

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: May 3, 2018

J. Peterson
Neustar
E. Rescorla
R. Barnes
Mozilla
R. Housley
Vigilsec
October 30, 2017

Best Practices for Securing RTP Media Signaled with SIP
draft-ietf-sipbrandy-rtpsec-03.txt

Abstract

Although the Session Initiation 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 best practices for negotiating confidential media with SIP, including both comprehensive protection 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 <https://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 May 3, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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
2.	Terminology	3
3.	Security at the SIP and SDP layer	3
3.1.	Comprehensive Protection	3
3.2.	Opportunistic Security	4
4.	STIR Profile for Endpoint Authentication and Verification Services	4
4.1.	Credentials	5
4.2.	Anonymous Communications	6
4.3.	Connected Identity Usage	7
4.4.	Authorization Decisions	8
5.	Media Security Protocols	8
6.	Relayed Media and Conferencing	9
7.	ICE and Connected Identity	9
8.	Best Current Practices	10
9.	Acknowledgments	10
10.	IANA Considerations	11
11.	Security Considerations	11
12.	Informative References	11
	Authors' Addresses	14

[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 (optionally) 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 protection, 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.

Various specifications that user agents must implement to support media confidentiality are given in the sections below; a summary of the best current practices appears in [Section 8](#).

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 protection and opportunistic security (as defined in [[RFC7435](#)]).

3.1. Comprehensive Protection

Comprehensive protection 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 protection 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(s) that SDP negotiates. In that way, once SIP and SDP have exchanged the necessary information to initiate a session, 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 Identity header 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 protection for SIP sessions, in concert with SDP and SRTP, when DTLS-SRTP is the media security service. The underlying PASSport [I-D.ietf-stir-passport] object used by STIR is extensible, however, and it would be possible to provide signatures over other SDP attributes that contain alternate keying material. A profile for using STIR to provide media confidentiality is given in [Section 4](#).

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. However, opportunistic confidentiality for media will prevent passive attacks of the form most common in the threat of pervasive monitoring.

4. STIR Profile for Endpoint Authentication and Verification Services

STIR [I-D.ietf-stir-rfc4474bis] defines the Identity header field for SIP, which provides a cryptographic attestation of the source of communications. This profile assumes that a STIR verification service will act in concert with an SRTP media endpoint to ensure that the key fingerprints, as given in SDP, match the keys exchanged to establish DTLS-SRTP. To satisfy this condition, 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 such interference. The SIPBRANDY deployment profile of STIR for media confidentiality thus shifts these responsibilities to the endpoints rather than the intermediaries.

In order to be compliant with best practices for SIP media confidentiality with comprehensive protection, user agent implementations MUST implement both the authentication service and verification service roles described in [[I-D.ietf-stir-rfc4474bis](#)]. STIR authentication services MUST signal their compliance with this specification by adding the "msec" header element defined in this specification to the PASSport header. Implementations MUST provide key fingerprints in SDP and the appropriate signatures over them per [[I-D.ietf-stir-passport](#)].

When generating either an offer or an answer, compliant implementations MUST include an "a=fingerprint" attribute containing the fingerprint of an appropriate key (see [Section 4.1](#)).

[4.1](#). Credentials

In order to implement the authentication service function in the user agent, SIP endpoints must acquire the credentials needed to sign for their own identity. That identity is typically carried in the From header field of a SIP request, and either contains a greenfield SIP URI (e.g. "sip:alice@example.com") or a telephone number, which can appear in a variety of ways (e.g. "sip:+17004561212@example.com;user=phone").

[[I-D.ietf-stir-rfc4474bis](#)] [Section 8](#) contains guidance for separating the two, and determining what sort of credential is needed to sign for each.

To date, few commercial certificate authorities issue certificates for SIP URIs or telephone numbers; though work is ongoing on systems for this purpose (such as [[I-D.peterson-acme-telephone](#)]) it is not mature enough to be recommended as a best practice. This is one reason why the STIR standard is architected to permit intermediaries to act as an authentication service on behalf of an entire domain, just as in SIP a proxy server can provide domain-level SIP service. While certificate authorities that offered proof-of-possession certificates similar to those used in the email world could be offered for SIP, either for greenfield identifiers or for telephone numbers, this specification does not require their use.

