, page 17-12
- Configure the Advanced Service Parameters for the Server Recovery Manager, page 17-18
- Perform a Manual Fallback to Primary Node, page 17-22

Cisco UP Server Recovery Manager (SRM)

The Cisco UP Server Recovery Manager (SRM) on Cisco Unified Presence manages the failover between nodes in a subcluster. The Cisco UP Server Recovery Manager manages all state changes in a node; state changes are either automatic or initiated by the administrator (manual).

After you turn on High Availability in a subcluster, the Cisco UP Server Recovery Manager on each node establishes heartbeat connections with the peer node, and begins to monitor the critical processes.

The SRM is responsible for the user move operations after it detects that failover has occurred. It is the SRM on the peer node, not on the failed node, that performs the user move operation. For example, if node A fails, the SRM on node B performs the user move operation. The SRM throttles the number of users moved to the peer node, it moves the users in batches or iterations. You can configure the number of users that the SRM moves per iteration (the default value is 25). On failover, the SRM will move users that are signed in first, and then move users that are not signed in. Note that if you initiate a fallback, or for Cisco Unified Presence Release 8.6(4) or later automatic fallback occurs, users that are not signed in are moved first, and then users that are signed in.

If the SRM is not turned on, it does not monitor any critical processes, nor does it monitor the heartbeat connections with the peer node.

Caution Before you turn on High Availability in a subcluster, you must configure the SRM service parameters to properly reflect your deployment, see High Availability Client Login Profiles, page I-1.

Related Topics

- High Availability Deployments, page 17-12
- Configure the Advanced Service Parameters for the Server Recovery Manager, page 17-18

Manual Failover and Fallback

From the Cluster Topology interface, you can perform the following procedures:

- Initiate a manual failover for a subcluster. When you initiate a manual failover, the Cisco UP Server Recovery Manager stops the critical services on the failed node, and moves all users to the backup node.
- Initiate a manual fallback from the Cluster Topology interface, where the Cisco UP Server Recovery Manager restarts critical services on the primary node and moves users back to the primary node.
- Perform a manual recovery for a subcluster (when both nodes in the subcluster are in a failed state). When you perform a manual recovery, Cisco Unified Presence restarts the Cisco UP Server Recovery Manager service on both nodes in the subcluster.

Related Topics

- Perform a Manual Failover to Backup Node, page 17-21
- Perform a Manual Fallback to Primary Node, page 17-22
- Perform a Manual Recovery of a Subcluster, page 17-23

Important Note about High Availability and Intercluster Deployments

When failover occurs, the Intercluster Sync Agent is responsible for communicating the user move information to other clusters. The Intercluster Sync Agent runs on both the publisher and subscriber nodes in a cluster. In an Active-Standby configuration, if the publisher node fails or the Intercluster Sync Agent on the publisher node fails, the Intercluster Sync Agent on the subscriber node becomes Active and resumes synchronization, meaning the other clusters will continue to receive the information that users have moved to a different node. Intercluster availability and IM continue to work. Users that have failed over will receive availability information for remote users. Remote users continue to receive availability information and IMs from users that have failed over, and all IMs they send to a failed over user are delivered. When the publisher node recovers, the publisher falls back to Active mode and the subscriber returns to Standby mode.

High Availability Configuration

- Turn On or Off High Availability for a Subcluster, page 17-17
- Configure the Advanced Service Parameters for the Server Recovery Manager, page 17-18
- Configuration Verification, page 17-20
- Perform a Manual Failover to Backup Node, page 17-21

- Perform a Manual Fallback to Primary Node, page 17-22
- Perform a Manual Recovery of a Subcluster, page 17-23

Turn On or Off High Availability for a Subcluster

Caution Before you turn on High Availability in a subcluster, you must configure the SRM service parameters to properly reflect your deployment, see High Availability Client Login Profiles, page I-1.

You have to manually turn on High Availability in a subcluster; Cisco Unified Presence does not turn on High Availability in a subcluster by default. You can turn on High Availability in a subcluster when:

- there are two nodes in the subcluster, *and*
- both nodes have IP addresses that are resolvable addresses, *and*
- both nodes are running Cisco Unified Presence Release 8.5 or higher.

You can either assign users to the nodes in the subcluster before or after you turn on High Availability for the subcluster.

