



This PDF is generated from authoritative online content, and is provided for convenience only. This PDF cannot be used for legal purposes. For authoritative understanding of what is and is not supported, always use the online content. To copy code samples, always use the online content.

Microsoft Skype for Business Deployment Guide

High-Availability Deployment

5/14/2025

High-Availability Deployment

Contents

- **1 High-Availability Deployment**
 - 1.1 Skype for Business High Availability
 - 1.2 T-Server High Availability
 - 1.3 UCMA Connector High-Availability

This section describes the general steps for setting up a high-availability (HA) environment for T-Server and UCMA Connector for Skype for Business.

Skype for Business High Availability

The main high-availability scheme for most server roles in Skype for Business is based on server redundancy via pooling. If a server running a certain server role fails, the other servers in the pool running the same role take the load of that server. This applies to Front End Servers, Edge Servers, Mediation Servers, and Directors. Skype for Business also enhances Back End Server high availability, by supporting synchronous SQL mirroring for your Back End databases.

For information about high availability, see Microsoft documentation:

- [Microsoft Skype for Business Server](#)

Note: Skype for Business high-availability features are available for the Enterprise edition only.

Deployment

For information about deploying high availability, see Microsoft documentation:

- [Microsoft Skype for Business Server](#)

T-Server High Availability

The high-availability architecture of T-Server for Skype for Business implies the existence of redundant applications, a primary and a backup. These are monitored by a management application so that, if one application fails, the other can take over its operations without any significant loss of contact center data. The Framework Management Layer currently supports two types of redundant configuration: warm standby and hot standby. Microsoft Skype for Business Server supports only hot standby.

Licensing HA Implementations

T-Servers operating with the hot standby redundancy type require a special CTI HA technical license, which allows for high-availability implementations, in addition to regular T-Server licenses. Neither T-Server in a redundant pair configured for hot standby starts if this license is unavailable. Moreover, the primary and backup T-Servers must use the same licenses to control the same pool of DN's. If your T-Servers are configured with the hot standby redundancy type, order licenses for CTI HA support.

Synchronization Between Redundant T-Servers

When T-Servers operate in a high-availability environment, the backup T-Server must be ready to take on the primary role when required. For this purpose, both T-Servers must be running and must

have the same information. When you configure redundant T-Servers to operate with the hot standby type, the primary T-Server uses the connection to the backup to deliver synchronization updates. Genesys recommends that you enable the Advanced Disconnect Detection Protocol (ADDP), for this connection. Do so using the configuration options in the “Backup-Synchronization Section” section. Refer to the [T-Server Common Configuration Options](#) chapter for option descriptions.

Configuration Warnings

When configuring T-Servers to support hot standby redundancy type, remember:

- When at least one of the two T-Servers that operate in a redundant mode is running, do not change a redundancy type, host, or port in either T-Server configuration.
- When both the primary and backup T-Servers are running, do not remove the backup T-Server Application object from the configuration.
- You are responsible for the option synchronization in the configuration of the primary and backup T-Servers; Configuration Server does not synchronize either options or their values in different T-Server Application objects. That is, you must configure both T-Servers to have the same options with the same values. If you change a value in one T-Server configuration, you must change it in the other T-Server configuration manually. The log options in the primary T-Server can differ from those in the backup T-Server configuration. The link configuration options in the primary T-Server can also differ from those in the backup T-Server configuration.

Important

Genesys strongly recommends that you install the backup and primary T-Servers on different host computers.