For users who do not possess such certificates, DTLS-SRTP [[RFC5763](#)] permits the use of self-signed keys. This profile of STIR for media confidentiality therefore relaxes the authority requirements of [[I-D.ietf-stir-rfc4474bis](#)] to allow the use of self-signed keys for authentication services that are composed with user agents, by

generating a certificate (per the guidance of [\[I-D.ietf-stir-certificates\]](#)) with a subject corresponding to the user's identity. Such a credential could be used for trust on first use (see [\[RFC7435\]](#)) by relying parties. Note that relying parties SHOULD NOT use certificate revocation mechanisms or real-time certificate verification systems for self-signed certificates as they will not increase confidence in the certificate.

Users who wish to remain anonymous can instead generate self-signed certificates as described in [Section 4.2](#).

Generally speaking, without access to out-of-band information about which certificates were issued to whom, it will be very difficult for relying parties to ascertain whether or not the signer of a SIP request is genuinely an "endpoint." Even the term "endpoint" is a problematic one, as SIP user agents can be composed in a variety of architectures and may not be devices under direct user control. While it is possible that techniques based on certificate transparency [\[RFC6962\]](#) or similar practices could help user agents to recognize one another's certificates, those operational systems will need to ramp up with the certificate authorities that issue credentials to end user devices going forward.

[4.2](#). 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 implementing an anonymization service as specified in [\[RFC3323\]](#).

Even for anonymous sessions, providing media confidentiality and partial SDP integrity is still desirable. This specification RECOMMENDS using one-time self-signed certificates for anonymous communications, with a subjectAltName of "sip:anonymous@anonymous.invalid". After a session is terminated, the certificate should be discarded, and a new one, with new keying material, should be generated before each future anonymous call. As with self-signed certificates, relying parties SHOULD NOT use certificate revocation mechanisms or real-time certificate verification systems for anonymous certificates as they will not increase confidence in the certificate.

Note that when using one-time anonymous self-signed certificates, any man in the middle could strip the Identity header and replace it with one signed by its own one-time certificate, changing the "mkey" parameters of PASSporT and any "a=fingerprint" attributes in SDP as it chooses. This signature only provides protection against non-Identity aware entities that might modify SDP without altering the PASSporT conveyed in the Identity header.

4.3. Connected Identity Usage

STIR [[I-D.ietf-stir-rfc4474bis](#)] 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 request sent in the backwards direction.

The problem of providing "Connected Identity" for the original [RFC4474](#) was explored in [[RFC4916](#)], which uses a provisional or mid-dialog UPDATE request in the backwards direction to convey an Identity header for the recipient of an INVITE. The procedures in that specification are largely compatible with the revision of the Identity header in [[I-D.ietf-stir-rfc4474bis](#)]. However, the following updates to [[RFC4916](#)] are required:

The UPDATE carrying signed SDP with a fingerprint in the backwards direction MUST be sent during dialog establishment, following the receipt of a PRACK after a provisional 1xx response.

For use with this STIR Profile for media confidentiality, the UAS that responds to the INVITE request MUST act as an authentication service for the UPDATE sent in the backwards direction.

The text in [RFC4916 Section 4.4.1](#) regarding the receipt at a UAC of error codes 428, 436, 437 and 438 in response to a mid-dialog request RECOMMENDS treating the dialog as terminated. [[I-D.ietf-stir-rfc4474bis](#)] allows the retransmission of requests with repairable error conditions (see [section 6.1.1](#)) in a way that can override that SHOULD in [RFC4916](#). In particular, an authentication service MAY retry a mid-dialog as [[I-D.ietf-stir-rfc4474bis](#)] allows rather than treating the dialog as terminated, though note that only one such retry is permitted.

The examples in [RFC4916](#) are based on the original [RFC4474](#), and will not match signatures using [[I-D.ietf-stir-rfc4474bis](#)].

Future work may be done to revise [RFC4916](#) for STIR; that work should take into account any impacts on the profile described in this document. The use of [RFC4916](#) has some further interactions with ICE; see [Section 7](#).

4.4. Authorization Decisions

[I-D.ietf-stir-rfc4474bis] grants STIR verification services a great deal of latitude when making authorization decisions based on the presence of the Identity header field. It is largely a matter of local policy whether an endpoint rejects a call based on absence of an Identity header field, or even the presence of a header that fails an integrity check against the request.