Before You Begin

- Configure the subclusters and nodes in your network, and assign nodes to the subclusters.
- Make sure critical services are running on both nodes in the subcluster before you turn on high-availability in a subcluster. If one or more critical services are not running on a node, when you turn on High Availability, that node will failover to the backup node. When one or more critical services are not running on one node in a subcluster, but all critical services are running on the second node, the subcluster will go into a failed state after you turn on High Availability.

Restriction

- You can only turn on High Availability in a subcluster when there are two nodes assigned to that subcluster. The High Availability checkbox does not display when there are no nodes, or one node, assigned to the subcluster.
- You can only turn off High Availability when the nodes in the subcluster are not in a transition state (Failing Over, Falling Back). If you turn off High Availability in a subcluster when either node is in a failed over scenario (Failed Over, Failed), users that Cisco Unified Presence fails over to the backup node are homed to the backup node. Cisco Unified Presence will not move these users back to the primary node, they remain on the secondary node.

Procedure

Step 1 **Cisco Unified Presence Administration > System > Cluster Topology.**

Step 2 Click the **edit** link on the appropriate subcluster.

Step 3 Check **Enable High Availability**.

Note If the nodes in the subcluster are not in a transition state (Failing Over, Falling Back), you can turn off High Availability for the sublcluster by unchecking **Enable High Availability**.

Step 4 Click **Save**.

Cisco Unified Presence displays the following information about High Availability for the subcluster

Field	Description
Monitored Node	The node in the subcluster that Cisco Unified Presence is monitoring for failover detection.
Node State	The state of the node. See High Availability Deployments, page 17-12 for definitions of the states.
Node Reason	The reason for the node state.
Node Action	The action you can take to change the state of the node: • **Fallback** - This option is displayed for nodes that are in Idle or Failed Over states. Choose to manually initiate a fallback to this node. • **Failover** - This option is displayed for nodes that are in Normal state. Choose to manually initiate a failover to this node. • **Recovery** - This option is displayed if both nodes in the subcluster are in a failed state. Choose to manually initiate a recovery of the subcluster where Cisco Unified Presence restarts the SRM service on both nodes.

Configure the Advanced Service Parameters for the Server Recovery Manager

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Service Parameters.**

Step 2 Choose a Cisco Unified Presence server from the Server menu.

Step 3 Choose Cisco UP Server Recovery Manager from the Service menu.

Step 4 Configure these service parameters:

Parameter	Description	Additional Information
Service Port	This parameter specifies the port that Cisco UP Server Recovery Manager uses to communicate with its peer.	If you modify this parameter, Cisco Unified Presence restarts the Cisco UP Server Recovery Manager on all nodes in the cluster.
Admin RPC Port	This parameter specifies the port that Cisco UP Server Recovery Manager uses to provide admin RPC requests.	If you modify this parameter, Cisco Unified Presence restarts the Cisco UP Server Recovery Manager on all nodes in the cluster.

Parameter	Description	Additional Information
Critical Service Down Delay	This parameter determines the duration a critical service can be down before Cisco Unified Presence initiates an automatic failover.	If you change this value, this affects how long a critical service can be down before Cisco Unified Presence initiates an automatic failover.
Enable Automatic Failover	This parameter turns automatic failover on or off on Cisco Unified Presence.	This parameter is on by default. Turn this parameter off only if you do not want automatic failover on Cisco Unified Presence and you only want to perform manual failover. **Note**: This parameter applies to Cisco Unified Presence Release 8.6(3) and earlier only.
Enable Automatic Fallback	This parameter turns automatic fallback on or off on Cisco Unified Presence.	This parameter is off by default. Turn this parameter on only if you want to enable automatic fallback on Cisco Unified Presence. **Note**: This parameter applies to Cisco Unified Presence Release 8.6(4) and later only.
Initialization Keep Alive (Heartbeat) Timeout	This parameter specifies the duration that the heartbeat is lost with the peer node (SRM) when the peer SRM restarts and is in the initialization state.	Cisco recommends that you configure this value to at least twice the value of the Keep Alive (Heartbeat) Timeout in order to avoid unnecessary failovers.
Keep Alive (Heartbeat) Timeout	This parameter specifies the duration that the heartbeat is lost with the peer node (SRM) before Cisco Unified Presence initiates an automatic failover.	Cisco recommends that you configure this value to at least twice the value of KeepAliveInterval value. If this value is too close to the KeepAliveInterval value, this can cause a failover to occur.
Keep Alive (Heart Beat) Interval	This parameter specifies the interval between keep alive (heartbeat) messages sent to the peer node.	N/A
Users Moved Per Iteration	This parameter specifies the number of users that Cisco Unified Presence moves for each iteration when it performs a failover or a fallback. There is a delay of one second between each iteration.	Increasing this value will shorten the failover time at the expense of CPU. Lowering the value will lengthen failover time, but have less impact on the CPU. **Caution** Before you configure the Users Moved Per Iteration parameter value, refer to the High Availability Client Login Profiles, page I-1.

