



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.

SIP Endpoint SDK Overview

SIP Endpoint SDK 8.5.20SX

2/11/2022

Table of Contents

Overview	3
SIP Endpoint SDK for OS X Components	6
SIP-Based Third-Party Call Control	8
T-Lib-Initiated MakeCall re-INVITE	9
T-Lib-Initiated MakeCall INVITE	11
T-Lib Termination TReleaseCall	12
T-Lib-Initiated Hold-Retrieve	15
T-Lib-Initiated Hold-Retrieve Music on Hold	17
T-Lib-Initiated SingleStepTransfer re-INVITE	19
T-Lib-Initiated SingleStepTransfer REFER	20
T-Lib-Initiated SingleStepTransfer User Busy	21
T-Lib-Initiated Consult Call Single EndPoint	22
T-Lib-Initiated Consult Transfer re-INVITE	25
T-Lib-Initiated 3-Way Conf Central Mixing	28
Answer Call Functionality	31

Overview

Important

The following features are not supported:

- Video
- IPv6
- NAT, ICE, TURN, STUN
- SIP Cluster
- SIP Proxy

The SIP Endpoint SDK for OS X enables you to build a SIP endpoint that can seamlessly connect agent desktop applications with the Genesys SIP Server in order to handle audio and video calls. This overview will help you understand its design goals and its architecture.

Design Goals of SIP Endpoint SDK

Many SIP softphones that are currently available require an agent to interact with their own separate user interface in addition to that of the agent desktop application. For instance, the agent might have to use the SIP phone interface to answer a call, while other actions — such as holding or releasing the call — would have to be done via agent desktop interface.

In contrast to that, the SIP Endpoint SDK is designed to be integrated into an agent desktop so the agent can use a single user interface to control calls. Genesys recommends that this be done in a way that leaves actual control in the hands of a T-Lib-based agent desktop application, which has a fuller feature set and is also fully supported by Genesys.

The SIP Endpoint SDK is also designed to integrate with the Genesys SIP Server. It supports the SIP, SDP, and RTP/RTCP protocols.

In addition to these principal design goals, the SIP Endpoint SDK supports the following features:

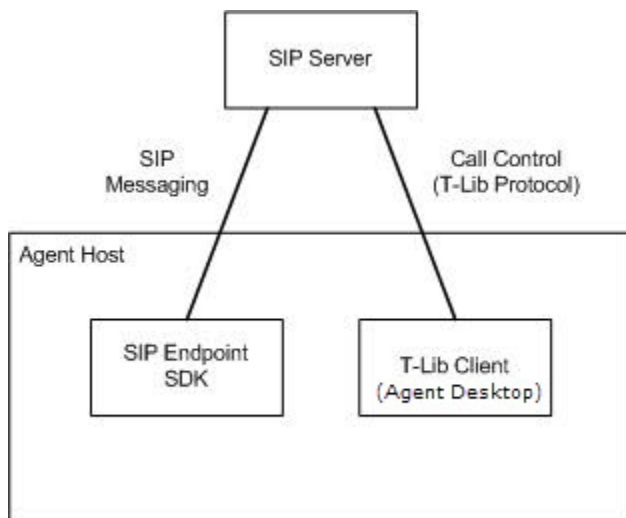
- Traditional call control functionality such as:
 - Establishing inbound and outbound calls
 - Hold and retrieve
 - Transfers and conference calls (Note that SIP Endpoint SDK does not handle conference calls directly. You must use third-party call control in order to enable these features.)
- Multi-line call handling

Note: Multi-line call handling is available up to the limitations on the number of simultaneous connections supported by the licensing available to Genesys at this time.

Architecture

The SIP Endpoint SDK is a software component that resides on the agent's computer. It uses the SIP protocol to communicate with the Genesys SIP Server and supports both thin (or 3-tier) and thick (or rich) T-Lib clients.

In SIP messaging terms, the SIP Endpoint is a user agent. The contact center agent's computer should also host a Genesys T-Lib client that works with SIP Server and provides Genesys agent-related functionality, as shown below.



As you can see, the SIP Endpoint SDK handles the SIP messaging, and the T-Lib Client (Agent Desktop) uses the T-Lib protocol to handle call control.

SDK Components

The components of the SIP Endpoint SDK are described on the following page:

- [SIP Endpoint SDK for OS X Components](#)

Supported Codecs

SIP Endpoint SDK 8.5.2 for Apple OS supports the following codecs:

- G.711 (PCMA, PCMU)
- G.722
- iLBC (Internet Low Bitrate Codec)
- iSAC (Internet Speech Audio Codec)

- [Opus](#)

Learning More

To continue learning about the SIP Endpoint SDK, we recommend you read the pages describing [SIP-Based Third-Party Call Control](#) to understand messaging patterns. After that, you should be ready to use the [Deployment Guide](#) to install the SDK on your system.

