

STRAW
Internet-Draft
Intended status: Standards Track
Expires: November 20, 2015

R. Ravindranath
T. Reddy
G. Salgueiro
Cisco
V. Pascual
Quobis
Parthasarathi. Ravindran
Nokia Networks
May 19, 2015

**DTLS-SRTP Handling in Session Initiation Protocol (SIP) Back-to-Back
User Agents (B2BUAs)
draft-ietf-straw-b2bua-dtls-srtp-01**

Abstract

Session Initiation Protocol (SIP) Back-to-Back User Agents (B2BUAs) often function on the media plane, rather than just on the signaling path. This document describes the behavior B2BUAs should follow when acting on the media plane that use Secure Real-time Transport (SRTP) security context setup with Datagram Transport Layer Security (DTLS) protocol.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 20, 2015.

Copyright Notice

Copyright (c) 2015 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 | 2 |
| 1.1. | Overview | 2 |
| 1.2. | Goals | 3 |
| 2. | Terminology | 3 |
| 3. | Media Plane B2BUAs | 4 |
| 3.1. | Media Relay | 4 |
| 3.2. | Media Aware Relay | 6 |
| 3.2.1. | RTP and RTCP Header Inspection | 6 |
| 3.2.2. | RTP and RTCP Header Modification | 6 |
| 3.3. | Media Plane B2BUA with NAT handling | 7 |
| 4. | Forking | 7 |
| 5. | Security Considerations | 7 |
| 6. | IANA Considerations | 7 |
| 7. | Acknowledgments | 7 |
| 8. | Contributors | 7 |
| 9. | References | 8 |
| 9.1. | Normative References | 8 |
| 9.2. | Informative References | 8 |
| | Authors' Addresses | 9 |

[1.](#) Introduction

[1.1.](#) Overview

[RFC5763] describes how Session Initiation Protocol (SIP) [[RFC3261](#)] can be used to establish a Secure Real-time Transport Protocol (SRTP) [[RFC3711](#)] security context with Datagram Transport Layer Security (DTLS) [[RFC6347](#)] protocol. It describes a mechanism of transporting a certificate fingerprint in the Session Description Protocol (SDP) [[RFC4566](#)], which identifies the certificate that will be presented during the DTLS handshake. DTLS-SRTP is 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 association, 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 entities exist in the SIP signaling path between the originating and final terminating endpoints. These SIP entities, as described in [[RFC7092](#)], modify SIP and SDP bodies and also are likely to be on the media path. Such entities, when present in the signaling/media path, are likely to do several things. For example, some B2BUAs modify parts of the SDP body (like IP address, port) and subsequently modify the RTP headers as well.

[1.2.](#) Goals

[RFC7092] describes two different categories of such B2BUAs, according to the level of activities performed on the media plane:

A B2BUA that act as a simple media relay effectively unaware of anything that is transported and only modifies the UDP/IP header of the packets.

A B2BUA that performs a media-aware role. It inspects and potentially modifies RTP or RTP Control Protocol (RTCP) headers; but it does not modify the payload of RTP/RTCP.

The following sections describe the behaviour B2BUAs should follow in order to avoid any impact on end-to-end DTLS-SRTP streams.

[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](#)].

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 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 is based on [[RFC7092](#)].

It is assumed the reader is already familiar with the fundamental concepts of the RTP protocol [[RFC3550](#)] and its taxonomy [[I-D.ietf-avtext-rtp-grouping-taxonomy](#)], as well as those of SRTP [[RFC3711](#)], and DTLS [[RFC6347](#)].

3. Media Plane B2BUAs

3.1. Media Relay

A media relay, as defined in [section 3.2.1 of \[RFC7092\]](#), from an application layer point-of-view, forwards all packets it receives on a negotiated UDP connection, without inspecting or modifying them. It forwards the UDP payload as-is changing only the UDP/IP header.

A media relay B2BUA MUST forward the certificate fingerprint and setup attribute it receives in the SDP from the originating endpoint as-is to the remote side and vice-versa. The example below shows an "INVITE with SDP" SIP call flow, with both SIP user agents doing DTLS-SRTP and a media relay B2BUA that changes only the IP address/port.

