Internet Engineering Task Force (IETF)
Request for Comments: 7092
Category: Informational
ISSN: 2070-1721

H. Kaplan
Oracle
V. Pascual
Quobis

December 2013

A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents

Abstract

In many SIP deployments, SIP entities exist in the SIP signaling path between the originating and final terminating endpoints, which go beyond the definition of a SIP proxy, performing functions not defined in Standards Track RFCs. The only term for such devices provided in RFC 3261 is for a Back-to-Back User Agent (B2BUA), which is defined as the logical concatenation of a SIP User Agent Server (UAS) and User Agent Client (UAC).

There are numerous types of SIP B2BUAs performing different roles in different ways; for example, IP Private Branch Exchanges (IPBXs), Session Border Controllers (SBCs), and Application Servers (ASs). This document identifies several common B2BUA roles in order to provide taxonomy other documents can use and reference.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7092.

Table of Contents

1. Introduction

In current SIP deployments, there are numerous forms of Back-to-Back User Agents (B2BUAs), operating at various levels of transparency and for various purposes, with widely varying behavior from a SIP perspective. Some act as pure SIP proxies and only change to the role of B2BUA in order to generate BYEs to terminate dead sessions. Some are full User Agent stacks with only high-level event and application logic binding the User Agent Server (UAS) and User Agent Client (UAC) sides. Some B2BUAs operate only in the SIP signaling plane, while others participate in the media plane as well.

As more SIP domains are deployed and interconnected, the probability of a single SIP session crossing multiple B2BUAs at both the signaling and media planes increases significantly.

This document provides a taxonomy of several common B2BUA roles so that other documents may refer to them using their given names without redefining them in each document.

2. Terminology

The following terms are defined in [RFC3261], Section 6.

B2BUA: a SIP Back-to-Back User Agent, which is the logical combination of a User Agent Server (UAS) and User Agent Client (UAC).

UAS: a SIP User Agent Server.

UAC: a SIP User Agent Client.

3. B2BUA Role Types

Within the context of this document, the classification refers to a B2BUA role, not a particular system type. A given system type may change its role in the middle of a SIP session, for example, when a stateful proxy tears down a session by generating BYEs or when an SBC [RFC5853] performs transcoding or REFER termination.

Furthermore, this document defines "B2BUA" following the definition provided in [RFC3261], which is the logical concatenation of a UAS and UAC. A typical centralized conference server, for example, is not a B2BUA because it is the target UAS of multiple UACs, whereby the UACs individually and independently initiate separate SIP sessions to the central conference server. Likewise, a third-party call control transcoder, as described in Section 3.1 of [RFC5369], is

not a B2BUA, whereas an inline (conference bridge) transcoder based on [RFC5370] is a B2BUA.

3.1. Signaling Plane B2BUA Roles

A signaling plane B2BUA is one that operates only on the SIP message protocol layer and only with SIP messages and headers but not with the media itself in any way. This implies that it does not modify the Session Description Protocol (SDP) bodies, although it may save them and/or operate on other MIME body types. This category is further subdivided into specific roles as described in this section.

It should be noted that there is a large variety of modifications made by "signaling plane B2BUAs".

3.1.1. Proxy-B2BUA

A Proxy-B2BUA is one that appears, from a SIP perspective, to be a SIP proxy based on [RFC3261] and its extensions, except that it maintains a sufficient dialog state to generate in-dialog SIP messages on its own and does so in specific cases. The most common example of this is a SIP proxy that can generate BYE requests to tear down a dead session.

A Proxy-B2BUA does not modify the received header fields such as To, From, or Contact, and it only modifies the Via and Record-Route header fields following the rules in [RFC3261] and its extensions. If a Proxy-B2BUA can generate in-dialog messages, then it will also need to modify the CSeq header after it has generated its own. A Proxy-B2BUA neither modifies nor inspects MIME bodies (including SDP), does not have any awareness of media, will forward any method type, etc.

3.1.2. Signaling-only

A Signaling-only B2BUA is one that operates at the SIP layer but in ways beyond those of [RFC3261] proxies, even for normally forwarded requests. For example, such a B2BUA might replace the Contact URI, modify or remove all Via and Record-Route headers, modify the To and From header fields, modify or inspect specific MIME bodies, etc. No SIP header field is guaranteed to be copied from the received request on the UAS side to the generated request on the UAC side.

An example of such a B2BUA would be some form of an Application Server and a PBX, such as a server that locally processes REFER requests and generates new INVITEs on behalf of the REFER’s target. Another example would be a privacy service proxy [RFC3323] performing the ’header’ privacy function.