For this profile, however, a compliant verification service that receives a dialog-forming SIP request containing an Identity header with a PASSport type of "msec", after validating the request per the steps described in [I-D.ietf-stir-rfc4474bis] [Section 6.2](#), MUST reject the request if there is any failure in that validation process with the appropriate status code per [Section 6.2.2](#). If the request is valid, then if a terminating user accepts the request, it MUST then follow the steps in [Section 4.3](#) to act as an authentication service and send a signed request with the "msec" PASSport type in its Identity header as well, in order to enable end-to-end bidirectional confidentiality.

For the purposes of this profile, the "msec" PASSport type can be used by authentication services in one of two ways: as a mandatory request for media security, or as a merely opportunistic request for media security. As any verification service that receives an Identity header in a SIP request with an unrecognized PASSport type will simply ignore that Identity header, an authentication service will know whether or not the terminating side supports "msec" based on whether or not its user agent receives a signed request in the backwards direction per [Section 4.3](#). If no such requests are received, the UA may do one or two things: shut down the dialog, if the policy of the UA requires that "msec" be supported by the terminating side for this dialog; or, if policy permits, allow the dialog to continue without media security.

5. Media Security Protocols

As there are several ways to negotiate media security with SDP, any of which might be used with either opportunistic or comprehensive protection, further guidance to implementers is needed. In [I-D.johnston-dispatch-osrtp], opportunistic approaches considered include DTLS-SRTP, security descriptions [[RFC4568](#)], and ZRTP [[RFC6189](#)].

DTLS-SRTP is the only Standards Track Internet protocol for media security. For that reason, this specification **REQUIRES** support for DTLS-SRTP.

The "mkey" claim of PASSport provides integrity protection for "a=fingerprint" attributes in SDP, including cases where multiple "a=fingerprint" attributes appear in the same SDP.

6. Relayed Media and Conferencing

Providing end-to-end media confidentiality for SIP is complicated by the presence of many forms of media relays. While many media relays merely proxy media to a destination, others present themselves as media endpoints and terminate security associations before re-originating media to its destination.

Centralized conference bridges are one type of entity that typically terminates a media session in order to mux media from multiple sources and then to re-originate the muxed media to conference participants. In many such implementations, only hop-by-hop media confidentiality is possible. Work is ongoing to specify a means to encrypt both the hop-by-hop media between a user agent and a centralized server as well as the end-to-end media between user agents, but is not sufficiently mature at this time to make a recommendation for a best practice here. Those protocols are expected to identify their own best practice recommendations as they mature.

Another class of entities that might relay SIP media are back-to-back user agents (B2BUAs). If a B2BUA follows the guidance in [\[RFC7879\]](#), it may be possible for those devices to act as media relays while still permitting end-to-end confidentiality between user agents.

Ultimately, if an endpoint can decrypt media it receives, then that endpoint can forward the decrypted media without the knowledge or consent of the media's originator. No media confidentiality mechanism can protect against these sorts of relayed disclosures, or trusted entities that can decrypt media and then record a copy to be sent elsewhere (see [\[RFC7245\]](#)).

7. ICE and Connected Identity

Providing confidentiality for media with comprehensive protection requires careful timing of when media streams should be sent and when a user interface should signify that confidentiality is in place.

In order to best enable end-to-end connectivity between user agents, and to avoid media relays as much as possible, implementations of

this specification must support ICE [[I-D.ietf-ice-rfc5245bis](#)]. To speed up call establishment, it is RECOMMENDED that implementations support trickle ICE [[I-D.ietf-mmusic-trickle-ice-sip](#)].

Note that in the comprehensive protection case, the use of Connected Identity [[RFC4916](#)] with ICE entails that the answer containing the key fingerprints, and thus the STIR signature, will come in an UPDATE sent in the backwards direction a provisional response and acknowledgment (PRACK), rather than in any earlier SDP body. Only at such a time as that UPDATE is received will the media keys be considered exchanged in this case.

Similarly, in order to prevent, or at least mitigate, the denial-of-service attack envisioned in [[RFC5245](#)] [Section 18.5.1](#), this specification incorporates best practices for ensuring that recipients of media flows have consented to receive such flows. Implementations of this specification MUST implement the STUN usage for consent freshness defined in [[RFC7675](#)].