Once you have installed the SDK, you may want to check the following page from the [Developer's Guide](#):

- [SIP Endpoint SDK Configuration for Apple OS](#)

SIP Endpoint SDK for OS X Components

Important

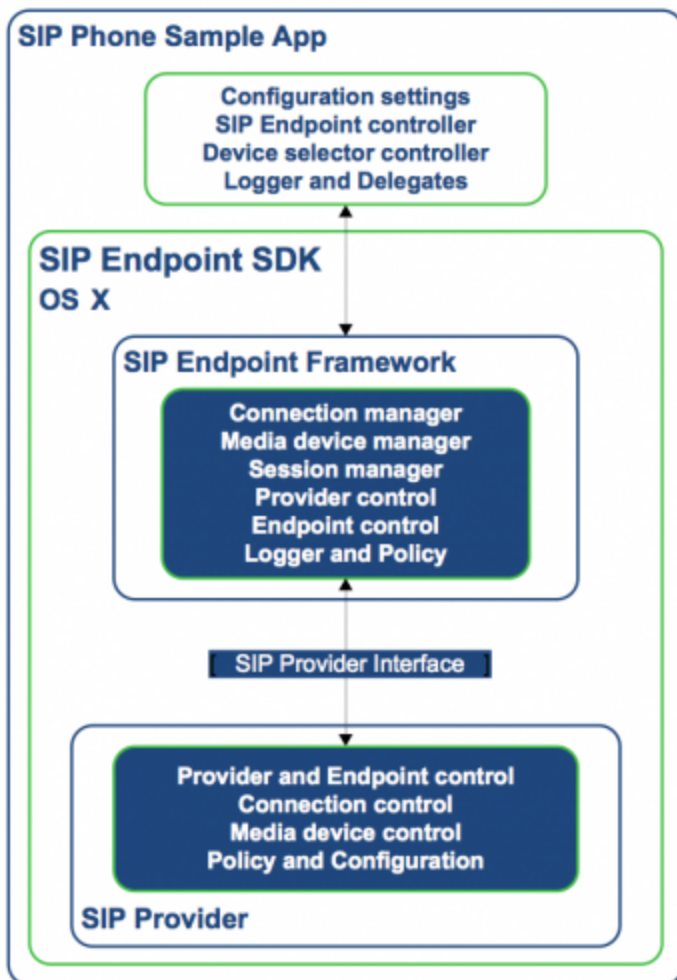
The following features are not supported:

- Video
- IPv6
- NAT, ICE, TURN, STUN
- SIP Cluster
- SIP Proxy

The SIP Endpoint SDK for OS X distribution consists of the following main components:

- The SDK itself, which provides all of the SIP-related functionality by way of the SIP Endpoint Framework. The SDK also provides an interface that you can use to integrate it into different GUI-based applications.
- A [QuickStart application](#) which is built on the SIP Endpoint SDK.

The SDK itself communicates with a SIP Provider interface, which links the SDK to the SIP Provider, as shown in the following diagram.



SIP-Based Third-Party Call Control

This section provides sequence diagrams and descriptions of the call scenarios supported by SIP Server and the SIP Endpoint SDK, which are achieved using the T-Lib API. It also demonstrates how the T-Lib API is mapped into SIP messaging.

Note that the SIP Endpoint SDK supports the talk/hold NOTIFY extension.

