

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 18, 2014

C.S. Perkins
University of Glasgow
M. Westerlund
Ericsson
October 15, 2013

Securing the RTP Protocol Framework: Why RTP Does Not Mandate a Single
Media Security Solution
draft-ietf-avt-srtp-not-mandatory-14.txt

Abstract

This memo discusses the problem of securing real-time multimedia sessions, and explains why the Real-time Transport Protocol (RTP), and the associated RTP Control Protocol (RTCP), do not mandate a single media security mechanism. Guidelines for designers and reviewers of future RTP extensions are provided, to ensure that appropriate security mechanisms are mandated, and that any such mechanisms are specified in a manner that conforms with the RTP architecture.

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 April 18, 2014.

Copyright Notice

Copyright (c) 2013 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
2. RTP Applications and Deployment Scenarios	3
3. RTP Media Security	4
4. RTP Session Establishment and Key Management	4
5. On the Requirement for Strong Security in Framework protocols	5
6. Guidelines for Securing the RTP Protocol Framework	6
7. Conclusions	7
8. Security Considerations	7
9. IANA Considerations	8
10. Acknowledgements	8
11. Informative References	8
Authors' Addresses	9

1. Introduction

The Real-time Transport Protocol (RTP) [RFC3550] is widely used for voice over IP, Internet television, video conferencing, and other real-time and streaming media applications. Despite this use, the basic RTP specification provides only limited options for media security, and defines no standard key exchange mechanism. Rather, a number of extensions are defined that can provide confidentiality and authentication of RTP media streams and RTP Control Protocol (RTCP) messages. Other mechanisms define key exchange protocols. This memo outlines why it is appropriate that multiple extension mechanisms are defined rather than mandating a single security and keying mechanism for all users of RTP.

The IETF policy on Strong Security Requirements for IETF Standard Protocols [RFC3365] (the so-called "Danvers Doctrine") states that "we MUST implement strong security in all protocols to provide for the all too frequent day when the protocol comes into widespread use in the global Internet". The security mechanisms defined for use with RTP allow these requirements to be met. However, since RTP is a protocol framework that is suitable for a wide variety of use cases, there is no single security mechanism that is suitable for every scenario. This memo outlines why this is the case, and discusses how users of RTP can meet the requirement for strong security.

This memo provides information for the community and for reviewers of future RTP-related work in the IETF. It does not specify a standard of any kind.

2. RTP Applications and Deployment Scenarios

The range of application and deployment scenarios where RTP has been used includes, but is not limited to, the following:

- o Point-to-point voice telephony;
- o Point-to-point video conferencing and telepresence;
- o Centralised group video conferencing and telepresence, using a Multipoint Conference Unit (MCU) or similar central middlebox;
- o Any Source Multicast (ASM) video conferencing using the light-weight sessions model (e.g., the Mbone conferencing tools);
- o Point-to-point streaming audio and/or video (e.g., on-demand TV or movie streaming);
- o Source-Specific Multicast (SSM) streaming to large receiver groups (e.g., IPTV streaming by residential ISPs, or the 3GPP Multimedia Broadcast Multicast Service [MBMS]);
- o Replicated unicast streaming to a group of receivers;
- o Interconnecting components in music production studios and video editing suites;
- o Interconnecting components of distributed simulation systems; and
- o Streaming real-time sensor data (e.g., e-VLBI radio astronomy).

As can be seen, these scenarios vary from point-to-point sessions to very large multicast groups, from interactive to non-interactive, and from low bandwidth (kilobits per second) telephony to high bandwidth (multiple gigabits per second) video and data streaming. While most of these applications run over UDP [RFC0768], some use TCP [RFC0793], [RFC4614] or DCCP [RFC4340] as their underlying transport. Some run on highly reliable optical networks, others use low rate unreliable wireless networks. Some applications of RTP operate entirely within a single trust domain, others run inter-domain, with untrusted (and, in some cases, potentially unknown) users. The range of scenarios is wide, and growing both in number and in heterogeneity.

3. RTP Media Security

The wide range of application scenarios where RTP is used has led to the development of multiple solutions for securing RTP media streams and RTCP control messages, considering different requirements.

Perhaps the most widely applicable of these security options is the Secure RTP (SRTP) framework [RFC3711]. This is an application-level media security solution, encrypting the media payload data (but not the RTP headers) to provide confidentiality, and supporting source origin authentication as an option. SRTP was carefully designed to be low overhead, including operating on links subject to RTP header compression, and to support the group communication and third-party performance monitoring features of RTP, across a range of networks.

SRTP is not the only media security solution for RTP, however, and alternatives can be more appropriate in some scenarios, perhaps due to ease of integration with other parts of the complete system. In addition, SRTP does not address all possible security requirements, and other solutions are needed in cases where SRTP is not suitable. For example, ISMACryp payload-level confidentiality [ISMACrypt2] is appropriate for some types of streaming video application, but is not suitable for voice telephony, and uses features that are not provided by SRTP.

The range of available RTP security options, and their applicability to different scenarios, is outlined in [I-D.ietf-avtcore-rtp-security-options]. At the time of this writing, there is no media security protocol that is appropriate for all the environments where RTP is used. Multiple RTP media security protocols are expected to remain in wide use for the foreseeable future.

