

Main page: [Cisco Unified Presence, Release 7.x](#)

Contents

- [1 Previous Topic](#)
- [2 Scalability Enhancements](#)
- [3 Supported Endpoints](#)
- [4 Performance Recommendations](#)
 - ◆ [4.1 Related Topics](#)
- [5 High-Availability](#)
 - ◆ [5.1 Related Topics](#)
- [6 Subclusters](#)
 - ◆ [6.1 Related Topics](#)
- [7 User Assignment](#)
 - ◆ [7.1 Related Topics](#)
- [8 Server-Side Failback for Cisco Unified Personal Communicator Clients](#)

Previous Topic

- [Planning a Cisco Unified Presence Multi-Node Deployment](#)

- [Scalability Enhancements](#)
- [Supported Endpoints](#)
- [Performance Recommendations](#)
- [High-Availability](#)
- [Subclusters](#)
- [User Assignment](#)
- [Server-Side Failback for Cisco Unified Personal Communicator Clients](#)

Scalability Enhancements

The Cisco Unified Presence multi-node scalability feature supports the following:

- A Cisco Unified Presence cluster containing up to six nodes:
 - ◆ These six nodes can support thirty thousand users, where each user has an average of one hundred contacts.
 - ◆ This allows Cisco Unified Presence deployments to scale along with the thirty-thousand user deployment level of Cisco Unified Communications Manager.
- Five thousand users per node with a Cisco Unified IP Phone and a presence client (two end points per user, or 10,000 end points per node).
- A user can have a maximum contact list size of two hundred.
- Cisco Unified Presence continues to support intercluster deployments with the multi-node feature.

Supported Endpoints

The multi-node scalability feature supports the following end points:

- Cisco Unified Communications Manager (desk phone)
- Cisco Unified Personal Communicator (soft client)
- Cisco Unified Mobile Communicator
- Cisco Unified Mobile Communicator (soft client)
- Microsoft Office Communicator (Microsoft soft client)
- Lotus Sametime (Lotus soft client)
- Third Party Interface clients

Performance Recommendations

You can achieve optimum performance with the multi-node feature when:

- The resources on all Cisco Unified Presence servers are equivalent in terms of memory, disk size, and age. Mixing hardware classes results in servers that are under-powered, therefore resulting in poor performance.
- You deploy hardware that complies with the hardware recommendations.
- You configure a Balanced Mode deployment model. In this case the total number of users is equally divided across all nodes in all subclusters. Cisco Unified Presence defaults to Balanced Mode user assignment to achieve optimum performance.

Related Topics

- [Multi-Node Hardware Recommendations](#)
- [Balanced Non-Redundant High-Availability Deployment](#)
- [Balanced Redundant High-Availability Deployment](#)

High-Availability

We recommend that you configure your Cisco Unified Presence deployments as high-availability deployments. Although mixed mode deployments are permitted, for example high-availability subclusters and non high-availability subclusters in a single deployment, we do not recommend this configuration.

You can achieve a high-availability deployment by configuring the Balanced Mode (Redundant High-Availability) or the Active/Standby Redundant High-Availability deployment models.

Related Topics

- [Subclusters](#)
- [Balanced Redundant High-Availability Deployment](#)

- [Active/Standby Redundant High-Availability Deployment](#)

Subclusters

The multi-node feature introduces the concept of a *subcluster*. A subcluster is a single Cisco Unified Presence server, or a pair of Cisco Unified Presence servers operating with a shared presence database that is able to support common users.

In a single-node deployment within the subcluster, there is no high-availability failover protection for the users assigned to that node. In a dual-node deployment within a subcluster, users have high-availability failover protection as each node acts as a backup for the other node allowing clients to fail over in case of outages of components or nodes. In this high-availability mode, all users in the subcluster have redundancy and full failover capabilities.

Related Topics

- [About the Multi-Node Deployment Models](#)
- [Subcluster-wide DNS SRV](#)
- [Configuring Subclusters](#)

User Assignment

You can manually or automatically assign users in a Cisco Unified Presence deployment. You can manage user assignment using the User Assignment Mode parameter on the Sync Agent on Cisco Unified Presence.

If you configure automatic user assignment, the Sync Agent assigns the users to all nodes in all subclusters in an attempt to balance the user assignment evenly across all nodes. You can also configure the Sync Agent to assign the total number of users to only the first (active) node of an subcluster, with a standby node being used for failover protection.

If you configure manual user assignment, you must manually assign your users to nodes, and subclusters, using the system topology interface on Cisco Unified Presence Administration.

Related Topics

- [About the Multi-Node Deployment Models](#)
- [User Assignment Mode Recommendations](#)
- [How to Work with User Assignments in System Topology](#)

Server-Side Failback for Cisco Unified Personal Communicator 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

Cisco Unified Presence, Release 7.x -- About the Multi-Node Scalability Feature

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.