Configure Presence Redundancy Groups

Presence Redundancy Group Overview

A presence redundancy group is comprised of two IM and Presence Service nodes from the same cluster. Each node in the presence redundancy group monitors the status, or heartbeat, of the peer node. You can configure a presence redundancy group to provide both redundancy and recovery for IM and Presence Service clients and applications.

- Failover—Occurs in a presence redundancy group when one or more critical services fails on an IM and Presence Service node in the group or a node in the group fails. Clients automatically connect to the other IM and Presence Service node in that group.
- Fallback—Occurs when a fallback command is issued from the CLI or Cisco Unified Communications Manager during either of these conditions:
  - The failed IM and Presence Service node comes back into service and all critical services are running. The failed-over clients in that group reconnect with the recovered node when it becomes available.
  - The backup activated IM and Presence Service node fails due to a critical service failure, and the peer node is in the Failed Over state and supports the automatic recovery fallback.

For example, if you are using presence redundancy groups, Cisco Jabber clients will fail over to a backup IM and Presence Service node if the services or hardware fail on the local IM and Presence Service node. When the failed node comes online again, the clients automatically reconnect to the local IM and Presence Service node if you have configured automatic fallback. If you have not configured automatic fallback, you can manually initiate the fallback when the failed node comes online.

In addition to redundancy and recovery, presence redundancy groups also allow you to configure high availability for your cluster.

High Availability

The IM and Presence Service supports high availability for multiple-node deployments.

After you configure a presence redundancy group, you can enable high availability for the group. A pair of nodes is required for high availability. Each node has an independent database and set of users operating with a shared availability database that is able to support common users.

All IM and Presence Service nodes must belong to a presence redundancy group, which can consist of a single IM and Presence Service node or a pair of IM and Presence Service nodes.

You can configure high availability using two different modes:

- Balanced mode: This mode provides redundant high availability with automatic user load balancing and user failover in the event that one nodes fails because of component failure or power outage.
- Active/standby mode: The standby node automatically takes over for the active node if the active node fails. It does not provide automatic load balancing.

We recommend that you configure your IM and Presence Service deployments as high availability deployments. Although you are permitted to have both high availability and non-high availability presence redundancy groups configured in a single deployment, this configuration is not recommended.

Presence Redundancy Group Prerequisites

For deployments over the WAN, a minimum of 10 megabits per second of dedicated bandwidth is required for each IM and Presence Service cluster, with no more than an 80-millisecond round-trip latency. Any bandwidth less than this recommendation can adversely impact performance.

Presence Redundancy Group Task Flow

An IM and Presence Service node can be assigned to only one presence redundancy group. For high availability, you must assign two nodes from the same cluster to the presence redundancy group and enable high availability for the group.

Procedure

	Command or Action	Purpose
Step 1	Verify Database Replication, on page 3	Ensure that database replication is setup in the IM and Presence Service cluster.
Step 2	Verify Services, on page 3	Make sure critical services are running on the nodes that you plan to add to a presence redundancy group.
Step 3	Configure a Presence Redundancy Group, on page 4	Provide redundancy and recovery for IM and Presence Service clients and applications.
Step 4	Configure Heartbeat Interval for Failover, on page 5	Optional. Each node in the presence redundancy group monitors the status, or heartbeat, of its

	Command or Action	Purpose
		peer node. You can configure the intervals by which each node monitors its peer.
Step 5	Enable High Availability, on page 6	Optional. Follow this procedure if you did not enable high availability when you configured the presence redundancy group.
Step 6	Configure User Assignment Mode, on page 7	Configure how you want the Sync Agent to distribute users across various nodes in the IM and Presence Service cluster. This setting affects how your system handles failover and load balancing.

Verify Database Replication

Ensure that database replication is setup in the IM and Presence Service cluster before you enable high availability for a presence redundancy group.

Procedure

Step 1 Start a CLI session using one of the following methods:

- From a remote system, use SSH to connect securely to the Cisco Unified Operating System. In your SSH client, enter your **ssh *adminname@hostname*** and enter your password.
- From a direct connection to the serial port, enter your credentials at the prompt that displays automatically.

Step 2 Execute the **utils dbreplication status** command to check for errors or mismatches in the database tables.

Step 3 Execute the **utils dbreplication runtimestate** command to check if the database replication is active on the node.

The output lists all the nodes and if database replication is set up and in a good state, the **replication setup** value for each node is **2**.

If a value other than 2 is returned, you must resolve the errors before proceeding.

