Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: September 22, 2016

J. Peterson
Neustar
E. Rescorla
R. Barnes
Mozilla
R. Housley
Vigilsec

March 21, 2016

Best Practices for Securing RTP Media Signaled with SIP
draft-peterson-dispatch-rtpsec-00.txt

Abstract

Although the Session Initital Protocol (SIP) includes a suite of security services that has been expanded by numerous specifications over the years, there is no single place that explains how to use SIP to establish confidential media sessions. Additionally, existing mechanisms have some feature gaps that need to be identified and resolved in order for them to address the pervasive monitoring threat model. This specification describes practices for negotiating confidential media with SIP, including both comprehensive security solutions which bind the media to SIP-layer identities as well as opportunistic security solutions.

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 September 22, 2016.

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

The Session Initiation Protocol (SIP) [RFC3261] includes a suite of security services, ranging from Digest authentication for authenticating entities with a shared secret, to TLS for transport security, to S/MIME (optional) for body security. SIP is frequently used to establish media sessions, in particular audio or audiovisual sessions, which have their own security mechanisms available, such as Secure RTP [RFC3711]. However, the practices needed to bind security at the media layer to security at the SIP layer, to provide an assurance that protection is in place all the way up the stack, rely on a great many external security mechanisms and practices, and require a central point of documentation to explain their optimal use as a best practice.

Revelations about widespread pervasive monitoring of the Internet have led to a reevaluation of the threat model for Internet communications [RFC7258]. In order to maximize the use of security features, especially of media confidentiality, opportunistic measures must often serve as a stopgap when a full suite of services cannot be negotiated all the way up the stack. This document explains the limitations that may inhibit the use of comprehensive security, and provides recommendations for which external security mechanisms

implementers should use to negotiate secure media with SIP. It moreover gives a gap analysis of the limitations of existing solutions, and specifies solutions to address them.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] and RFC 6919 [RFC6919].

3. Security at the SIP and SDP layer

There are two approaches to providing confidentiality for media sessions set up with SIP: comprehensive security and opportunistic security.

3.1. Comprehensive Security

Comprehensive security for media sessions established by SIP requires the interaction of three protocols: SIP, the Session Description Protocol (SDP), and the Real-time Protocol, in particular its secure profile SRTP. Broadly, it is the responsibility of SIP to provide integrity for the media keying attributes conveyed by SDP, and those attributes will in turn identify the keys used by endpoints in the RTP media session that SDP negotiates. In that way, once SIP and SDP have exchanged the necessary information to initiate a session, the media endpoints will have a strong assurance that the keys they exchange have not been tampered with by third parties, and that end-to-end confidentiality is available.

Our current target mechanism for establishing the identity of the endpoints of a SIP session is the use of STIR [I-D.ietf-stir-rfc4474bis]. The STIR signature has been designed to prevent a class of impersonation attacks that are commonly used in robocalling, voicemail hacking, and related threats. STIR generates a signature over certain features of SIP requests, including header field values that contain an identity for the originator of the request, such as the From header field or P-Asserted-Identity field, and also over the media keys in SDP if they are present. As currently defined, STIR only provides a signature over the "a=fingerprint" attribute, which is a key fingerprint utilized by DTLS-SRTP [RFC5763]; consequently, STIR only offers comprehensive security for SIP sessions, in concert with SDP and SRTP, when DTLS-SRTP is the media security service. The underlying security object of STIR is extensible, however, and it would be possible to provide signatures over other SDP attributes that contain alternate keying material.

A STIR verification service can act in concept with an SRTP media endpoint to ensure that the key fingerprints, as given in SDP, match the keys exchanged to establish DTLS-SRTP. Typically, the verification service function would in this case be implemented in the SIP UAS, which would be composed with the media endpoint. If the STIR authentication service or verification service functions are implemented at an intermediary rather than an endpoint, this introduces the possibility that the intermediary could act as a man-in-the-middle, altering key fingerprints. As this attack is not in STIR's core threat model, which focuses on impersonation rather than man-in-the-middle attacks, STIR offers no specific protections against it. However, it would be possible to build a deployment profile of STIR for media confidentiality which shifts these responsibilities to the endpoints rather than the intermediaries.

Note that STIR provides integrity protection for the SDP bodies of SIP requests, but not SIP responses. When a session is established, therefore, any SDP body carried by a 200 class response in the backwards direction will not be protected by an authentication service and cannot be verified. Thus, sending a secured SDP body in the backwards direction will require an extra RTT, typically a re-INVITE in the backwards direction. Again, this could be specified as a component of a secure media profile for STIR.

Future versions of this specification will show in detail how those gaps can be filled.

3.1.1. Anonymous Communications

In some cases, the identity of the initiator of a SIP session may be withheld due to user or provider policy. Per the recommendations of [RFC3323], this may involve using an identity such as "anonymous@anonymous.invalid" in the identity fields of a SIP request. [I-D.ietf-stir-rfc4474bis] does not currently permit authentication services to sign for requests that supply this identity. It does however permit signing for valid domains, such as "anonymous@example.com," as a way of implementation an anonymization service as specified in [RFC3323].