8. Best Current Practices

The following are the best practices for SIP user agents to provide media confidentiality for SIP sessions.

Implementations MUST support the STIR endpoint profile given in [Section 4](#), and signal that in PASSport with the "msec" header element.

Implementations MUST follow the authorization decision behavior in [Section 4.4](#).

Implementations MUST support DTLS-SRTP for key-management, as described in [Section 5](#).

Implementations MUST support the ICE, and the STUN consent freshness mechanism, as specified in [Section 7](#).

9. Acknowledgments

We would like to thank Adam Roach, Andrew Hutton, and Ben Campbell for contributions to this problem statement and framework.

10. IANA Considerations

This specification defines a new values for the PASSport Type registry called "msec," and the IANA is requested to add that to the registry with a value pointing to [RFCThis].

11. Security Considerations

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

12. Informative References

[I-D.ietf-ice-rfc5245bis]

Keranen, A., Holmberg, C., and J. Rosenberg, "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal", [draft-ietf-ice-rfc5245bis-14](#) (work in progress), October 2017.

[I-D.ietf-mmusic-trickle-ice-sip]

Ivov, E., Stach, T., Marocco, E., and C. Holmberg, "A Session Initiation Protocol (SIP) usage for Trickle ICE", [draft-ietf-mmusic-trickle-ice-sip-10](#) (work in progress), October 2017.

[I-D.ietf-stir-certificates]

Peterson, J. and S. Turner, "Secure Telephone Identity Credentials: Certificates", [draft-ietf-stir-certificates-14](#) (work in progress), May 2017.

[I-D.ietf-stir-passport]

Wendt, C. and J. Peterson, "Personal Assertion Token (PASSporT)", [draft-ietf-stir-passport-11](#) (work in progress), February 2017.

[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-16](#) (work in progress), February 2017.

[I-D.johnston-dispatch-osrtp]

Johnston, A., Ph.D., D., Hutton, A., Liess, L., and T. Stach, "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.

[I-D.peterson-acme-telephone]

Peterson, J. and R. Barnes, "ACME Identifiers and Challenges for Telephone Numbers", [draft-peterson-acme-telephone-00](#) (work in progress), October 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://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, <<https://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, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc4568>>.

[RFC4916] Elwell, J., "Connected Identity in the Session Initiation Protocol (SIP)", [RFC 4916](#), DOI 10.17487/RFC4916, June 2007, <<https://www.rfc-editor.org/info/rfc4916>>.

- [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, <<https://www.rfc-editor.org/info/rfc5124>>.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [RFC 5245](#), DOI 10.17487/RFC5245, April 2010, <<https://www.rfc-editor.org/info/rfc5245>>.
- [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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc6919>>.
- [RFC6962] Laurie, B., Langley, A., and E. Kasper, "Certificate Transparency", [RFC 6962](#), DOI 10.17487/RFC6962, June 2013, <<https://www.rfc-editor.org/info/rfc6962>>.
- [RFC7245] Hutton, A., Ed., Portman, L., Ed., Jain, R., and K. Rehor, "An Architecture for Media Recording Using the Session Initiation Protocol", [RFC 7245](#), DOI 10.17487/RFC7245, May 2014, <<https://www.rfc-editor.org/info/rfc7245>>.
- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.
- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", [RFC 7435](#), DOI 10.17487/RFC7435, December 2014, <<https://www.rfc-editor.org/info/rfc7435>>.
- [RFC7675] Perumal, M., Wing, D., Ravindranath, R., Reddy, T., and M. Thomson, "Session Traversal Utilities for NAT (STUN) Usage for Consent Freshness", [RFC 7675](#), DOI 10.17487/RFC7675, October 2015, <<https://www.rfc-editor.org/info/rfc7675>>.

[RFC7879] Ravindranath, R., Reddy, T., Salgueiro, G., Pascual, V., and P. Ravindran, "DTLS-SRTP Handling in SIP Back-to-Back User Agents", [RFC 7879](https://www.rfc-editor.org/info/rfc7879), DOI 10.17487/RFC7879, May 2016, <<https://www.rfc-editor.org/info/rfc7879>>.

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: rlb@ipv.sx

Russ Housley
Vigilsec

Email: housley@vigilsec.com