What to do next

Verify Services, on page 3

Verify Services

Make sure critical services are running on the nodes that you plan to add to a presence redundancy group. Critical services must be running before you turn on high availability. If critical services are not running on either node, the presence redundancy group will go into a Failed state when you turn on high availability. If critical services are not running on one node, then that node fails over to the other node when you turn on high availability.

Before you begin

Verify Database Replication, on page 3

Procedure

Step 1 From Cisco Unified IM and Presence Serviceability, choose **Tools** > **Control Center - Network Services**.

Step 2 From the **Server** list, choose the appopriate node and click **Go**.

Step 3 In the **IM and Presence Services** area, ensure that the following services are started:

- **Cisco Client Profile Agent**
- **Cisco Sync Agent**
- **Cisco XCP Router**

Step 4 From the **Related Links** drop-down list, select **Control Center - Network Services** and click **Go**.

Step 5 In the **IM and Presence Services** area, ensure that the following services are started:

- **Cisco SIP Proxy**
- **Cisco Presence Engine**

What to do next

Configure a Presence Redundancy Group, on page 4

Configure a Presence Redundancy Group

Use Cisco Unified Communications Manager to configure redundancy for IM and Presence Service nodes.

Each presence redundancy group can contain two IM and Presence Service nodes. Each node can be assigned to only one presence redundancy group. Both nodes in the presence redundancy group must be on the same cluster and have the same IM and Presence Service database publisher node.

Before you begin

- Verify Services, on page 3
- Ensure that the IM and Presence Service nodes you are adding to a presence redundancy group are running the same software version.

Procedure

Step 1 From **Cisco Unified CM Administration**, choose **System** > **Presence Redundancy Groups**.

Step 2 Click **Add New**.

Step 3 Enter a unique name for the presence redundancy group.

You can enter a maximum of 128 alphanumeric characters, including underscore (_) and dash (-).

Step 4 Enter a description of the group.

You can enter a maximum of 128 alphanumeric characters including symbols, but it cannot include double-quotes ("), percentage sign (%), ampersand (&), forward slash (\), or angle brackets (<>).

Step 5 Choose two different IM and Presence Service nodes in the **Presence Server** fields to assign them to the group.

Step 6 (Optional) Check the **Enable High Availability** check box to enable high availability for the presence redundancy group.

Step 7 Click **Save**.

What to do next

Configure Heartbeat Interval for Failover, on page 5

Configure Heartbeat Interval for Failover

Configure optional service parameters that determine the keep alive settings by which each peer in a presence redundancy group monitors the heartbeat (i.e., the status) of its peer node in order to confirm that the peer is active. A failover can be initiated if the peer node is unresponsive after a configured timer expires.

Note Cisco recommends that you use the default values for these service parameters. However, you can also reconfigure the values to suit your needs.

Procedure

Step 1 In Cisco Unified CM IM and Presence Administration, choose **System** > **Service Parameters**.

Step 2 From the **Server** drop-down, select an IM and Presence node

Step 3 From the **Service** drop-down, select **Cisco Server Recovery Manager (Active)**.

Step 4 Under **General Server Recovery Manager Parameters (Clusterwide)**, configure the clusterwide Keep Alive settings that each node in a Presence Redundancy Group uses to monitor monitor the heartbeat of its peer node. A failover can be initiated if the peer node is unresponsive.

- **Service Port**— This parameter specifies the port that Cisco Server Recovery Manager uses to communicate with its peer. The default is 22001.
- **Admin RPC Port**—This parameter specifies the port that Cisco Server Recovery Manager uses to provide admin rpc requests. The default is 20075.
- **Critical Service Delay**—This parameter specifies the duration in seconds that a critical service can be down before failover is initiated. The default is 90.
- **Enable Automatic Fallback**—This parameter specifies whether to do automatic fallback. In the event of a failover, the IM and Presence Service moves users automatically from the backup node to the primary node thirty minutes after the primary node returns to a healthy state. The default value is False.
- **Initialization Keep Alive (Heartbeat) Timeout**—This parameter specifies the duration in seconds that the heartbeat can be lost with the peer during initialization before failover is initiated. The default is 120.

- **Keep Alive (Heartbeat) Timeout**—This parameter specifies the duration in seconds that the heartbeat can be lost with the peer before failover is initiated. the default is 60.
- **Keep Alive (HeartBeat) Interval**—This parameter specifies the interval in seconds between keep alive (heart beat) messages being sent to the peer. The default is 15.

