

Internet Engineering Task Force (IETF)
Request for Comments: 7879
Category: Standards Track
ISSN: 2070-1721

R. Ravindranath
T. Reddy
G. Salgueiro
Cisco
V. Pascual
Oracle
P. Ravindran
Nokia Networks
May 2016

DTLS-SRTP Handling in SIP Back-to-Back User Agents

Abstract

Session Initiation Protocol (SIP) Back-to-Back User Agents (B2BUAs) exist on the signaling and media paths between the endpoints. This document describes the behavior of B2BUAs when Secure Real-time Transport (SRTP) security context is set up with the Datagram Transport Layer Security (DTLS) protocol.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in Section 2 of RFC 7841.

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/rfc7879>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 3
 - 1.1. Overview 3
 - 1.2. Goals and Scope of this Document 4
- 2. Terminology 4
- 3. B2BUAs Procedures to Allow End-to-End DTLS-SRTP 5
- 4. Signaling-Plane B2BUA Handling of DTLS-SRTP 5
 - 4.1. Proxy-B2BUAs 6
 - 4.2. Signaling-Only and SDP-Modifying Signaling-Only B2BUAs . 6
- 5. Media-Plane B2BUA Handling of DTLS-SRTP 6
 - 5.1. General 6
 - 5.1.1. Media Relay 6
 - 5.1.2. RTP- and RTCP-Aware Media-Aware B2BUA 8
- 6. Forking Considerations 9
- 7. Security Considerations 10
- 8. References 11
 - 8.1. Normative References 11
 - 8.2. Informative References 11
- Acknowledgments 12
- Contributors 12
- Authors' Addresses 13

1. Introduction

1.1. Overview

[RFC5763] describes how the Session Initiation Protocol (SIP) [RFC3261] can be used to establish a Secure Real-time Transport Protocol (SRTP) [RFC3711] security context with the Datagram Transport Layer Security (DTLS) protocol [RFC6347]. It describes a mechanism for transporting a certificate fingerprint using the Session Description Protocol (SDP) [RFC4566]. The fingerprint identifies the certificate that will be presented during the DTLS handshake. DTLS-SRTP is currently defined for point-to-point media sessions, in which there are exactly two participants. Each DTLS-SRTP session (described in Section 3 of [RFC5764]) contains a single DTLS connection (if RTP and RTP Control Protocol (RTCP) are multiplexed) or two DTLS connections (if RTP and RTCP are not multiplexed), and either two SRTP contexts (if media traffic is flowing in both directions on the same 5-tuple) or one SRTP context (if media traffic is only flowing in one direction).

In many SIP deployments, SIP Back-to-Back User Agents (B2BUA) entities exist on the SIP-signaling path between the endpoints. As described in [RFC7092], these B2BUAs can modify SIP and SDP information. For example, as described in Section 3.1.3 of [RFC7092], SDP-modifying signaling-only B2BUAs can potentially modify the SDP. B2BUAs can also be present on the media path, in which case they modify parts of the SDP information (like IP address, port) and subsequently modify the RTP headers as well. Such B2BUAs are referred to as "media-plane B2BUAs". [RFC7092] describes two different categories of media-plane B2BUAs, according to the level of activities performed on the media plane.

When B2BUAs are present in a call between two SIP User Agents (UAs), they often make end-to-end DTLS-SRTP sessions impossible. An "end-to-end DTLS-SRTP session" means that man-in-the-middle devices cannot break the DTLS-SRTP session between the endpoints. In other words, the man-in-the-middle device cannot create a separate DTLS-SRTP session between the client and the middle device on one side, and the middle device and the remote peer on the other side. B2BUAs may be deployed for address hiding or media latching [RFC7362], although Traversal Using Relays around NAT (TURN) and Interactive Connectivity Establishment (ICE) are expected to be used more often for this purpose as it provides better security properties. Such B2BUAs are able to perform their functions without requiring termination of DTLS-SRTP sessions, i.e., these B2BUAs need not act as DTLS proxy and decrypt the RTP payload.

1.2. Goals and Scope of this Document

A B2BUA could be deployed for address hiding or media latching as described in [RFC7362]. Such B2BUAs only terminate the media plane at the IP and transport (UDP/TCP) layers and may inspect the RTP headers or RTP Control Protocol (RTCP) packets. The goal of this specification is to provide guidance on how such B2BUAs function without breaking the end-to-end DTLS-SRTP session. A B2BUA could also terminate the media, or modify the RTP headers or RTP Control Protocol (RTCP) packets. Such B2BUAs will not allow end-to-end DTLS-SRTP. The recommendations made in this document are not expected to be applied by B2BUAs terminating DTLS-SRTP sessions given deployment reality.