Parameter	Description	Additional Information
Client Re-Login Lower Limit	This parameter specifies the minimum number of seconds which Cisco Unified Personal Communicator will wait before attempting to re-login to this Cisco Unified Presence server. This waiting time occurs due to the failure of a node or a critical service on a node. **Note** This parameter is per node.	This parameter only applies to Cisco Unified Personal Communicator Release 8.5 or higher 8.x releases. **Caution** Refer to the High Availability Client Login Profiles, page I-1 for guidelines on defining the client re-login lower and upper limits.
Client Re-Login Upper Limit	This parameter specifies the maximum number of seconds which Cisco Unified Personal Communicator will wait before attempting to re-login to this Cisco Unified Presence server. This waiting time occurs due to the failure of a node or a critical service on a node. **Note** This parameter is per node.	This parameter only applies to Cisco Unified Personal Communicator Release 8.5 or higher 8.x releases. **Caution** Refer to the High Availability Client Login Profiles, page I-1 for guidelines on defining the client re-login lower and upper limits.

Step 5 Click **Save**.

Related Topic

High Availability Client Login Profiles, page I-1

Configuration Verification

After you have configured High Availability, you can check the status of nodes.

Verify the Cisco UP Service Recovery Manager Service is Running

When you turn on High Availability in a subcluster, Cisco Unified Presence restarts the Cisco UP Service Recovery Manager service and it begins to monitor for failover detection. To verify this service is running, choose **Cisco Unified Serviceability > Tools > Control Center - Network Services**.

Node State Definitions

Table 17-1 describes the different node states and associated reasons. You can view the state of an existing node by either viewing the node details or the subcluster details on the Cluster Topology interface

Table 17-1 Node State Descriptions

State	Description
Initializing	This is the initial (transition) state when the Cisco UP Server Recovery Manager service starts; it is a temporary state.
Idle	Cisco Unified Presence is in Idle state when failover occurs and services are stopped. In Idle state, the Cisco Unified Presence node does not provide any availability or Instant Messaging services. In Idle state, you can manually initiate a fallback to this node from the Cluster Topology interface.
Normal	This is a stable state. The Cisco Unified Presence node is operating normally. In this state, you can manually initiate a failover to this node from the Cluster Topology interface.
Running in Backup Mode	This is a stable state. The Cisco Unified Presence node is acting as the backup for its peer node. Users have moved to this (backup) node.
Taking Over	This is a transition state. The Cisco Unified Presence node is taking over for its peer node.
Failing Over	This is a transition state. The Cisco Unified Presence node is being taken over by its peer node.
Failed Over	This is a stable state. The Cisco Unified Presence node has failed over, but no critical services are down. In this state, you can manually initiate a fallback to this node from the Cluster Topology interface.
Failed Over with Critical Services Not Running	This is a stable state. Some of the critical services on the Cisco Unified Presence node have either stopped or failed.
Falling Back	This is a transition state. The system is falling back to this Cisco Unified Presence node from the node running in Backup Mode.
Taking Back	This is a transition state. The failed Cisco Unified Presence node is taking back over from its peer.
Running in Failed Mode	An error occurs during the transition states or Running in Backup Mode state.
Unknown	State unknown.

Perform a Manual Failover to Backup Node

You can perform a manual failover to the backup node in the subcluster using the Cluster Topology interface. When you initiate a manual failover, the Cisco UP Server Recovery Manager stops the critical services on that node, and moves all users to the backup node.

The Cisco UP Server Recovery Manager stops the following critical services on the node:

- Cisco UP SIP Proxy
- Cisco UP Presence Engine

- Cisco UP XCP Router (this causes all XCP processes to stop)
- Cisco UP Client Profile Agent

The Cisco UP Server Recovery Manager then moves all users to the backup node

Restriction

You can only initiate a failover for a node that is in 'Normal' state.

