Main page: Cisco Unified Presence, Release 7.x

Contents

- 1 Previous Topic
- 2 Balanced Non-Redundant High-Availability Deployment
  - 2.1 Figure: Balanced Non- Redundant High-Availability Deployment
  - 2.2 Related Topics
- 3 Balanced Redundant High-Availability Deployment
  - 3.1 Related Topics
- 4 Active/Standby Redundant High-Availability Deployment
  - 4.1 Figure: Active/Standby High Availability Deployment
  - 4.2 Related Topics
- 5 Non High-Availability Deployment
  - 5.1 Figure: Non High-Availability Deployment

Previous Topic

- Planning a Cisco Unified Presence Multi-Node Deployment

You need to consider how you are going to deploy the multi-node feature in your network. The Cisco Unified Presence system topology interface allows you to configure the desired multi-node deployment model for your network. You access the system topology interface in Cisco Unified Presence Administration by selecting **System > Topology**.

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

Note: The system topology interface is only used to configure your *local* Cisco Unified Presence cluster. See the intercluster peer module for information on configuring intercluster peer relationships with remote Cisco Unified Presence clusters.

- Balanced Non-Redundant High-Availability Deployment
- Balanced Redundant High-Availability Deployment
- Active/Standby Redundant High-Availability Deployment
- Non High-Availability Deployment

Balanced Non-Redundant High-Availability Deployment

You can achieve a balanced mode high-availability deployment by evenly balancing your users across all nodes in the deployment, using up to 70% of the CPU of each Cisco Unified Presence server.

The balanced mode high-availability deployment option in a non-redundant manner supports up to thirty

thousand users per cluster. For this deployment model, distribute the users evenly to each of the Cisco Unified Presence nodes. For example, if you have six Cisco Unified Presence nodes in your deployment, and thirty thousand users, you assign five thousand users to each Cisco Unified Presence node.

When you use the balanced mode high-availability deployment option in a non-redundant manner, you can assign twice the number of users compared to the balanced redundant high-availability deployment option, but you use the Cisco Unified Presence hardware to maximum capacity.

If one node fails, the other node *may* handle the full load of the additional 50% of users in the subcluster, but may not be able to sustain this full load at peak traffic. This mode can provide failover protection in low traffic periods, but not during peak traffic periods.

See Figure: Balanced Non- Redundant High-Availability Deployment for an example of this deployment model on the system topology interface.