Step 5 Configure the following additional parameters, which tell CUPC 8.5 and higher clients how long to wait before attempting to relogin. Unlike the above parameters, these parameters must be configured separately for each cluster node.

- **Client Re-Login Lower Limit**—This parameter specifies the minimum number of seconds which CUPC 8.5 (and higher) should wait before attempting to re-login to this server. The default is 120.
- **Client Re-Login Upper Limit**—This parameter specifies the maximum number of seconds which CUPC 8.5 (and higher) should wait before attempting to re-login to this server. The default is 537.

Step 6 Click **Save**.

What to do next

If you did not enable high availability when you configured the presence redundancy group, Enable High Availability, on page 6 now.

Enable High Availability

Caution Failure to set up replication in the IM and Presence Service cluster and ensure that all critical services are running may result in an immediate failover when high availability is enabled for the presence redundancy group.

Before you begin

- Configure a Presence Redundancy Group, on page 4
- Ensure that replication is set up in the IM and Presence Service cluster.
- Ensure that all critical services are running.

Procedure

Step 1 From **Cisco Unified CM Administration**, choose **System** > **Presence Redundancy Groups**.

Step 2 Specify search criteria and then click **Find**.

Step 3 Choose the presence redundancy group that you configured.

Step 4 To enable high availability, check the **Enable High Availability** check box.

Step 5 Click **Save**.

Configure User Assignment Mode

Use this procedure to configure the way in which the sync agent distributes users to the nodes in the cluster. This setting helps to manage failover and load balancing.

Procedure

Step 1 From Cisco Unified CM Administration, choose **System** > **Enterprise Parameters**.

Step 2 In the **User Management Parameters** Area, choose one of the following options for the **User Assignment Mode for Presence Server** parameter:

- **Balanced**—This mode assigns users equally to each node in each subcluster and attempts to balance the total number of users equally across each node. This is the default option.
- **Active-Standby**—This mode assigns all users to the first node of the subcluster, leaving the secondary server as a backup.
- **None**—This mode results in no assignment of the users to the nodes in the cluster by the sync agent.

Step 3 Click **Save**.

Redundancy Interactions and Restrictions

Feature	Interaction
Adding Users	You cannot add new users to an IM and Presence Service cluster while one of the cluster nodes is in a failover state.
Multiple Device Messaging	The Multiple Device Messaging feature causes a delay with server recovery on the IM and Presence Service if failover occurs. If server failover occurs on a system where Multiple Device Messaging is configured, the failover times generally are twice as long as the times specified with the **Cisco Server Recovery** Manager service parameters.

Feature	Interaction
Temporary presence status of a user	The temporary presence status of a user displays the stale presence status after Failover, Fallback, and user moves. This is because the subscription to temporary presence will be deleted and the user must re-subscribe to temporary presence to see the valid temporary presence status of the user. For example, If User A is subscribed to user B's temporary presence and a failover occurs on the IM and Presence node where User B is assigned, then user B displays offline to User A even after User B re-logins to the backup node. It is because the subscription to temporary presence of User B is deleted and User A is not aware of the deletion. User A must re-subscribe to temporary presence of User B again. When User A deletes search of User B from Jabber client, User A needs to wait at least 30 seconds before It tries to search the temporary presence of User B. If not, then User A sees the stale presence of User B. Jabber client must wait for at least 30 seconds between two searches for same user to get a valid temporary presence status.
IM and Presence status	When a user is moved from one Presence Redundancy Group to another, The user has to be logged out from Jabber session, for the IM and Presence status to be visible in the current Presence Redundancy Group which the user has moved into.

Initiate Manual Failover, Fallback, or Recovery

Use this procedure to initiate manual failover, fallback, or recovery of IM and Presence Service nodes within a presence redundancy group.

- Manual failover—When you initiate a manual failover, the **Cisco Server Recovery Manager** stops the critical services on the failed node. All users from the failed node are disconnected and must re-login to the backup node. Critical services will not be restarted unless we invoke manual fallback.
- Manual fallback—When you initiate a manual fallback, the **Cisco Server Recovery Manager** restarts critical services on the primary node and disconnects all users that had been failed over. Those users must then re-login to their assigned node.
- Manual recovery—A manual recovery is necessary when both nodes in the presence redundancy group are in the failed state. In this case, the IM and Presence Service restarts the **Cisco Server Recovery Manager** service on both nodes in the presence redundancy group.

Procedure

Step 1 From Cisco Unified CM Administration, choose **System** > **Presence Redundancy Groups**.

Step 2 Click **Find** and select the Presence Redundancy Group with the applicable nodes.