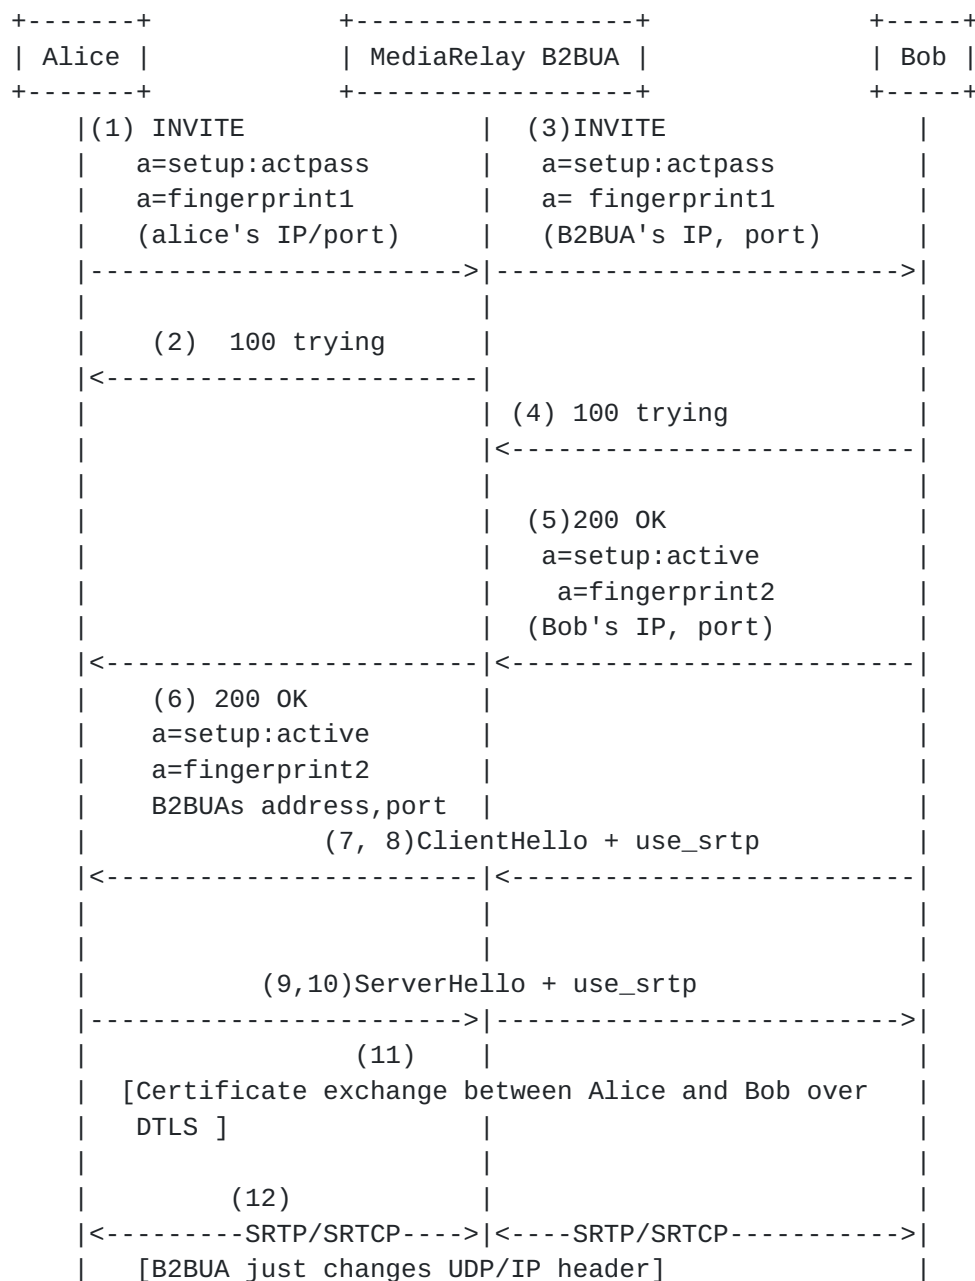


Figure 1: INVITE with SDP callflow for Media Relay B2BUA

NOTE: For brevity the entire fingerprint attribute is not shown.

For each RTP or RTCP flow, the peers do a DTLS handshake on the same source and destination port pair to establish a DTLS association. In this case, Bob, after he receives an INVITE, triggers a DTLS connection. Note the DTLS handshake and the response to the INVITE may happen in parallel; thus, the B2BUA SHOULD be prepared to receive media on the ports it advertised to Bob in the OFFER. Since a media relay B2BUA does not differentiate between a DTLS, RTP or any packet

sent it receives, it just changes the UDP/IP addresses and forwards the packet on either leg.

[I-D.ietf-stir-rfc4474bis] provides a means for signing portions of SIP requests in order to provide identity assurance and certificate pinning by providing a signature over the fingerprint of keying material in SDP for DTLS-SRTP [RFC5763]. A media relay B2BUA MUST ensure that it does not modify any of the headers used to construct the signature.

In the above example Alice may be authorized by the authorization server (SIP proxy) in its domain using the procedures in section 5 of [I-D.ietf-stir-rfc4474bis]. In such a case, if B2BUA changes some of the SIP headers or SDP content that was used by Alice's authorization server to generate the identity, it would break the identity verification procedure explained in section 4.2 of [I-D.ietf-stir-rfc4474bis] resulting in a 438 error response being returned.