Figure: Balanced Non- Redundant High-Availability Deployment

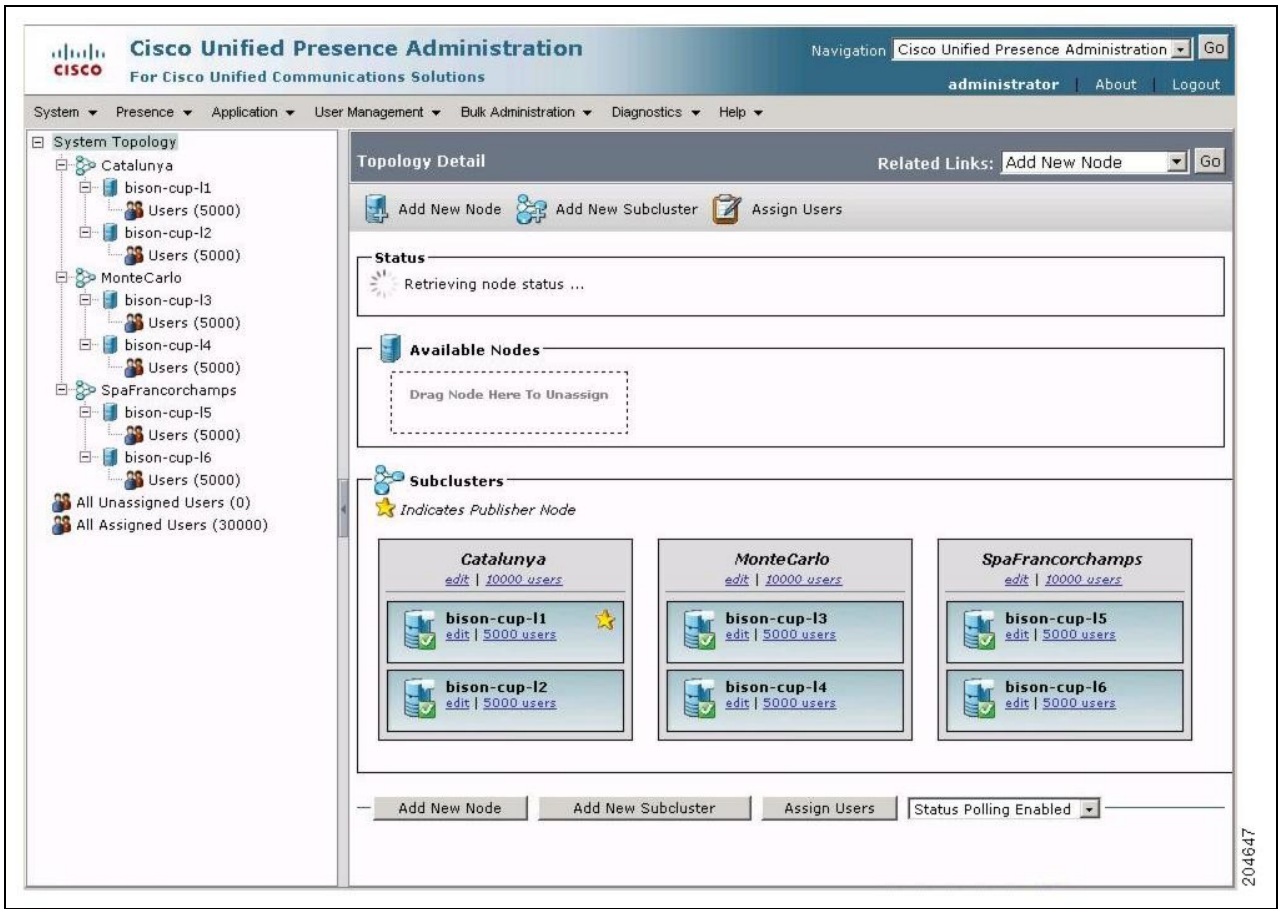

Related Topics

- How to Configure the System Topology on Cisco Unified Presence
- User Assignment Mode Recommendations
- Scalability Options for Your Deployment

- 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

Balanced 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.

Related Topics

- How to Configure the System Topology on Cisco Unified Presence
- User Assignment Mode Recommendations
- Scalability Options for Your Deployment
- 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 Redundant High-Availability Deployment

For this deployment model, assign all your users to the active Cisco Unified Presence node, and none to the backup node. The standby node can handle all traffic from the active node if the active node fails.

See Figure: Active/Standby High Availability Deployment for an example configuration for this deployment model.

Figure: Active/Standby High Availability Deployment

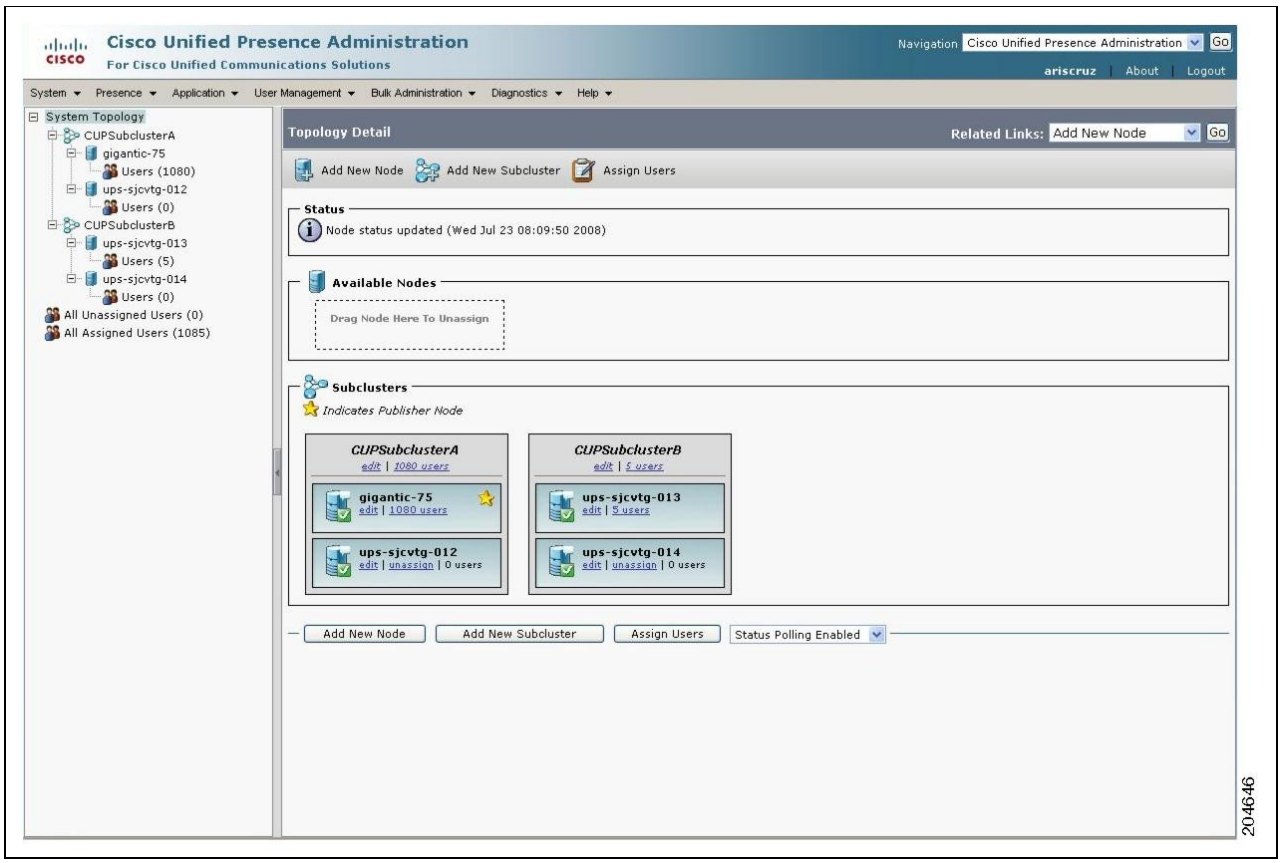

Related Topics

- Subcluster-wide DNS SRV
- How to Configure the System Topology on Cisco Unified Presence
- User Assignment Mode Recommendations
- Scalability Options for Your Deployment
- 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

Non High-Availability Deployment

For this deployment, assign all your users to one node. With this deployment users have no failover protection if services or hardware failure on this node. In order to achieve failover protection, you must configure a redundant node in the subcluster.

Note: We do not recommend this deployment model, but it is permitted.

See Figure: Non High-Availability Deployment for an example configuration for this deployment model.

Figure: Active/Standby High Availability Deployment

Figure: Non High-Availability Deployment