Even for anonymous sessions, providing media confidentiality and partial SDP integrity is still desirable. Barring the use of an anonymization service, this can only be accomplished with opportunistic security; the value of trying to provide an intermediate level between comprehensive and opportunistic security for this use case is a matter for futher discussion and study.

3.2. Opportunistic Security

Work is already underway on defining approaches to opportunistic media security for SIP in [I-D.johnston-dispatch-osrtp], which builds on the prior efforts of [I-D.kaplan-mmusic-best-effort-srtp]. The major protocol change proposed by that draft is to signal the use of opportunistic encryption by negotiating the AVP profile in SDP, rather than the SAVP profile (as specified in [RFC3711]) that would ordinarily be used when negotiating SRTP.

Opportunistic encryption approaches typically have no integrity protection for the keying material in SDP. Sending SIP over TLS hop-by-hop between user agents and any intermediaries will reduce the prospect that active attackers can alter keys for session requests on the wire.

4. Media Security

As there are several ways to negotiate media security with SDP, any of which might be used with either opportunistic or comprehensive security, further guidance to implementers is needed. In [I-D.johnston-dispatch-osrtp], opportunistic approaches considered include DTLS-SRTP, security descriptions [RFC4568], and ZRTP [RFC6189]. In order to prevent men-in-the-middle from decrypting media traffic, the "a=crypto" SDP parameter of security descriptions requires signaling confidentiality which STIR and related comprehensive security approaches cannot provide, so delivering keys by value in SDP in this fashion is NOT RECOMMENDED. Both DTLS-SRTP and ZRTP instead provide hashes which are carried in SDP, and thus require only integrity protection rather than confidentiality.

Of DTLS-SRTP and ZRTP, only DTLS-SRTP is a Standards Track Internet protocol. Future versions of this specification will give specific recommendations on support for media security protocols.

Future versions of this specification will explore the issue of multiple fingerprints appearing in the message, and offers that include both DTLS-SRTP and ZRTP security.

5. Acknowledgments

We would like to thank YOU for contributions to this problem statement and framework.

6. IANA Considerations

This memo includes no requests to the IANA.

7. Security Considerations

This document describes the security features that provide media sessions established with SIP with confidentiality, integrity, and authentication.

8. Informative References

[I-D.ietf-stir-rfc4474bis]
Peterson, J., Jennings, C., Rescorla, E., and C. Wendt, "Authenticated Identity Management in the Session Initiation Protocol (SIP)", draft-ietf-stir-rfc4474bis-07 (work in progress), February 2016.

[I-D.johnston-dispatch-osrtp]
Johnston, A., Aboba, B., Hutton, A., Liess, L., and T. Thomas, "An Opportunistic Approach for Secure Real-time Transport Protocol (OSRTP)", draft-johnston-dispatch-osrtp-02 (work in progress), February 2016.

[I-D.kaplan-mmusic-best-effort-srtp]
Audet, F. and H. Kaplan, "Session Description Protocol (SDP) Offer/Answer Negotiation For Best-Effort Secure Real-Time Transport Protocol", draft-kaplan-mmusic-best-effort-srtp-01 (work in progress), October 2006.

[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>.

[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>.

[RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, DOI 10.17487/RFC3264, June 2002, <http://www.rfc-editor.org/info/rfc3264>.

[RFC3323] Peterson, J., "A Privacy Mechanism for the Session Initiation Protocol (SIP)", RFC 3323, DOI 10.17487/RFC3323, November 2002, <http://www.rfc-editor.org/info/rfc3323>.

[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>.

[RFC4568] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, DOI 10.17487/RFC4568, July 2006, <http://www.rfc-editor.org/info/rfc4568>.

[RFC5124] Ott, J. and E. Carrara, "Extended Secure RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/SAVPF)", RFC 5124, DOI 10.17487/RFC5124, February 2008, <http://www.rfc-editor.org/info/rfc5124>.

[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>.

[RFC6189] Zimmermann, P., Johnston, A., Ed., and J. Callas, "ZRTP: Media Path Key Agreement for Unicast Secure RTP", RFC 6189, DOI 10.17487/RFC6189, April 2011, <http://www.rfc-editor.org/info/rfc6189>.

[RFC6919] Barnes, R., Kent, S., and E. Rescorla, "Further Key Words for Use in RFCs to Indicate Requirement Levels", RFC 6919, DOI 10.17487/RFC6919, April 2013, <http://www.rfc-editor.org/info/rfc6919>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, DOI 10.17487/RFC7258, May 2014, <http://www.rfc-editor.org/info/rfc7258>.

Authors' Addresses

Jon Peterson
Neustar, Inc.
1800 Sutter St Suite 570
Concord, CA 94520
US

Email: jon.peterson@neustar.biz

Eric Rescorla
Mozilla

Email: ekr@rtfm.com

Richard Barnes
Mozilla

Email: rbarnes@mozilla.com

Russ Housley
Vigilsec

Email: rhousley@vigilsec.com