Step 3 Do one of the following. Note that the available button depends on the current state of the node:

- Click **Failover** to initiate failover of an active node.

- Click **Fallback** to initate fallback of a failed over node.
- Click **Recover** if both nodes are failed over and you want to recover them.

Note You can also initiate these actions from Cisco Unified Communications Manager or IM and Presence Service using the CLI. See the *Command Line Interface Guide for Cisco Unified Communications Solutions* for details.

Note You cannot add end users to an IM and Presence Service cluster while one of the nodes is in a failover state.

Node State Definitions

Table 1: Presence Redundancy Group Node State Definitions

State	Description
Initializing	This is the initial (transition) state when the Cisco Server Recovery Manager service starts; it is a temporary state.
Idle	IM and Presence Service is in Idle state when failover occurs and services are stopped. In Idle state, the IM and Presence Service node does not provide any availability or Instant Messaging services. In Idle state, you can manually initiate a fallback to this node using the **Cisco Unified CM Administration** user interface.
Normal	This is a stable state. The IM and Presence Service node is operating normally. In this state, you can manually initiate a failover to this node using the **Cisco Unified CM Administration** user interface.
Running in Backup Mode	This is a stable state. The IM and Presence Service 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 IM and Presence Service node is taking over for its peer node.
Failing Over	This is a transition state. The IM and Presence Service node is being taken over by its peer node.
Failed Over	This is a steady state. The IM and Presence Service node has failed over, but no critical services are down. In this state, you can manually initiate a fallback to this node using the **Cisco Unified CM Administration** user interface.
Failed Over with Critical Services Not Running	This is a steady state. Some of the critical services on the IM and Presence Service node have either stopped or failed.
Falling Back	This is a transition state. The system is falling back to this IM and Presence Service node from the node that is running in backup mode.

State	Description
Taking Back	This is a transition state. The failed IM and Presence Service 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	Node state is unknown. A possible cause is that high availability was not enabled properly on the IM and Presence Service node. Restart the Server Recovery Manager service on both nodes in the presence redundancy group.

Node States, Causes, and Recommended Actions

You can view the status of nodes in a presence redundancy group on the **Presence Redundancy Group Configuration** window when you choose a group using the **Cisco Unified CM Administration** user interface.

Table 2: Presence Redundancy Group Node High-Availability States, Causes, and Recommended Actions

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Normal	Normal	Normal	Normal	Normal
Failing Over	On Admin Request	Taking Over	On Admin Request	The administrator initiated a manual failover from node 1 to node 2. The manual failover is in progress.
Idle	On Admin Request	Running in Backup Mode	On Admin Request	The manual failover from node 1 to node 2 that the administrator initiated is complete.
Taking Back	On Admin Request	Falling Back	On Admin Request	The administrator initiated a manual fallback from node 2 to node 1. The manual fallback is in progress.
Idle	Initialization	Running in Backup Mode	On Admin Request	The administrator restarts the SRM service on node 1 while node 1 is in “Idle” state.
Idle	Initialization	Running in Backup Mode	Initialization	The administrator either restarts both nodes in the presence redundancy group, or restarts the SRM service on both nodes while the presence redundancy group was in manual failover mode.
Idle	On Admin Request	Running in Backup Mode	Initialization	The administrator restarts the SRM service on node 2 while node 2 is running in backup mode, but before the heartbeat on node 1 times out.
Failing Over	On Admin Request	Taking Over	Initialization	The administrator restarts the SRM service on node 2 while node 2 is taking over, but before the heartbeat on node1 times out.