This specification assumes that a B2BUA is not providing identity assurance and is not authorized to terminate the DTLS-SRTP session. A B2BUA that provides identity assurance on behalf of endpoints behind it can modify any portion of SIP and SDP before it generates the identity signature. As the B2BUA is generating the identity signature, it is not possible to detect if a B2BUA has terminated the DTLS-SRTP session. B2BUAs providing identity assurance and terminating DTLS-SRTP sessions are out of scope of this document.

The following sections describe the behavior B2BUAs can follow to avoid breaking end-to-end DTLS-SRTP sessions.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Transport Address: The combination of an IP address and port number.

The following generalized 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 a User Agent Client (UAC).

UAS: A SIP User Agent Server.

UAC: A SIP User Agent Client.

All of the pertinent B2BUA terminology and taxonomy used in this document are based on [RFC7092].

It is assumed the reader is already familiar with the fundamental concepts of the RTP protocol [RFC3550] and its taxonomy [RFC7656], as well as those of SRTP [RFC3711] and DTLS [RFC6347].

3. B2BUAs Procedures to Allow End-to-End DTLS-SRTP

A B2BUA MUST follow the rules mentioned below to allow end-to-end DTLS-SRTP sessions.

1. B2BUAs MUST forward the certificate fingerprint and SDP setup attribute it receives from one endpoint unmodified towards the other endpoint and vice versa.
2. The enhancements described in [RFC4474] provide a means for signing portions of SIP requests in order to provide identity assurance and certificate pinning by providing an identity signature over the SDP that carries the fingerprint of keying for DTLS-SRTP [RFC5763]. B2BUAs can identify that the enhancements in [RFC4474] are used for identity assurance if the SIP request contains both Identity and Identity-Info headers. In cases where endpoints use [RFC4474], B2BUAs MUST ensure that it does not modify any of the information used to construct the identity signature. This includes the entire SDP body and portions of the SIP header as described in [RFC4474]. In this case, a B2BUA cannot act as a media-relay B2BUA.
3. [SIP-ID] is introduced to overcome the limitations of [RFC4474] (discussed in Section 1 of [SIP-ID]). Unlike [RFC4474], [SIP-ID] does not generate an identity signature over material that intermediaries in the field commonly alter. In this case, a B2BUA can act as a media-relay B2BUA. B2BUAs can identify that [SIP-ID] is used for identity assurance if the SIP request contains an Identity header but does not include an Identity-Info header. The Identity-Info header is deprecated in [SIP-ID]. A B2BUA MUST ensure that it does not modify any of the headers used to construct the identity signature.
4. Both media relays and media-aware relays MUST NOT modify the authenticated portion of RTP and RTCP packets, and MUST NOT modify the authentication tag in the RTP and RTCP packets.

4. Signaling-Plane B2BUA Handling of DTLS-SRTP

Section 3.1 of [RFC7092] describes different categories of signaling-plane B2BUAs. This section explains how these B2BUAs are expected to comply with the recommendations in Section 3.

4.1. Proxy-B2BUAs

Proxy-B2BUAs, as defined in Section 3.1.1 of [RFC7092], modify only the Via and Record-Route SIP headers. These B2BUAs can continue to perform their function and still allow end-to-end DTLS-SRTP sessions since none of the headers used to construct the identity signature are modified.

4.2. Signaling-Only and SDP-Modifying Signaling-Only B2BUAs

These categories of B2BUAs are likely to modify headers that are used to construct the identity signature. For example, a signaling-only B2BUA can modify the Contact URI. Such B2BUAs are likely to violate rule 2 or rule 3 in Section 3. Depending upon the application requirements, such a B2BUA may be able to limit modification of header fields to those allowed to be modified by [RFC4474] or [SIP-ID].

5. Media-Plane B2BUA Handling of DTLS-SRTP

5.1. General

This section describes how the different types of media-plane B2BUAs defined in [RFC7092] are expected to comply with the recommendations in Section 3.

5.1.1. Media Relay

From an application-layer point of view, a media relay (as defined in Section 3.2.1 of [RFC7092]) forwards all packets it receives on a negotiated connection, without inspecting or modifying the packet contents. A media relay only modifies the transport layer (UDP/TCP) and IP headers.