Hot Standby Redundancy

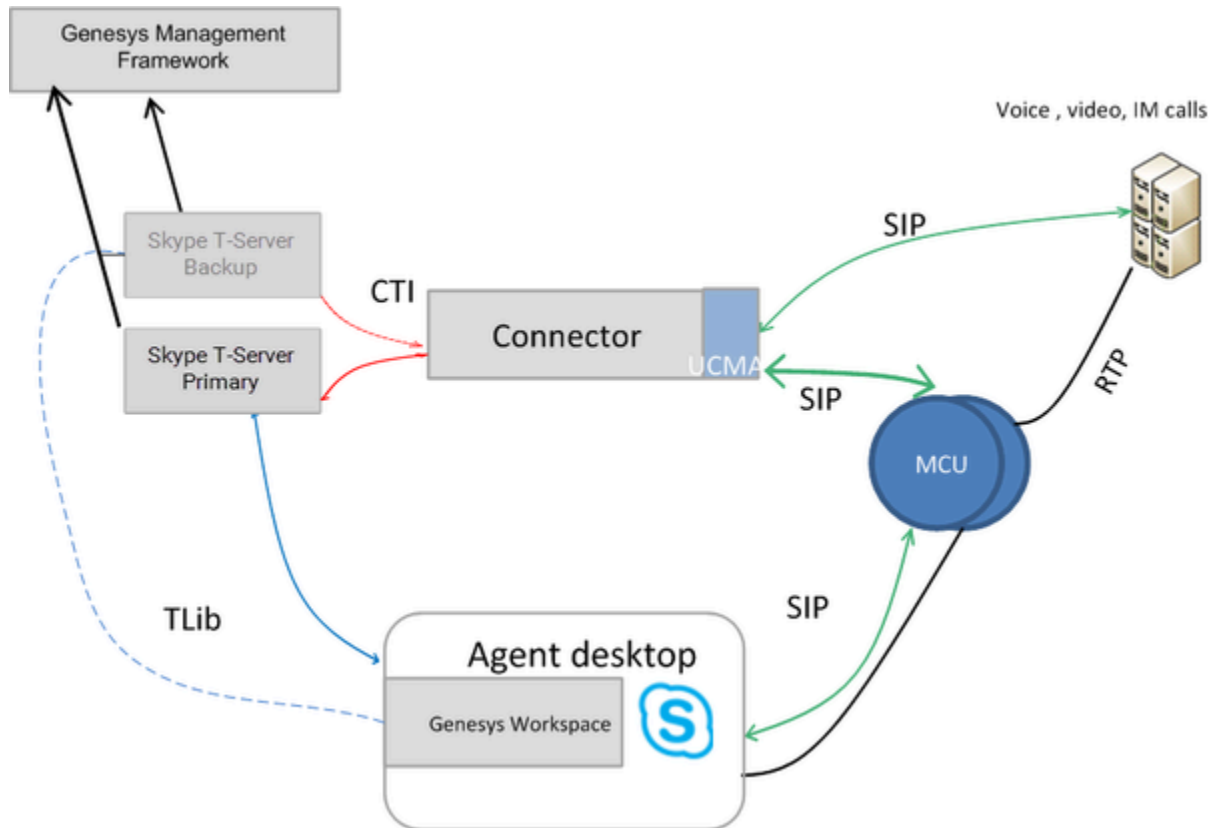
Genesys uses the expression *hot standby* to describe the redundancy type in which a backup server application remains initialized, clients connect to both the primary and backup servers at startup, and the backup server data is synchronized from the primary server. Data synchronization and existing client connections to the backup guarantee higher availability of a component.

T-Servers start simultaneously and connect to the switch via Connector component. At T-Server startup, the Management Layer assigns the role of the primary server to T-Server 1, and the role of backup to T-Server 2. T-Server clients register with both T-Servers, but only the primary T-Server handles client requests other than the registration requests. The internal T-Server information is synchronized between the primary and backup T-Servers, such as:

- Calls (all necessary information including UCMA data)
- Device info
- Monitoring subscriptions
- Agent states
- Remote Treatment sessions

Therefore, the backup T-Server has the same information as the primary T-Server. If T-Server 1 fails,

the Management Layer makes T-Server 2 the new primary server, and it starts processing client requests. The Management Layer attempts to restart T-Server 1, and if it is successful, it makes T-Server 1 the new backup server.



Benefits of Hot Standby Redundancy

The hot standby redundancy type provides the following benefits:

- Ensures the processing of interactions in progress if a failure occurs. After the primary T-Server (T-Server 1) fails, T-Server 2 handles all new interactions and takes over the processing of interactions that are currently in progress.
- T-Servers perform one-way (from primary to backup) synchronization of call-associated data, including, but not limited to:
 - Connection IDs.
 - Attached user data.
 - Inter Server Call Control (ISCC; formerly called External Routing) call references to another site in a multi-site environment (to support the ISCC/COF feature).
 - Allocation of ISCC-controlled resources.

Deployment

The general guidelines for T-Server hot standby configuration are:

1. Configure two T-Server Application objects as described in [Configuring T-Server](#).
2. Make sure the Switch object is configured for the switch these T-Servers should serve, as described in [Switches](#).
3. Modify the configuration of the primary and backup T-Servers as instructed below.