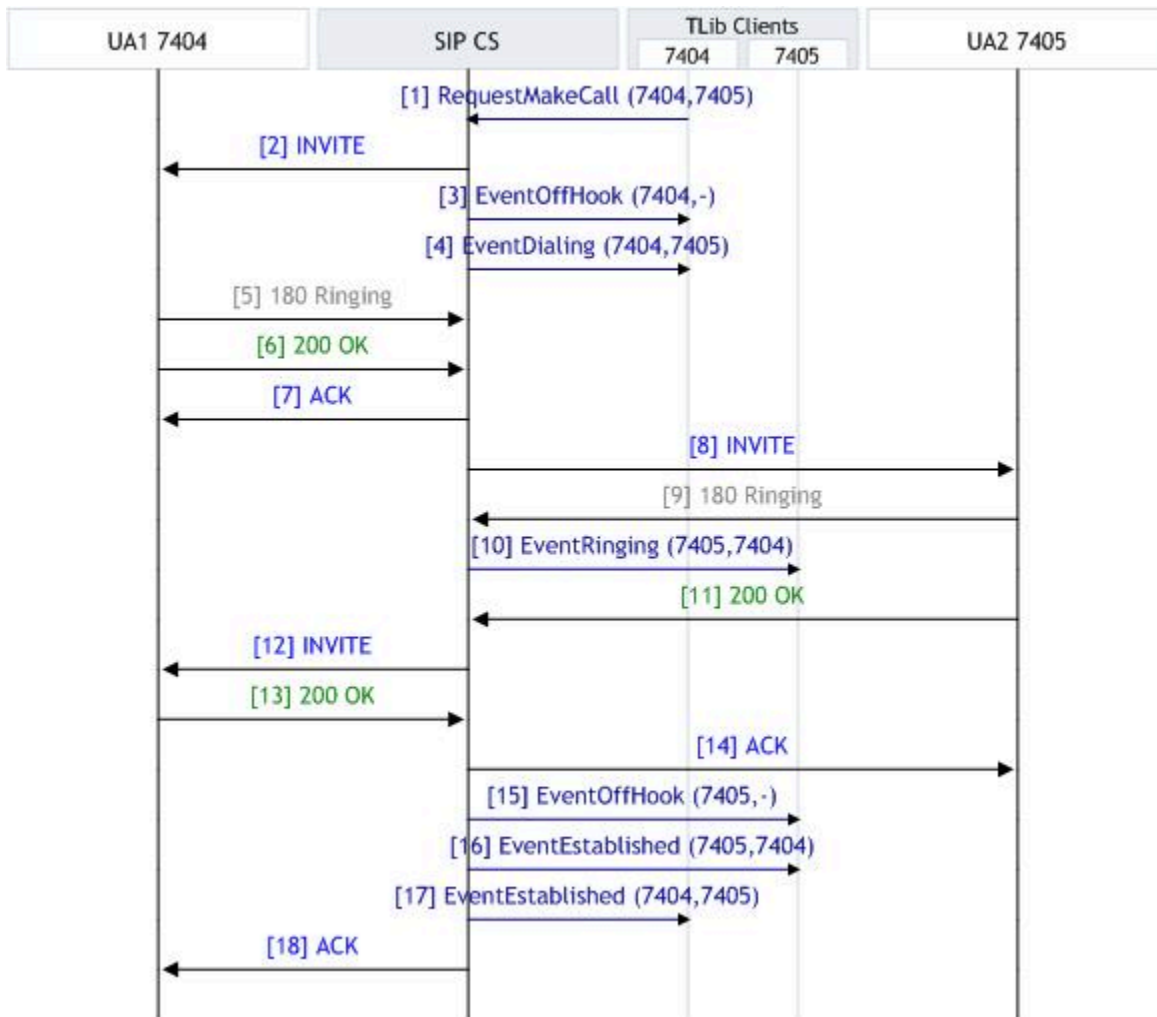
T-Lib-Initiated MakeCall re-INVITE

Description

UA1 calls UA2. If the UA1 endpoint does not support the REFER method, then the re-INVITE mechanism could be used. The main disadvantages of re-INVITE are:

- The phone may not reflect the real call progress unless early media is activated between both endpoints.
- Re-INVITE makes it next to impossible for third-party switches to track call topology and accumulate accurate call detail records (CDR).

Diagram




T-Lib-Initiated MakeCall INVITE

Description

RFC 3725 is a configuration flag that can be set to 2 (the default, which causes the **T-Lib-Initiated MakeCall re-INVITE** call scenario to be used) or 1. When the parameter is set to 1, the call flow will be as given below. Notice in this case that according to the RFC (best current practices for third-party call control) UA1 does not get an ACK until UA2 has responded with 200 OK. In this case, SIP Server does not send a re-INVITE to UA1.

Note: Usually, the **re-INVITE call flow** is the preferred use case. This is the default behavior of SIP Server, unless the RFC3725 parameter is set to 1.

Diagram




T-Lib Termination TReleaseCall

Description

The T-Lib client initiates the termination of the call by invoking TReleaseCall from the agent desktop at UA2. SIP Server terminates the existing call from UA1 to UA2 by sending two BYE requests.

Diagram



T-Lib-Initiated Hold-Retrieve

Description

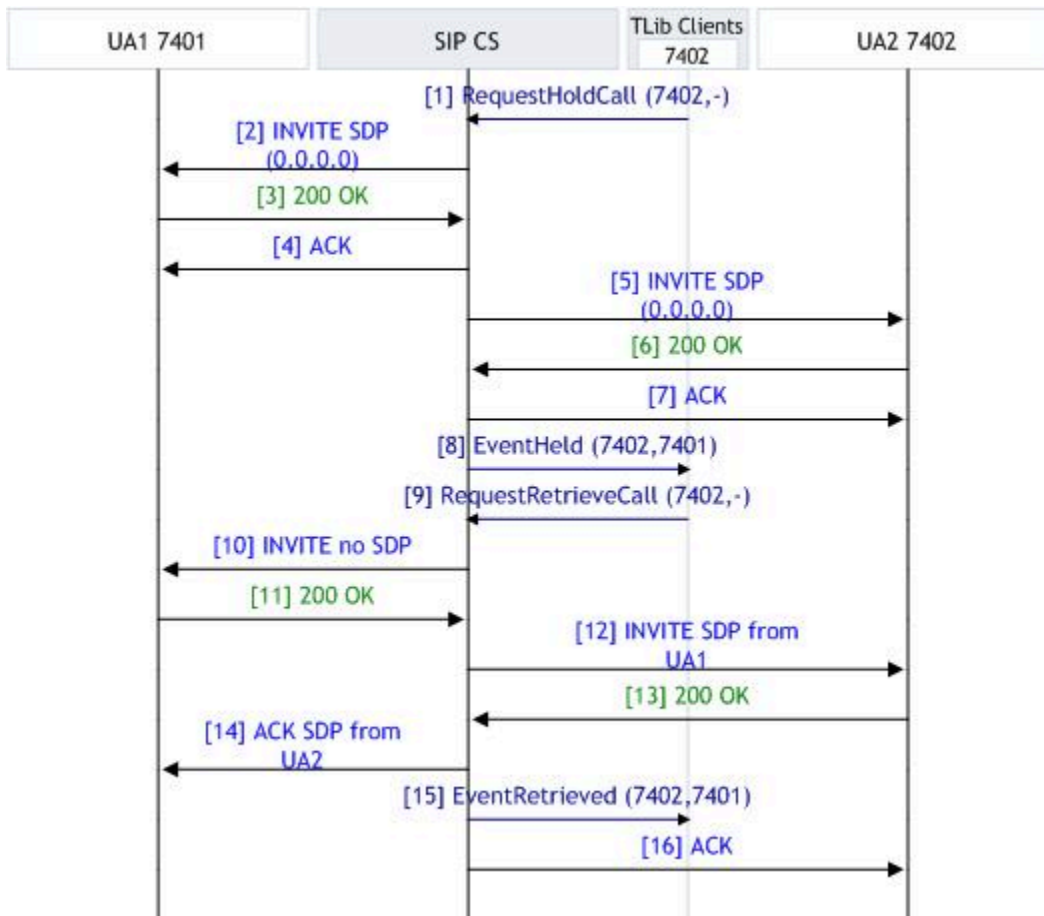
SIP Server's Hold operation can be configured to play music while the customer is placed on hold by an agent. The music-on-hold configuration is global — if it is enabled, all held parties hear music; if it is not enabled, all held parties hear silence.