A media-relay B2BUA follows rule 1 mentioned in Section 3 and forwards the certificate fingerprint and SDP setup attribute it receives from one endpoint unmodified towards the other endpoint and vice versa. The following example shows a SIP call establishment flow, with both SIP endpoints (user agents) using DTLS-SRTP, and a media-relay B2BUA.

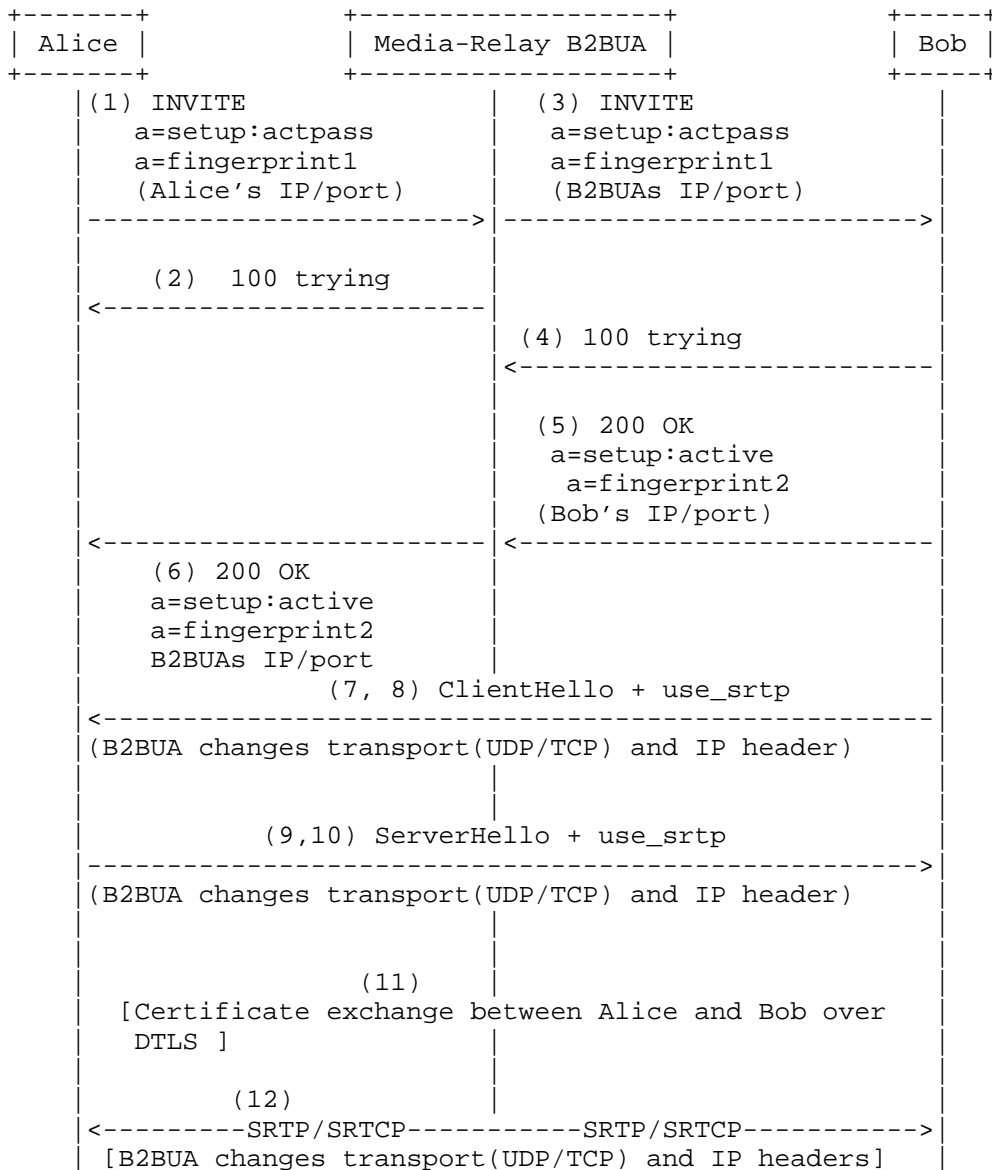


Figure 1: INVITE with SDP Call Flow for Media-Relay B2BUA

Note: For brevity, the entire value of the SDP fingerprint attribute is not shown. The example here shows only one DTLS connection for the sake of simplicity. In reality, depending on whether the RTP and RTCP flows are multiplexed or demultiplexed, there will be one or two DTLS connections.