3.1.3. SDP-Modifying Signaling-only

An SDP-Modifying Signaling-only B2BUA is one that operates in the signaling plane only and is not in the media path, but it does modify SDP bodies and is thus aware of and understands SDP syntax and semantics. In some cases, some Application Servers and PBXs act in this role, for example, to remove certain codec choices or merge two media endpoints into one SDP offer.

These B2BUAs don't do anything that changes the path that the media takes (in particular, they don't insert themselves on the media path), but they may make SDP changes that affect what is sent on the media plane.

3.2. Signaling/Media Plane B2BUA Roles

A signaling/media plane B2BUA is one that operates at both the SIP and media planes and not only on SIP messages but also on SDP and potentially the Real-time Transport Protocol (RTP) / the Real-Time Control Protocol (RTCP) [RFC3550] or other media. Such a B2BUA may or may not replace the Contact URI, modify or remove all Via and Record-Route headers, modify the To and From header fields, etc. No SIP header field is guaranteed to be copied from the received request on the UAS side to the generated request on the UAC side, and SDP will also be modified.

An example of such a B2BUA would be a Session Border Controller (SBC) performing the functions defined in [RFC5853], a B2BUA transcoder as defined in [RFC5370], a rich-ringtone Application Server, or a recording system. Another example would be a privacy service proxy [RFC3323] performing the 'session' privacy function.

Note that a signaling/media plane B2BUA need not be instantiated in a single physical system, but it may be decomposed into separate signaling and media functions.

The signaling/media plane B2BUA category is further subdivided into specific roles as described in this section.

3.2.1. Media Relay

A B2BUA that performs a media-relay role is one that terminates the media plane at the IP and transport (UDP/TCP) layers on its UAS and UAC sides, but neither modifies nor restricts which forms of payload are carried within the packets. Rather, the payload is transparently copied from one side to the other. Such a role may or may not support only UDP, only TCP, both UDP and TCP, as well as other transport types. It may also involve policing the IP packets to fit

within a bandwidth limit or converting from IPv4 to IPv6, or vice versa. This is typically similar to NAT behavior, except a NAT operating in both directions on both the source and destination information; it is often found as the default behavior in SBCs.

3.2.2. Media Aware

A B2BUA that performs a media-aware role is similar to a media relay except that it inspects and potentially modifies the payload carried in UDP or TCP (as it could be RTP or RTCP [RFC3550]), but it does not at a codec or higher layer. An example of such a role is a Secure Real-time Transport Protocol (SRTP) [RFC3711] terminator, which does not need to care about the RTP payload but does care about the RTP header; or, a device that monitors RTCP for QoS information; or, a device that multiplexes/demultiplexes RTP and RTCP packets on the same 5-tuple.

3.2.3. Media Termination

A B2BUA that performs a media-termination role is one that operates at the media payload layer, such as RTP/RTCP codec or the Message Session Relay Protocol (MSRP) message layer and higher. Such a role may only terminate/generate specific RTP media, such as dual-tone multi-frequency (DTMF) packets [RFC4733], or it may convert between media codecs or act as a Back-to-Back MSRP [RFC4975] agent. This is the role performed by transcoders, conference servers based on [RFC5366], etc.

4. Mapping SIP Device Types to B2BUA Roles

Although the B2BUA roles defined previously do not define system types, as discussed in Section 3, some discussion of what common system types perform which defined roles may be appropriate. This section provides such a ’mapping’ for general cases to aid in understanding of the roles.

4.1. SIP PBXs and Softswitches

SIP-enabled Private Branch Exchanges (SIP PBXs) and softswitches are market category terms and are not specified in any standard. In general, they can perform every role described in this document at any given time based on their architecture or local policy. Some are based on architectures that make them the equivalent of a SIP UAS and UAC connected with a logical Primary Rate Interface (PRI) loopback; others are built as a SIP proxy core with optional modules to "do more".

4.2. Application Servers

Application Servers are a broad marketing term and are not specified in any standard in general, although the Third Generation Partnership Project (3GPP) and other organizations specify some specific Application Server functions and behaviors. Common examples of Application Server functions are message-waiting indicators (MWIs), Find Me/Follow Me services, privacy services, call center Automatic Call Distributor (ACD) services, call screening, and Voice Call Continuity (VCC) services. Some only operate in the signaling plane in either Proxy-B2BUA or Signaling-only B2BUA roles; others operate as full Media-termination B2BUAs, such as when providing Interactive Voice Recognition (IVR), rich ringtones, or integrated voicemail services.