Hold-with-silence only involves stopping the RTP streams. The Retrieve operation resumes the RTP streams.

There are multiple methods for stopping the RTP streams. SIP Server uses IP address 0.0.0.0 in the `c=` attribute as the only universally accepted method.

Because some endpoints may change the RTP port address when trying to retrieve a call, the re-INVITE with delayed SDP offer-answer negotiation shown in the following diagram is necessary.

Diagram




T-Lib-Initiated Hold-Retrieve Music on Hold

Description

The Music-on-Hold (MOH) Server is a regular SIP UA that streams music in an RTP stream. Stream Manager is used to provide Music On Hold.

In the scenario illustrated below, the agent at UA1 places UA2 on hold by invoking the T-Lib request THoldCall. UA2 will hear music while on hold; UA1 will hear silence.

Diagram



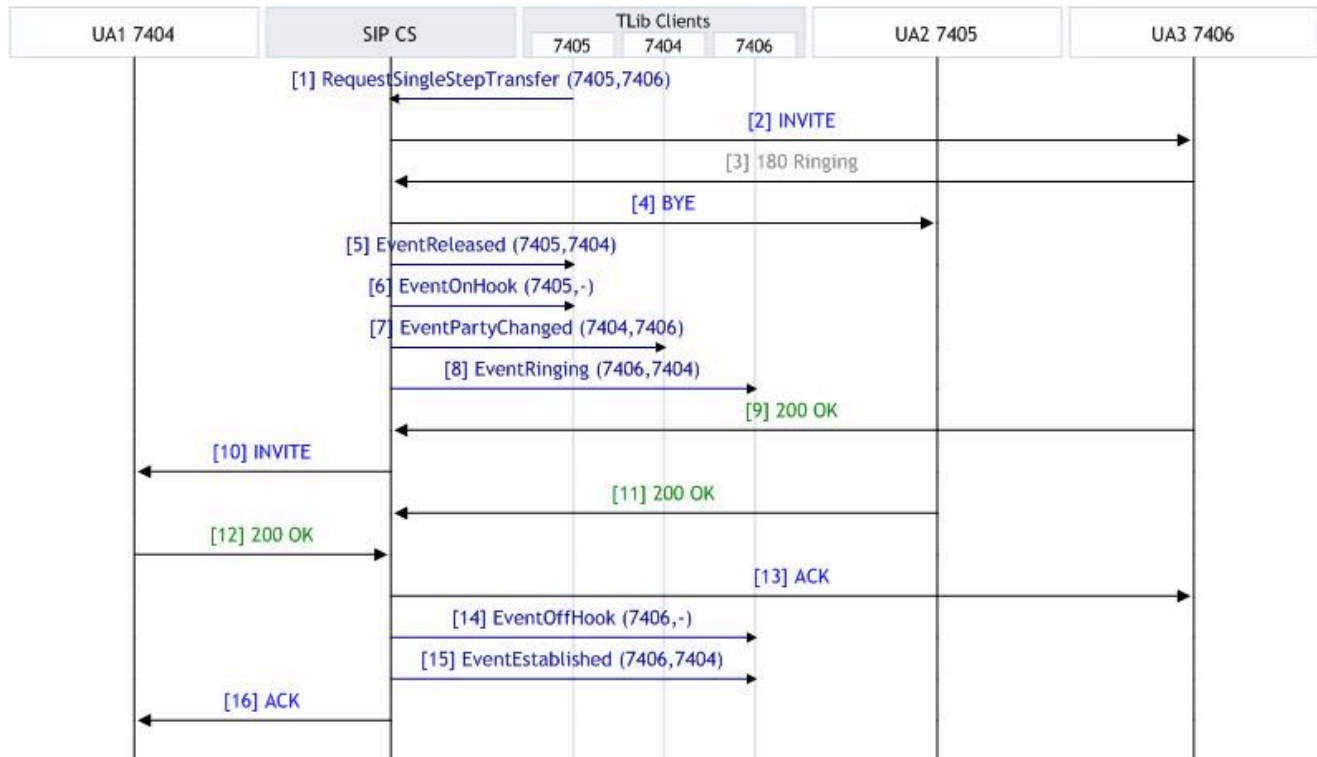
T-Lib-Initiated SingleStepTransfer re-INVITE

Description

UA1 calls UA2, and the agent at UA2 invokes a T-Lib request to transfer the call from UA2 to UA3.

For the advantages and disadvantages of the re-INVITE mechanism involved in this scenario, see the description in the section on [T-Lib-Initiated MakeCall re-INVITE](#).