If RTP and RTCP traffic is multiplexed on a single port as described in [RFC5761], then only a single DTLS connection is required between the peers. If RTP and RTCP are not multiplexed, then the peers would have to establish two DTLS connections. In this case, after receiving an INVITE request, Bob triggers the establishment of a DTLS connection. Note that the DTLS handshake and the sending of the INVITE response can happen in parallel; thus, the B2BUA has to be prepared to receive DTLS, Session Traversal Utilities for NAT (STUN), and media on the ports it advertised to Bob in the SDP offer before it receives an SDP answer from Bob. Since a media-relay B2BUA does not differentiate between a DTLS message, RTP, or any packet it receives, it only changes the transport layer (UDP/TCP) and IP headers and forwards the packet towards the other endpoint. The B2BUA cannot decrypt the RTP payload, as the payload is encrypted using the SRTP keys derived from the DTLS connection setup between Alice and Bob.

If the endpoints use [RFC4474], a B2BUA cannot function as a media-relay without violating rule 2 in Section 3. If [SIP-ID] is used, a B2BUA can modify the IP address in the c= line and the port in the m= line in the SDP as long as it does not otherwise violate rule 3 in Section 3.

5.1.2. RTP- and RTCP-Aware Media-Aware B2BUA

Unlike the media relay discussed in Section 5.1.1, a media-aware relay as defined in Section 3.2.2 of [RFC7092] is aware of the type of media traffic it is receiving. There are two types of media-aware relays, those that merely inspect the RTP headers and unencrypted portions of RTCP packets, and those that inspect and modify the RTP headers and unencrypted portions of RTCP packets.

5.1.2.1. RTP Header and RTCP Packets Inspection

An RTP-/RTCP-aware media relay does not modify the RTP headers and RTCP packets but only inspects the packets. Such B2BUAs follow rule 4 in Section 3 and can continue to do their function while allowing end-to-end DTLS-SRTP. Inspection by the B2BUA will not reveal the clear-text for encrypted parts of the SRTP/SRTCP packets.

5.1.2.2. RTP Header and RTCP Packet Modification

A B2BUA cannot modify RTP headers or RTCP packets, as to do so it would need to act as a DTLS endpoint, terminate the DTLS-SRTP session, and decrypt/re-encrypt RTP packets. If a B2BUA modifies unencrypted or encrypted portions of the RTP or RTCP packets, then the integrity check will fail and the packet will be dropped by the endpoint. The unencrypted and encrypted portions of the RTP or RTCP

packets are integrity protected using the HMAC algorithm negotiated during the DTLS handshake (discussed in Section 4.1.2 of [RFC5764]). B2BUAs have to follow the rules in Section 3 to avoid breaking the integrity of SRTP/SRTCP streams.

6. Forking Considerations

Due to forking [RFC3261], a SIP request carrying an SDP offer sent by an endpoint (offerer) can reach multiple remote endpoints. As a result, multiple DTLS-SRTP sessions can be established, one between the endpoint that sent the SIP request and each of the remote endpoints that received the request. B2BUAs have to follow rule 1 in Section 3 while handling offer/answer and forward the certificate fingerprints and SDP setup attributes it received in the SDP answer from each endpoint (answerer) unmodified towards the offerer. Since each DTLS connection is set up on a unique 5-tuple, B2BUA replaces the answerer's transport addresses in each answer with its unique transport addresses so that the offerer can establish a DTLS connection with each answerer. The B2BUA, acting as a media relay here, follows rule 4 mentioned in Section 3.




Figure 2: B2BUA Handling Multiple Answers

For instance, as shown in Figure 2, Alice sends a request with an offer and the request is forked. Alice receives answers from both Bob and Charlie. The B2BUA advertises different B2BUA transport addresses in each answer, as shown in Figure 2, where XXX and YYY represent different DTLS-SRTP sessions. The B2BUA replaces Bob's transport address (192.0.2.1:6666) in the answer with its transport address (192.0.2.3:7777) and Charlie's transport address

(192.0.2.2:6666) in the answer with its transport address (192.0.2.3:8888). The B2BUA tracks the remote sources (Bob and Charlie) and associates them to the local sources that are used to send packets to Alice.