4.3. Session Border Controllers

Session Border Controllers (SBCs) are a market category term and are not specified in any standard. Some of the common functions performed by SBCs are described in [RFC5853], but in general, they can perform every role described in this document at any given time based on local policy. By default, most SBCs are either Media-relay or Media-aware B2BUAs and replace the Contact URI; remove the Via and Record-Route headers; modify Call-ID, To, From, and various other headers; and modify SDP. Some SBCs remove all headers, all bodies, and reject all method types unless explicitly allowed by local policy; other SBCs pass all such elements through unless explicitly forbidden by local policy.

4.4. Transcoders

Transcoders perform the function of transcoding one audio or video media codec type to another, such as G.711 to G.729. As such, they perform the Media-termination role, although they may only terminate media in specific cases of codec mismatch between the two ends. Although [RFC5369] and [RFC5370] define two types of SIP transcoders, in practice, a popular variant of the inline conference bridge model [RFC5370] is to behave as a SIP B2BUA without using the resource-list mechanism but rather simply routing a normal INVITE request through a B2BUA with a built-in transcoder. SIP transcoder architectures are based on everything from SIP media servers and SBCs to looped-back Time Division Multiplexing (TDM) gateways, and thus run the gamut from replacing only specific headers/bodies and SDP content needed to perform their function to replacing almost all SIP headers and SDP content. Some transcoders save and remove SDP offers in INVITEs from the UAC, and wait for an offer in the response from the UAS, similar

to a Third Party Call Control (3PCC) model; others just insert additional codecs in SDP offers and only transcode if the inserted codec(s) is selected in the answer.

4.5. Conference Servers

In general, conference servers do not fall under the term "B2BUA" as defined by this document, since typically they involve multiple SIP UACs initiating independent SIP sessions to the single conference UAS. However, a conference server supporting [RFC5366], whereby the received INVITE triggers the conference focus UAS to initiate multiple INVITEs as a UAC, would be in a Media-termination B2BUA role when performing that function.

4.6. P-CSCF and IBCF Functions

The Proxy-Call Session Control Function (P-CSCF) and the Interconnection Border Control Function (IBCF) are defined by 3GPP [IMS] standards, and when coupled with the IP Multimedia Subsystems (IMS) media plane gateways (IMS Access Gateway (AGW), Transition Gateway (TrGW), etc.), they typically form a logical Media-relay or Media-aware B2BUA role.

4.7. S-CSCF Function

The Serving-Call Session Control Function (S-CSCF) is defined by 3GPP [IMS] standards and typically follows a Proxy-B2BUA role.

5. Security Considerations

Security risks are specific to each type of B2BUA, so little can be said in general. Of course, adding extra systems in the communication path creates extra points of attack, reduces or eliminates the ability to perform end-to-end encryption, decreases the privacy of SIP communications, and complicates trust models. Thus, every B2BUA design requires particular attention to security analysis.

A few general points can be made:

1. The B2BUA processing of SDP and media packets is an impediment to the deployment of end-to-end SRTP and reduces the ability to deploy new, stronger forms of SRTP key exchange.

2. The ability for B2BUAs to modify any SIP header field value and body impacts the ability to deploy SIP identity and message integrity.

3. The management and configuration mechanisms of B2BUAs are a tempting point of attack and must be strongly defended.

Further security considerations related to the functionality described in this document are addressed in the relevant references.

6. Informative References

[RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

[RFC3323] Peterson, J., "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, November 2002.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

[RFC4733] Schulzrinne, H. and T. Taylor, "RTP Payload for DTMF Digits, Telephony Tones, and Telephony Signals", RFC 4733, December 2006.

[RFC4975] Campbell, B., Mahy, R., and C. Jennings, "The Message Session Relay Protocol (MSRP)", RFC 4975, September 2007.

[RFC5366] Camarillo, G. and A. Johnston, "Conference Establishment Using Request-Contained Lists in the Session Initiation Protocol (SIP)", RFC 5366, October 2008.

[RFC5369] Camarillo, G., "Framework for Transcoding with the Session Initiation Protocol (SIP)", RFC 5369, October 2008.

[RFC5370] Camarillo, G., "The Session Initiation Protocol (SIP) Conference Bridge Transcoding Model", RFC 5370, October 2008.

[RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.

[IMS] 3GPP, "IP Multimedia Subsystem (IMS); Stage 2, 3GPP TS 23.228", Version 12.2.0, September 2013.

Authors' Addresses

Hadriel Kaplan
Oracle

EMail: hadriel.kaplan@oracle.com

Victor Pascual
Quobis

EMail: victor.pascual@quobis.com