Diagram



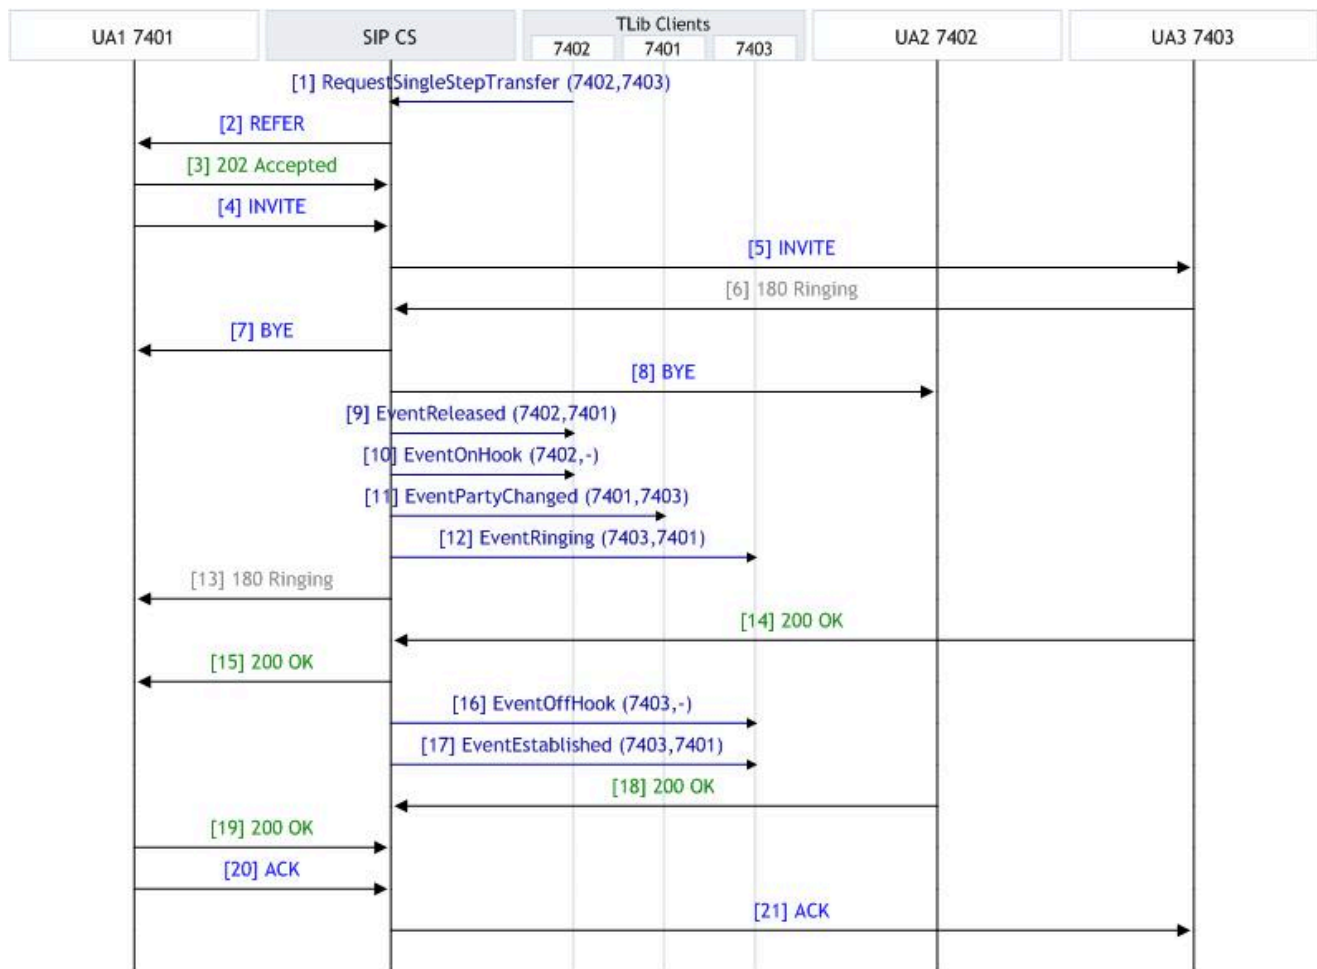
T-Lib-Initiated SingleStepTransfer REFER

Description

UA1 calls UA2, and the T-Lib request transfers the call from UA2 to UA3.

This scenario is equivalent to a single-step transfer using the REFER method executed from UA2, except that in this case, the REFER message originates on SIP Server, not on UA2. Also, the "202 Accepted" message comes back to SIP Server, not UA2, and SIP Server also disconnects the call between UA1 and UA2 using a third-party call control release.

Diagram




T-Lib-Initiated SingleStepTransfer User Busy

Description

If a single-step transfer is attempted and UA3 returns a busy signal or a no-answer event and rejects the attempted transfer, UA2 may or may not be available to pick up the dropped call. Note that due to the nature of single-step transfers, a call can be dropped in these cases. In addition to that, there is no guarantee that UA2 will still be available if UA3 reports busy or does not answer within a reasonable time. Thus, if your contact center requires that calls not be dropped, you should use consult transfer, not single-step transfer.

Diagram




T-Lib-Initiated Consult Call Single EndPoint

Description

There are some cases when it is not possible to use the normal means of initiating a consult call by setting the dialog of the main call on hold and creating a new dialog for the consult call. One such case involves a transfer initiated by a remote agent with a regular (non-VoIP) phone behind the media-gateway. In that case, a consult call can be initiated by "borrowing" the SIP Dialog from the main call and re-using it for the consult call.

The diagram below describes this kind of case. To provide Ring-Back to the agent during the progress of the call, the SIP Dialog is invited to Stream Manager, which will play a ring-back tone. If the consult call is finished by a TCompleteTransfer request (as shown in the diagram), that dialog will be terminated. If the agent decides to reconnect the call, the SIP Dialog will be returned to the main call and retrieved.