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Taking Back	Initialization	Falling Back	On Admin Request	The administrator restarts the SRM service on node 1 while taking back, but before the heartbeat on node 2 times out. After the taking back process is complete, both nodes are in Normal state.
Taking Back	Automatic Fallback	Falling Back	Automatic Fallback	Automatic Fallback has been initiated from node 2 to node 1 and is currently in progress.
Failed Over	Initialization or Critical Services Down	Running in Backup Mode	Critical Service Down	Node 1 transitions to Failed Over state when either of the following conditions occur: • Critical services come back up due to a reboot of node 1. • The administrator starts critical services on node 1 while node 1 is in Failed Over with Critical Services Not Running state. When node 1 transitions to Failed Over state the node is ready for the administrator to perform a manual fallback to restore the nodes in the presence redundancy group to Normal state.
Failed Over with Critical Services not Running	Critical Service Down	Running in Backup Mode	Critical Service Down	A critical service is down on node 1. IM and Presence Service performs an automatic failover to node 2. **Recommended Actions:** **1.** Check node 1 for any critical services that are down and try to manually start those services. **2.** If the critical services on node 1 do not start, then reboot node 1. **3.** When all the critical services are up and running after the reboot, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.
Failed Over with Critical Services not Running	Database Failure	Running in Backup Mode	Database Failure	A database service is down on node 1. IM and Presence Service performs an automatic failover to node 2. **Recommended Actions:** **1.** Reboot node 1. **2.** When all the critical services are up and running after the reboot, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Running in Failed Mode	Start of Critical Services Failed	Running in Failed Mode	Start of Critical Services Failed	Critical services fail to start while a node in the presence redundancy group is taking back from the other node. **Recommended Actions.** On the node that is taking back, perform the following actions: **1.** Check the node for critical services that are down. To manually start these services, click **Recovery** in the **Presence Redundancy Group Configuration** window. **2.** If the critical services do not start, reboot the node. **3.** When all the critical services are up and running after the reboot, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.
Running in Failed Mode	Critical Service Down	Running in Failed Mode	Critical Service Down	Critical services go down on the backup node. Both nodes enter the failed state. **Recommended Actions:** **1.** Check the backup node for critical services that are down. To start these services manually, click **Recovery** in the **Presence Redundancy Group Configuration** window. **2.** If the critical services do not start, reboot the node.

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Node 1 is down due to loss of network connectivity or the SRM service is not running.		Running in Backup Mode	Peer Down	Node 2 has lost the heartbeat from node 1. IM and Presence Service performs an automatic failover to node 2. **Recommended Action.** If node 1 is up, perform the following actions: **1.** Check and repair the network connectivity between nodes in the presence redundancy group. When you reestablish the network connection between the nodes, the node may go into a failed state. Click **Recovery** in the **Presence Redundancy Group Configuration** window to restore the nodes to the Normal state. **2.** Start the SRM service and perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state. **3.** (If the node is down) Repair and power up node 1. **4.** When the node is up and all critical services are running, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.
Node 1 is down (due to possible power down, hardware failure, shutdown, reboot)		Running in Backup Mode	Peer Reboot	IM and Presence Service performs an automatic failover to node 2 due to the following possible conditions on node 1: • hardware failure • power down • restart • shutdown **Recommended Actions:** **1.** Repair and power up node 1. **2.** When the node is up and all critical services are running, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Failed Over with Critical Services not Running OR Failed Over	Initialization	Backup Mode	Peer Down During Initialization	Node 2 does not see node 1 during startup. **Recommended Action:** When node1 is up and all critical services are running, perform a manual fallback to restore the nodes in the presence redundancy group to the Normal state.
Running in Failed Mode	Cisco Server Recovery Manager Take Over Users Failed	Running in Failed Mode	Cisco Server Recovery Manager Take Over Users Failed	User move fails during the taking over process. **Recommended Action:** Possible database error. Click **Recovery** in the **Presence Redundancy Group Configuration** window. If the problem persists, then reboot the nodes.
Running in Failed Mode	Cisco Server Recovery Manager Take Back Users Failed	Running in Failed Mode	Cisco Server Recovery Manager Take Back Users Failed	User move fails during falling back process. **Recommended Action:** Possible database error. Click **Recovery** in the **Presence Redundancy Group Configuration** window. If the problem persists, then reboot the nodes.
Running in Failed Mode	Unknown	Running in Failed Mode	Unknown	The SRM on a node restarts while the SRM on the other node is in a failed state, or an internal system error occurs. **Recommended Action:** Click **Recovery** in the **Presence Redundancy Group Configuration** window. If the problem persists, then reboot the nodes.
Backup Activated	Auto Recover Database Failure	Failover Affected Services	Auto Recovery Database Failure.	The database goes down on the backup node. The peer node is in failover mode and can take over for all users in the presence redundancy group. Auto-recovery operation automatically occurs and all users are moved over to the primary node.
Backup Activated	Auto Recover Database Failure	Failover Affected Services	Auto Recover Critical Service Down	A critical service goes down on the backup node. The peer node is in failover mode and can take over for all users in the presence redundancy group. Auto-recovery operation automatically occurs and all users are moved over to the peer node.

Node 1		Node 2		
State	**Reason**	**State**	**Reason**	**Cause/Recommended Actions**
Unknown		Unknown		Node state is unknown. A possible cause is that high availability was not enabled properly on the IM and Presence Service node. **Recommended Action:** Restart the Server Recovery Manager service on both nodes in the presence redundancy group.