Modifying the primary T-Server configuration for hot standby

1. Stop both primary and backup T-Servers if they are already running.
2. Using Genesys Administrator or Configuration Manager, open the Properties dialog box of the Application object for the T-Server that you want to configure as a primary server.
3. On the **Switches** tab, ensure that it specifies the Switch that this T-Server Application should serve. If necessary, select the correct Switch using the Browse button.
4. In the **Server Info** tab:
 - In the Ports section, select the port to which the backup server will connect for HA data synchronization and click **Edit Port**. In the Port Properties dialog box, on the **Port Info** tab, select the **HA sync** check box.
 - Specify the T-Server Application you want to use as the backup server. Use the Browse button next to the Backup Server field to locate the backup T-Server Application object.
 - Select **Hot Standby** as the Redundancy Type.
5. On the **Start Info** tab, select **Auto-Restart**.
6. To enable ADDP between the primary and backup T-Servers, click the **Options** tab. Open or create the **[backup-sync]** section and configure corresponding options.

Modifying the backup T-Server configuration for hot standby

1. Ensure the two T-Servers are not running.
2. Using Genesys Administrator or Configuration Manager, open the Properties dialog box of the Application object for the T-Server that you want to configure as a backup server.
3. On the **Switches** tab, using the Browse button, select the same Switch object you associated with the primary T-Server Application.
4. On the **Server Info** tab:
 - In the Ports section, select the port to which the primary server will connect for HA data synchronization and click **Edit Port**. In the Port Properties dialog box, on the **Port Info** tab, select the **HA sync** check box.
5. On the **Start Info** tab, select **Auto-Restart**.
6. On the **Options** tab, modify the values for all necessary configuration options. Genesys recommends that you set all configuration options for the backup T-Server to the same values as for the primary T-Server; the only exceptions are the log options and the **server-id** option.

Known T-Server HA Limitations

The following limitations apply to T-Server HA:

- Client requests sent during the failure and switchover might be lost.
- Treatment sessions are not synchronized between the primary and backup T-Servers and are interrupted if the primary T-Server goes down or is switched over.
- Routing requests sent by the switch during the failure and switchover might be lost.
- T-Server does not synchronize interactions that begin before it starts.
- Some T-Library events might be duplicated or lost.
- Reference IDs from client requests can be lost in events.
- Only *hot standby* redundancy type is supported.
- During a T-Server switchover, an unfinished single-step transfer of an instant messaging or audio/video call can lead to lost information about a transfer controller in T-Library messaging. Queries performed for this call might return all parties participating in the single-step transfer as being in a conference call.

UCMA Connector High-Availability

The high-availability architecture for UCMA Connector for Skype for Business implies:

- UCMA HA data storage
- Synchronization of HA data between UCMA Connector and T-Server
- Existence of a pool of UCMA Connectors

UCMA HA Data Storage

To prevent loss of calls if the server disconnects, the Connector stores all necessary HA call data, such as:

- Conferences
- Conversations
- Calls and parties
- B2B calls

This information allows UCMA Connector to recover calls after the restoration of communication with Skype for Business.

Synchronization of HA data between UCMA Connector and T-Server

To provide high-availability of existing calls, UCMA Connector synchronizes all necessary HA call data with T-Server. Data synchronization comprises information about:

- Conferences info
- Conversations info
- Calls and parties info
- B2B calls info

This information allows T-Server and UCMA Connector for Skype for Business to recover calls after Connectors restart.

Pools of UCMA Connectors

A pool of UCMA Connectors allows the T-Server and UCMA Connector environment to proceed working normally, even if some Connectors are down. In this case, T-Servers redistribute DNs from non-working Connectors to remaining Connectors. Skype for Business distributes all incoming calls to remaining Connectors. After recovery, each Connector will try to restore calls according to HA information received from T-Server.

Note: Pooling of UCMA Connectors is available only with Skype for Business Server 2015.

Deployment

Storage of UCMA HA data and synchronization of HA data between UCMA Connector and T-Server work automatically and do not require additional configuration.

A pool of UCMA Connectors can be configured as described in [Provisioning for UCMA Connectors](#).

Known Connector HA Limitations

- Instant Message calls cannot be recovered after the Connector restarts.
- Calls can only be recovered on the Connector where they were originally created, after the Connector restarts.