Diagram



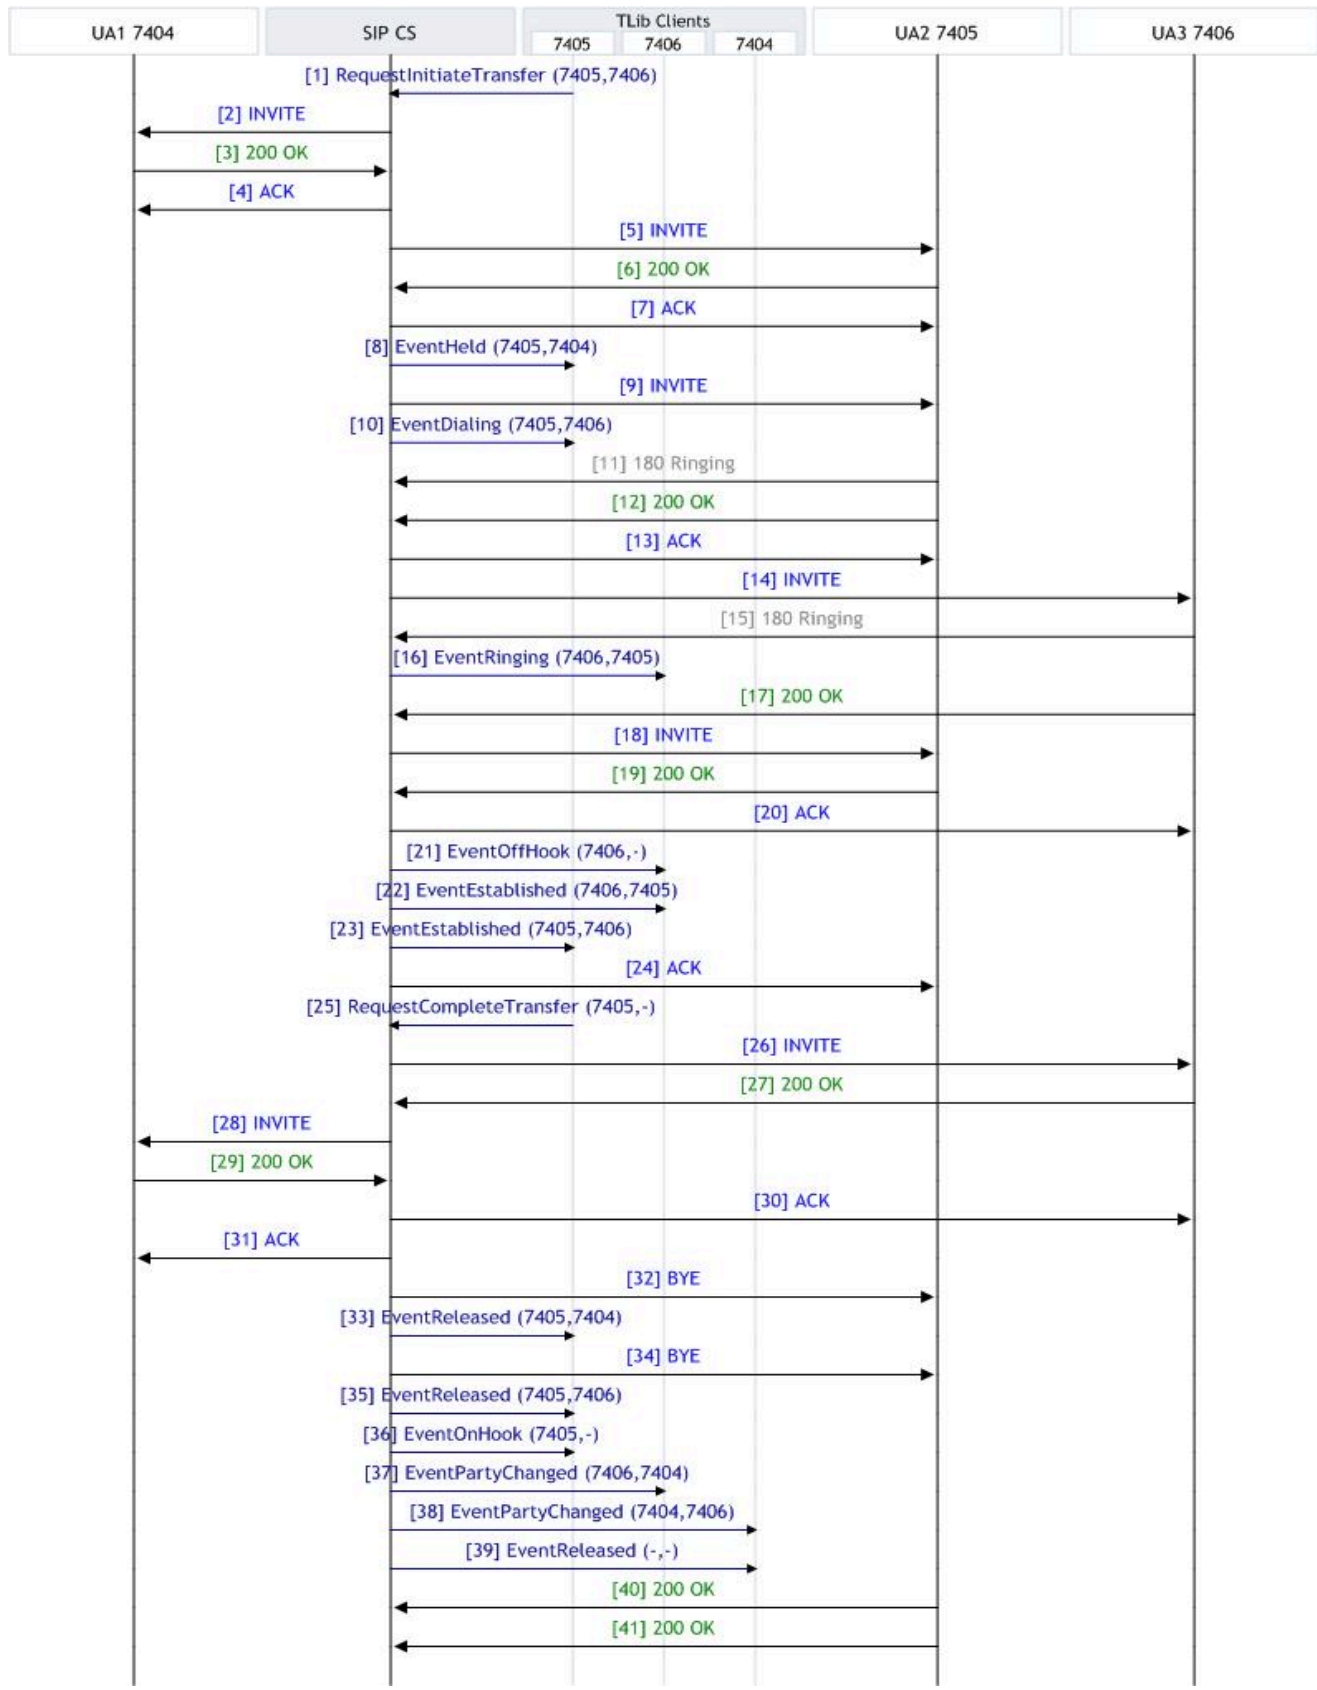
T-Lib-Initiated Consult Transfer re-INVITE