Before You Begin

Make sure that these services are running on the Failing Over node:

- Cisco UP XCP Connection Manager service
- Cisco UP XCP Router
- Cisco UP Presence Engine

Procedure

Step 1 **Cisco Unified Presence Administration > System > Cluster Topology.**

Step 2 Click the **edit** link on the appropriate subcluster.

Step 3 Choose **Failover** in the Node Action column.

Step 4 Click **Ok** to confirm the failover operation.

Step 5 To verify the failover operation is complete and successful:

- When the failover operation is in progress, the primary node should be in the "Failing Over" state, and the backup node should be in the "Taking Over" state. When the failover operation is complete, check that the backup node is in the state 'Running in Backup Mode', and the primary node is in "Idle' state. If the failover is unsuccessful, and the nodes are in a failed state, see Node States, Causes and Recommended Actions, page 21-1 for a recommended action.
- Check that the users have failed over to the backup node:
  - On the subcluster details screen, check that all users are now assigned to the backup node, and no users are assigned to the primary node.
  - On the node details screen, the 'Failed Over' column indicates the users that have failed over to the backup node.

Related Topic

High Availability Deployments, page 17-12

Perform a Manual Fallback to Primary Node

You can perform a manual fallback to the primary node in the Cluster Topology interface. When you initiate a manual fallback, the Cisco UP Server Recovery Manager restarts any critical services that are not already running on the primary node, and moves the failed over users back to the primary node.

When you manually initiate a fallback, the Cisco UP Server Recovery Manager restarts the following services on the primary node (if they are not already running):

- Cisco UP SIP Proxy

- Cisco UP Presence Engine
- Cisco UP XCP Router
- Any XCP services that were activated
- Cisco UP Client Profile Agent

The Cisco UP Server Recovery Manager then moves all failed over users back to the primary node.

Restriction

You can only initiate fallback for a node that is in 'Idle' or 'Failed Over' state.

Procedure

Step 1 **Cisco Unified Presence Administration > System > Cluster Topology.**

Step 2 Click the **edit** link on the appropriate subcluster.

Step 3 Choose **Fallback** in the Node Action column.

Step 4 Click **Ok** to confirm the fallback operation.

Step 5 To verify the fallback operation is complete and successful:

- When fallback operation is in progress, the primary node should be in the "Taking Back" state, and the backup node should be in the "Falling Back" state.When the fallback operation is complete, check that both nodes are in 'Normal' state. If the fallback is unsuccessful, and the nodes are in a failed state, see Node States, Causes and Recommended Actions, page 21-1 for a recommended action.
- Check that the users have fallen back to the primary node.
  - On the subcluster details screen, check that all users are now assigned to the primary node, and no users are assigned to the backup node.
  - On the node details screen, the 'Failed Over' column should be empty.

Related Topic

High Availability Deployments, page 17-12

Perform a Manual Recovery of a Subcluster

When you perform a manual recovery of a subcluster, Cisco Unified Presence restarts the Cisco UP Server Recovery Manager service on both nodes in the subcluster. You may experience IDS replication delays. You can check the status of the IDS replication on a node using this CLI command:

utils dbreplication runtimestate

Restriction

You can only initiate a recovery for a subcluster if both nodes are in a failed state.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Cluster Topology.**

Step 2 Click the **edit** link on the appropriate subcluster.

Step 3 Choose **Recovery** in the Node Action column.

Step 4 See Node States, Causes and Recommended Actions, page 21-1 to verify the status of the subcluster after you perform the manual recovery.

Related Topic

High Availability Deployments, page 17-12

Cluster-wide Routing Information on Cisco Unified Presence

Configure a Cluster-wide Cisco Unified Presence Address

This procedure is only applicable if you are configuring a multi-node deployment. Configure the cluster-wide Cisco Unified Presence address on the publisher node, and Cisco Unified Presence will replicate the address on all nodes in the cluster.

Note When you configure a cluster-wide Cisco Unified Presence address, set the port of SRV to 5060.

Before You Begin

Read the cluster-wide DNS SRV topic.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Service Parameters.**

Step 2 Choose the Cisco Unified Presence server from the Server menu.

Step 3 Choose **Cisco UP Sip Proxy** from the Service menu.

Step 4 Edit the **SRV Cluster Name** field in the General Proxy Parameters (Clusterwide) section.