7. Security Considerations

This document describes the behavior B2BUAs must follow to avoid breaking end-to-end DTLS-SRTP. Media relays that modify RTP or RTCP, or modify SIP header fields or SDP fields that are protected by the identity signature, are incompatible with end-to-end DTLS-SRTP. Such relays are out of scope for this document. Security considerations discussed in [RFC5763] are also applicable to this document. In addition, the B2BUA behaviors outlined in this document do not impact the security and integrity of a DTLS-SRTP session or the data exchanged over it. A malicious B2BUA can try to break into the DTLS connection, but such an attack can be prevented using the identity validation mechanism discussed in [RFC4474] or [SIP-ID]. Either the endpoints or the authentication service proxies involved in the call can use the identity validation mechanisms discussed in [RFC4474] or [SIP-ID] to validate the identity of peers and detect malicious B2BUAs that can attempt to terminate the DTLS connection to decrypt the RTP payload.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <<http://www.rfc-editor.org/info/rfc3711>>.
- [RFC5763] Fischl, J., Tschofenig, H., and E. Rescorla, "Framework for Establishing a Secure Real-time Transport Protocol (SRTP) Security Context Using Datagram Transport Layer Security (DTLS)", RFC 5763, DOI 10.17487/RFC5763, May 2010, <<http://www.rfc-editor.org/info/rfc5763>>.
- [RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, DOI 10.17487/RFC5764, May 2010, <<http://www.rfc-editor.org/info/rfc5764>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<http://www.rfc-editor.org/info/rfc6347>>.

8.2. 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, DOI 10.17487/RFC3261, June 2002, <<http://www.rfc-editor.org/info/rfc3261>>.
- [RFC4474] Peterson, J. and C. Jennings, "Enhancements for Authenticated Identity Management in the Session Initiation Protocol (SIP)", RFC 4474, DOI 10.17487/RFC4474, August 2006, <<http://www.rfc-editor.org/info/rfc4474>>.

- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <<http://www.rfc-editor.org/info/rfc4566>>.
- [RFC5761] Perkins, C. and M. Westerlund, "Multiplexing RTP Data and Control Packets on a Single Port", RFC 5761, DOI 10.17487/RFC5761, April 2010, <<http://www.rfc-editor.org/info/rfc5761>>.
- [RFC7092] Kaplan, H. and V. Pascual, "A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents", RFC 7092, DOI 10.17487/RFC7092, December 2013, <<http://www.rfc-editor.org/info/rfc7092>>.
- [RFC7362] Ivov, E., Kaplan, H., and D. Wing, "Latching: Hosted NAT Traversal (HNT) for Media in Real-Time Communication", RFC 7362, DOI 10.17487/RFC7362, September 2014, <<http://www.rfc-editor.org/info/rfc7362>>.
- [RFC7656] Lennox, J., Gross, K., Nandakumar, S., Salgueiro, G., and B. Burman, Ed., "A Taxonomy of Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources", RFC 7656, DOI 10.17487/RFC7656, November 2015, <<http://www.rfc-editor.org/info/rfc7656>>.
- [SIP-ID] Peterson, J., Jennings, C., Rescorla, E., and C. Wendt, "Authenticated Identity Management in the Session Initiation Protocol (SIP)", Work in Progress, draft-ietf-stir-rfc4474bis-09, May 2016

Acknowledgments

Special thanks to Lorenzo Miniero, Ranjit Avarsala, Hadriel Kaplan, Muthu Arul Mozhi, Paul Kyzivat, Peter Dawes, Brett Tate, Dan Wing, Charles Eckel, Simon Perreault, Albrecht Schwarz, Jens Guballa, Christer Holmberg, Colin Perkins, Ben Campbell, and Alissa Cooper for their constructive comments, suggestions, and early reviews that were critical to the formulation and refinement of this document. The authors would also like to thank Dan Romascanu, Vijay K. Gurbani, Francis Dupont, Paul Wouters, and Stephen Farrell for their review and feedback of this document.

Contributors

Rajeev Seth provided substantial contributions to this document.

Authors' Addresses

Ram Mohan Ravindranath
Cisco
Cessna Business Park
Sarjapur-Marathahalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: rmohanr@cisco.com

Tirumaleswar Reddy
Cisco
Cessna Business Park
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tiredy@cisco.com

Gonzalo Salgueiro
Cisco Systems, Inc.
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States

Email: gsalguei@cisco.com

Victor Pascual
Oracle
Barcelona, Spain

Email: victor.pascual.avila@oracle.com

Parthasarathi Ravindran
Nokia Networks
Bangalore, Karnataka
India

Email: partha@parthasarathi.co.in