Description

SIP Server implements a `TInitiateTransfer` by invoking `THoldCall` and `TMakeCall` in a sequence.

The call scenario for `TCompleteTransfer` is shown here. It assumes a call from UA1 to UA2 is placed on hold and that a call from UA2 to UA3 is in the Established state.

Diagram



T-Lib-Initiated 3-Way Conf Central Mixing

Description

T-Lib-controlled three-way conferences are handled through an external multimedia conferencing unit (MCU) device that performs the audio mixing.

The call scenario begins when a call is established between UA1 and UA2. UA2 places UA1 on hold and dials a consult call to UA3. When UA2 and UA3 are ready for a conference, SIP Server re-Invites everybody to the mixing device. A DN (that is, the user part of a SIP URI) identifies a conference — all participants will use the same number to identify the conference to the mixing device.

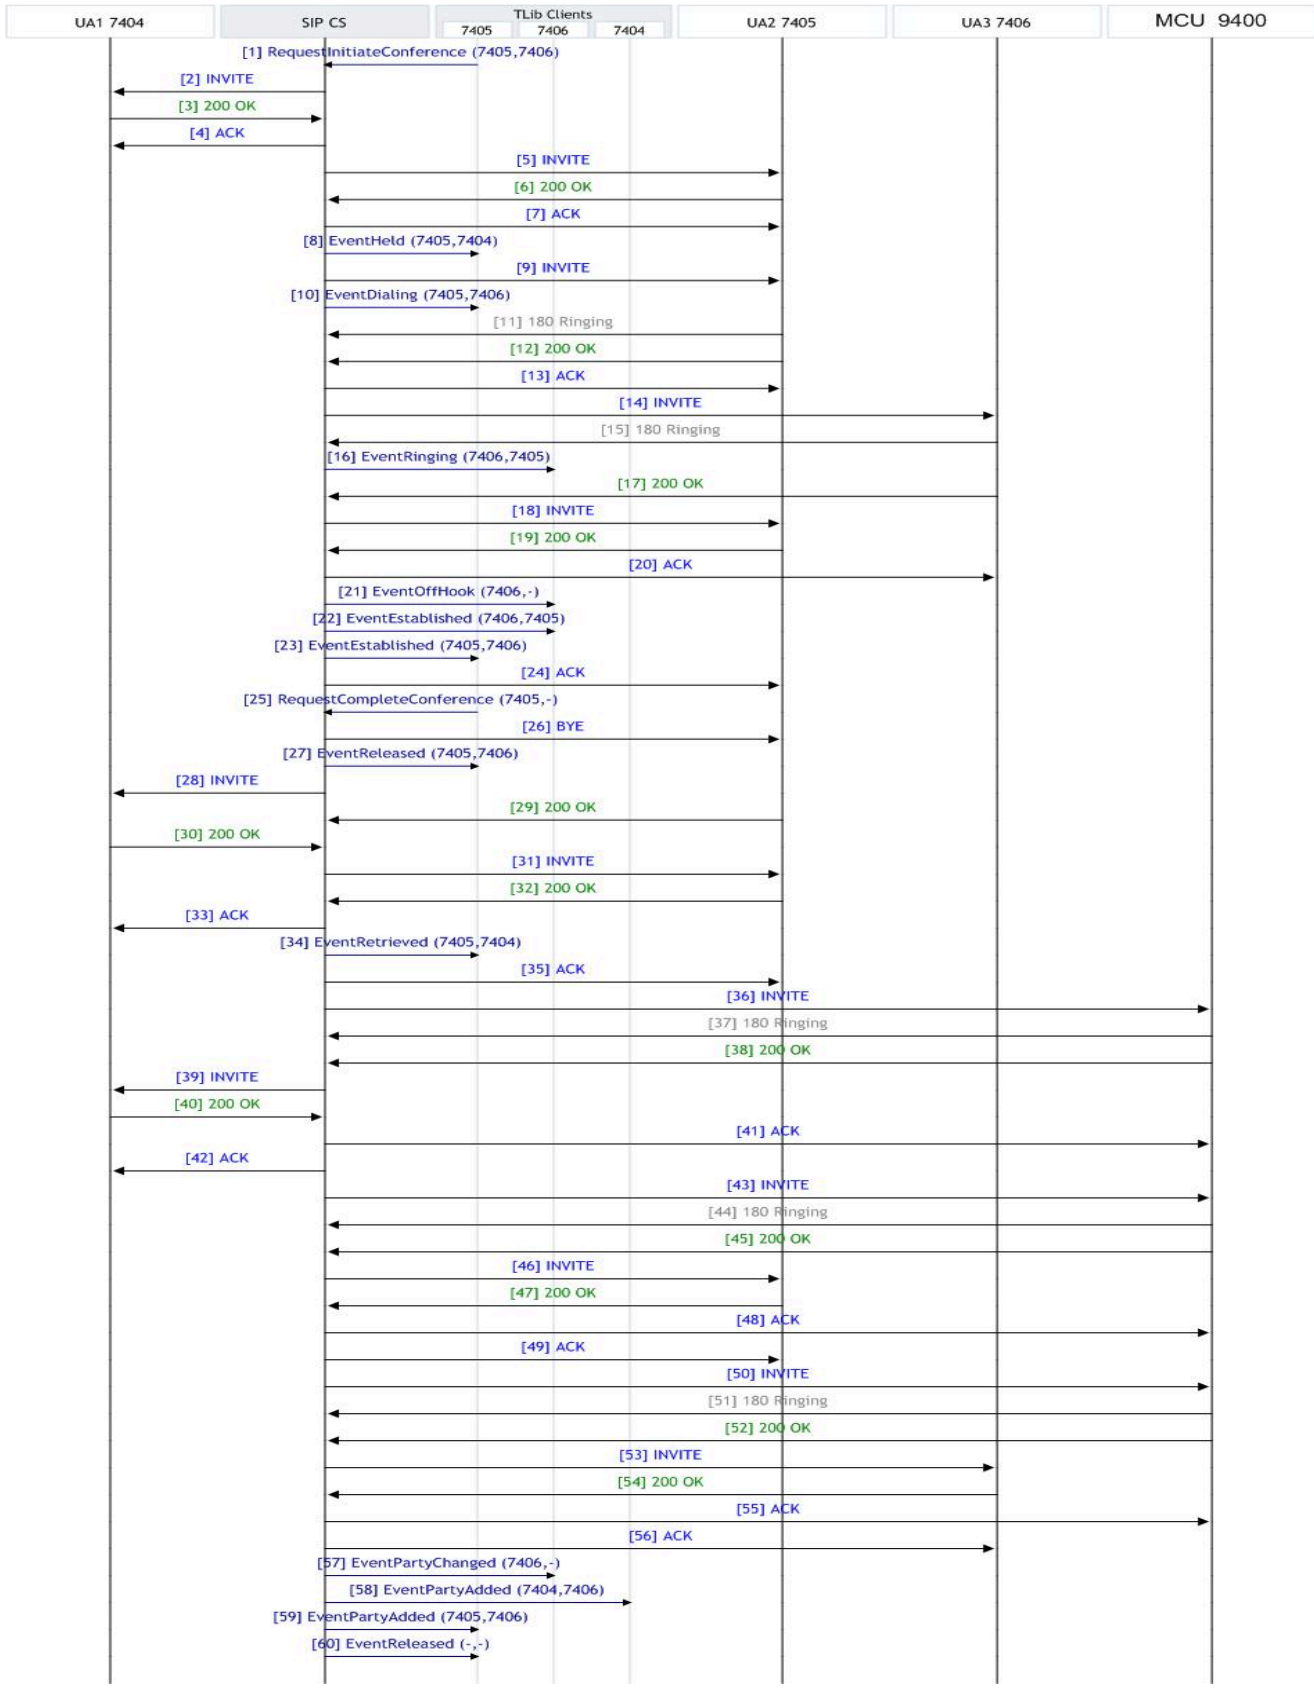
The diagram begins with a call between UA1 and UA2 in a Held state and a call from UA2 to UA3 in an Established state. Agent UA2 invokes the T-Lib request `TCompleteConference`.

Here is how this process works:

1. To allocate MCU resources, SIP Server starts by sending the MCU one INVITE request for each participant request without an SDP.
2. SIP Server sends INVITEs without the SDP offering to get the list of codecs from the MCU.
3. SIP Server does not modify any of the existing dialogs until it receives an OK for all INVITEs sent to the MCU. If any of the INVITEs fail, SIP Server rolls back to the initial state and all SIP dialogs remain intact.