By default this parameter is empty.

Step 5 Click **Save**.

What To Do Next

Upload the licenses on Cisco Unified Presence. For more information, see the *Upgrade Guide for Cisco Unified Presence Release 8.6*.

Static Route Configuration

If you configure a static route for SIP proxy server traffic, consider the following:

- A dynamic route represents a path through the network that is automatically calculated according to routing protocols and routing update messages.
- A static route represents a fixed path that you explicitly configure through the network.
- Static routes take precedence over dynamic routes.

This section contains the following subsections:

- Route Embed Templates, page 17-25
- Configure Route Embed Templates, page 17-26
- Configure Static Routes, page 17-26

Route Embed Templates

You must define a route embed template for any static route pattern that contains embedded wildcards. The route embed template contains information about the leading digits, the digit length, and location of the embedded wildcards. Before you define a route embed template, consider the sample templates we provide below.

When you define a route embed template, the characters that follow the '.' must match actual telephony digits in the static route. In the sample route embed templates below, we represent these characters with 'x'.

Sample Route Embed Template A

Route embed template: **74..78xxxxx***

With this template, Cisco Unified Presence will enable this set of static routes with embedded wildcards:

Destination Pattern	Next Hop Destination
74..7812345*	1.2.3.4:5060
74..7867890*	5.6.7.8.9:5060
74..7811993*	10.10.11.37:5060

With this template, Cisco Unified Presence will NOT enable these static route entries:

- 73..7812345* (The initial string is not ‘74’ as the template defines)
- 74..781* (The destination pattern digit length does not match the template)
- 74…7812345* (The number of wildcards does not match the template)

Sample Route Embed Template B

Route embed template: **471….xx***

With this template, Cisco Unified Presence will enable this set of static routes with embedded wildcards:

Destination Pattern	Next Hop Destination
471….34*	20.20.21.22
471…55*	21.21.55.79

With this template, Cisco Unified Presence will NOT enable these static route entries:

- 47…344* (The initial string is not ‘471’ as the template defines)
- 471…4* (The string length does not match template)
- 471.450* (The number of wildcards does not match template)

Configure Route Embed Templates

You can define up to five route embed templates. However, there is no limit to the number of static routes that you can define for any route embed template.

A static route that contains an embedded wildcard must match at least one of the route embed templates.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > System > Service Parameters**.

Step 2 Choose a Cisco Unified Presence server.

Step 3 Choose the Cisco UP SIP Proxy service.

Step 4 Define a route embed templates in the RouteEmbedTemplate field in the Routing Parameters (Clusterwide) section. You can define up to five route embed templates.

Step 5 Click **Save**.

What To Do Next

Configure Static Routes, page 17-26

Configure Static Routes

Procedure

Step 1 Choose **Cisco Unified Presence Administration > Routing > Static Routes**.

Step 2 Click **Add New**.

Step 3 Configure these static route settings:

Field	Description
Destination Pattern	This field specifies the pattern of the incoming number, up to a maximum of 255 characters. The SIP proxy allows only 100 static routes to have an identical route pattern. If you exceed this limit, Cisco Unified Presence logs an error. **Wildcard Usage** You can use "." as a wildcard for a single character and "*" as a wildcard for multiple characters. Cisco Unified Presence supports embedded '.' wildcard characters in static routes. However, you must define route embed templates for static routes that contain embedded wildcards. Any static route that contains an embedded wildcard must match at least one route embed template. See the route embed template topic (referenced in the Related Topics section below) for information about defining route embed templates. For phones: • A dot can exist at the end of the pattern, or embedded in a pattern. If you embed the dot in a pattern, you must create a route embed template to match the pattern. • An asterisk can only exist at the end of the pattern. For IP addresses and host names: • You can use an asterisk as part of the a host name. • The dot acts as a literal value in a host name. An escaped asterisk sequence, *, matches a literal * and can exist anywhere.
Description	Specifies the description of a particular static route, up to a maximum of 255 characters.
Next Hop	Specifies the domain name or IP address of the destination (next hop) and can be either a Fully Qualified Domain Name (FQDN) or dotted IP address. Cisco Unified Presence supports DNS SRV-based call routing. To specify DNS SRV as the next hop for a static route, set this parameter to the DNS SRV name.
Next Hop Port	Specifies the port number of the destination (next hop). The default port is 5060. Cisco Unified Presence supports DNS SRV-based call routing. To specify DNS SRV as the next hop for a static route, set the next hop port parameter to 0.