3.2. Media Aware Relay

A media-aware relay, unlike the media relay discussed in the previous section, is actually aware of the media traffic it is handling. A media-aware relay inspects SRTP and SRTCP packets flowing through it, and may or may not modify the headers of the packets before forwarding them.

3.2.1. RTP and RTCP Header Inspection

B2BUAs explained in [Section 3.2.2 of \[RFC7092\]](#) do not modify the RTP and RTCP headers but only inspect the headers. Such B2BUA MUST NOT terminate the DTLS-SRTP session.

3.2.2. RTP and RTCP Header Modification

In addition to inspecting the RTP and RTCP headers, the B2BUAs explained in [section 3.2.2 \[RFC7092\]](#), can also potentially modify them. To modify media headers a B2BUA needs to act as a DTLS intermediary and terminate the DTLS connection so it can decrypt/re-encrypt RTP packets. This breaks end-to-end security. This security and privacy problem can be addressed by having separate keys for encrypting the RTP header and media payload as discussed in [I-D.jones-avtcore-private-media-reqts], in which case the B2BUA is not aware of the keys used to decrypt the media payload.

3.3. Media Plane B2BUA with NAT handling

DTLS-SRTP handshakes and offer/answer can happen in parallel. If a UA is behind a NAT and acting as a DTLS server, the ClientHello message from a B2BUA(DTLS client) is likely to be lost, as described in [section 7.3 of \[RFC5763\]](#). In order to overcome this problem, a UA and B2BUA must support ICE as discussed in [section 7.3 of \[RFC5763\]](#). If ICE check is successful then UA will receive ClientHello packet from B2BUA.

4. Forking

In SIP, it is possible that a request can get forked and multiple answers might be received for that request. So a single endpoint may end up negotiating multiple DTLS-SRTP sessions due to forking. B2BUA in both media relay and media aware relay modes MUST forward the certificate fingerprints and setup attributes it receives from each answerer as-is to the offerer.

5. Security Considerations

This document describes the behavior media plane B2BUAs (media-aware and media-unaware) should follow when acting on the media plane that uses SRTP security context setup with the DTLS protocol. It does not introduce any specific security considerations beyond those detailed in [\[RFC5763\]](#). The B2BUA behaviors outlined here also do not impact the security and integrity of the DTLS-SRTP session nor the data exchanged over it. A malicious B2BUA can try to break into the DTLS session, but such an attack can be prevented using the identity validation mechanism discussed in [\[I-D.ietf-stir-rfc4474bis\]](#).

6. IANA Considerations

This document makes no request of IANA.

7. Acknowledgments

Special thanks to Lorenzo Miniero, Ranjit Avarsala, Hadriel Kaplan, Muthu Arul Mozhi, Paul Kyzivat, Peter Dawes, Brett Tate, Dan Wing, Charles Eckel and Simon Perreault for their constructive comments, suggestions, and early reviews that were critical to the formulation and refinement of this document.

8. Contributors

Rajeev Seth provided substantial contributions to this document.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.
- [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](#), May 2010.
- [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](#), May 2010.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", [RFC 6347](#), January 2012.

9.2. Informative References

- [I-D.ietf-avtext-rtp-grouping-taxonomy]
Lennox, J., Gross, K., Nandakumar, S., Salgueiro, G., and B. Burman, "A Taxonomy of Grouping Semantics and Mechanisms for Real-Time Transport Protocol (RTP) Sources", [draft-ietf-avtext-rtp-grouping-taxonomy-06](#) (work in progress), March 2015.
- [I-D.ietf-stir-rfc4474bis]
Peterson, J., Jennings, C., and E. Rescorla, "Authenticated Identity Management in the Session Initiation Protocol (SIP)", [draft-ietf-stir-rfc4474bis-03](#) (work in progress), March 2015.
- [I-D.jones-avtcore-private-media-reqts]
Jones, P., Ismail, N., Benham, D., Buckles, N., Mattsson, J., Cheng, Y., and R. Barnes, "Requirements for Private Media in a Switched Conferencing Environment", [draft-jones-avtcore-private-media-reqts-01](#) (work in progress), March 2015.

- [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.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [RFC7092] Kaplan, H. and V. Pascual, "A Taxonomy of Session Initiation Protocol (SIP) Back-to-Back User Agents", [RFC 7092](#), December 2013.

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, Varthur Hobli
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
US

Email: gsalguei@cisco.com

Victor Pascual
Quobis
Spain

Email: victor.pascual@quobis.com

Parthasarathi Ravindran
Nokia Networks
Bangalore, Karnataka
India

Email: partha@parthasarathi.co.in