4. Once SIP Server receives all the OKs, it sends re-INVITE dialogs to the three participants, using the SDPs it received from the MCU.
5. SIP Server uses the answer SDPs from the participants in the ACKs it sends to the MCU.
6. Before the conference is completed, UA2 will have had two dialogs — one for the primary call and one for the consult call. SIP Server re-INVITEs the primary call to the MCU; it sends BYE to terminate the consult call.

Implementation Note: The scenario just described is the most efficient way to establish a three-way conference using the T-Lib API. However, some MCU devices do not support INVITE without an SDP. To accommodate these devices, SIP Server sends INVITEs to MCU with valid non-hold SDPs (the actual SDP received from the corresponding devices). After that, it re-INVITEs the devices.

Diagram



Answer Call Functionality

Description

If an endpoint (SIP phone) supports remote CTI control by providing the SIP extensions event package known as "BroadSoft" SIP extensions, SIP Server allows SIP Server clients to answer the call by means of the TAnswerCall request. When responding to the TAnswerCall request, SIP Server will send an event with the header "Event" and a header value of "talk" to the SIP phone, in a ringing state of NOTIFY. This event instructs the SIP phone to answer the call without human interference. Note that SIP phones that support such an event package should send a 180 Ringing message header of "Allow-Events:", which should contain "talk" inside its value.

Diagram