Field	Description
Route Type	Specifies the route type: User or Domain. The default value is user. For example, in the SIP URI "sip:19194762030@myhost.com" request, the user part is '19194762030', and the host part is 'myhost.com'. If you choose User as the route type, Cisco Unified Presence uses the user-part value '19194762030' for routing SIP traffic. If you choose the Domain as the route type, Cisco Unified Presence uses 'myhost.com' for routing SIP traffic.
Protocol Type	Specifies the protocol type for this route, TCP, UDP, or TLS. The default value is TCP.
Priority	Specifies the route priority level. Lower values indicate higher priority. The default value is 1. Value range: 1-65535
Weight	Specifies the route weight. Use this parameter only if two or more routes have the same priority. Higher values indicate which route has the higher priority. Value range: 1-65535 **Example:** Consider these three routes with associated priorities and weights: • 1, 20 • 1, 10 • 2, 50 In this example, the static routes are listed in the correct order. The priority route is based on the lowest value priority, that is 1. Given that two routes share the same priority, the weight parameter with the highest value decides the priority route. In this example, Cisco Unified Presence directs SIP traffic to both routes configured with a priority value of 1, and distributes the traffic according to weight; The route with a weight of 20 receives twice as much traffic as the route with a weight of 10. Note that in this example, Cisco Unified Presence will only attempt to use the route with priority 2, if it has tried both priority 1 routes and both failed.
Allow Less-Specific Route	Specifies that the route can be less specific. The default setting is On.
In Service	Specifies whether this route has been taken out of service. This parameter allows the administrator to effectively take a route out of service (versus removing it completely and re-adding it).
Block Route Check Box	Check to block the static route. The default setting is Unblocked.

Step 4 Click **Save**.

Related Topics

- Route Embed Templates, page 17-25
- Configure Route Embed Templates, page 17-26

Presence Gateway Configuration on Cisco Unified Presence

- Presence Gateway Configuration Option, page 17-29
- Configuring the Presence Gateway, page 17-29

Presence Gateway Configuration Option

You must configure Cisco Unified Communications Manager as a Presence Gateway on Cisco Unified Presence to enable the SIP connection that handles the availability information exchange between Cisco Unified Communications Manager and Cisco Unified Presence.

When configuring the Presence Gateway, specify the FQDN (Fully Qualified Domain Name) or the IP address of the associated Cisco Unified Communications Manager server. Depending on your network this value can be one of the following:

- the FQDN address of the Cisco Unified Communications Manager publisher
- a DNS SRV FQDN that resolves to the Cisco Unified Communications Manager subscriber nodes
- the IP address of the Cisco Unified Communications Manager publisher

If DNS SRV is an option in your network, configure the following:

1. Configure the Presence Gateway on the Cisco Unified Presence server with a DNS SRV FQDN of the Cisco Unified Communications Manager subscriber nodes (equally weighted). This will enable Cisco Unified Presence to share availability messages equally among all the servers used for availability information exchange.
2. On Cisco Unified Communications Manager, configure the SIP trunk for the Cisco Unified Presence server with a DNS SRV FQDN of the Cisco Unified Presence publisher and subscriber.

If DNS SRV is not an option in your network, and you are using the IP address of the associated Cisco Unified Communications Manager server, you cannot share availability messaging traffic equally across multiple subscriber nodes because the IP address points to a single subscriber node.

Related Topic

SIP Trunk Configuration on Cisco Unified Communications Manager, page 6-4

Configuring the Presence Gateway

Before You Begin

- Read the Presence Gateway configuration options topic.
- Depending on your configuration requirements, obtain the FQDN, DNS SRV FQDN, or the IP address of the associated Cisco Unified Communications Manager server.

Procedure

Step 1 Choose **Cisco Unified Presence Administration > Presence > Gateways**.

Step 2 Click **Add New**.

Step 3 Choose **CUCM** for the Presence Gateway Type.

Step 4 Enter a description of the presence gateway in the Description field.

Step 5 Specify the FQDN, DNS SRV FQDN, or the IP address of the associated Cisco Unified Communications Manager server in the Presence Gateway field.

Step 6 Click **Save**.

Related Topic

Presence Gateway Configuration Option, page 17-29