4. RTP Session Establishment and Key Management

A range of different protocols for RTP session establishment and key exchange exist, matching the diverse range of use cases for the RTP framework. These mechanisms can be split into two categories: those that operate in-band on the media path, and those that are out-of-band and operate as part of the session establishment signalling channel. The requirements for these two classes of solution are different, and a wide range of solutions have been developed in this space.

A more detailed survey of requirements for media security management protocols can be found in [RFC5479]. As can be seen from that memo, the range of use cases is wide, and there is no single key management protocol that is appropriate for all scenarios. The solutions have

been further diversified by the existence of infrastructure elements, such as authentication systems, that are tied to the key management. The most important and widely used keying options for RTP sessions at the time of this writing are described in [I-D.ietf-avtcore-rtp-security-options].

5. On the Requirement for Strong Security in Framework protocols

The IETF requires that all protocols provide a strong, mandatory to implement, security solution [RFC3365]. This is essential for the overall security of the Internet, to ensure that all implementations of a protocol can interoperate in a secure way. Framework protocols offer a challenge for this mandate, however, since they are designed to be used by different classes of applications, in a wide range of different environments. The different use cases for the framework have different security requirements, and implementations designed for different environments are generally not expected to interwork.

RTP is an example of a framework protocol with wide applicability. The wide range of scenarios described in Section 2 show the issues that arise in mandating a single security mechanism for this type of framework. It would be desirable if a single media security solution, and a single key management solution, could be developed, suitable for applications across this range of use scenarios. The authors are not aware of any such solution, however, and believe it is unlikely that any such solution will be developed. In part, this is because applications in the different domains are not intended to interwork, so there is no incentive to develop a single mechanism. More importantly, though, the security requirements for the different usage scenarios vary widely, and an appropriate security mechanism in one scenario simply does not work for some other scenarios.

For a framework protocol, it appears that the only sensible solution to the strong security requirement of [RFC3365] is to develop and use building blocks for the basic security services of confidentiality, integrity protection, authorisation, authentication, and so on. When new uses for the framework protocol arise, they need to be studied to determine if the existing security building blocks can satisfy the requirements, or if new building blocks need to be developed. A mandatory to implement set of security building blocks can then be specified for that usage scenario of the framework.

Therefore, when considering the strong and mandatory to implement security mechanism for a specific class of applications, one has to consider what security building blocks need to be supported. To maximize interoperability it is important that common media security and key management mechanisms are defined for classes of application with similar requirements. The IETF needs to participate in this

selection of security building blocks for each class of applications that use the protocol framework and are expected to interoperate, in cases where the IETF has the appropriate knowledge of the class of applications.

6. Guidelines for Securing the RTP Protocol Framework

The IETF requires that protocols specify mandatory to implement (MTI) strong security [RFC3365]. This applies to the specification of each interoperable class of application that makes use of RTP. However, RTP is a framework protocol, so the arguments made in Section 5 also apply. Given the variability of the classes of application that use RTP, and the variety of the currently available security mechanisms described in [I-D.ietf-avtcore-rtp-security-options], no one set of MTI security options can realistically be specified that apply to all classes of RTP applications.

Documents that define an interoperable class of applications using RTP are subject to [RFC3365], and so need to specify MTI security mechanisms. This is because such specifications do fully specify interoperable applications that use RTP. Examples of such documents under development in the IETF at the time of this writing are the RTCWEB Security Architecture [I-D.ietf-rtcweb-security-arch] and the Real Time Streaming Protocol 2.0 (RTSP) [I-D.ietf-mmusic-rfc2326bis]. It is also expected that a similar document will be produced for voice-over-IP applications using SIP and RTP.

The RTP framework includes several extension points. Some extensions can significantly change the behaviour of the protocol, to the extent that applications using the extension form a separate interoperable class of applications to those that have not been extended. Other extension points are defined in such a manner that they can be used (largely) independently of the class of applications using RTP. Two important extension points that are independent of the class of applications are RTP Payload Formats and RTP Profiles.

An RTP Payload Format defines how the output of a media codec can be used with RTP. At the time of this writing, there are over 70 RTP Payload Formats defined in published RFCs, with more in development. It is appropriate for an RTP Payload Format to discuss the specific security implications of using that media codec with RTP. However, an RTP Payload Format does not specify an interoperable class of applications that use RTP since, in the vast majority of cases, a media codec and its associated RTP Payload Format can be used with many different classes of application. As such, an RTP Payload Format is neither secure in itself, nor something to which [RFC3365] applies. Future RTP Payload Format specifications need to explicitly state this, and include a reference to this memo for explanation. It

is not appropriate for an RTP Payload Format to mandate the use of SRTP [RFC3711], or any other security building blocks, since that RTP Payload Format might be used by different classes of application that use RTP, and that have different security requirements.

RTP Profiles are larger extensions that adapt the RTP framework for use with particular classes of application. In some cases, those classes of application might share common security requirements so that it could make sense for an RTP Profile to mandate particular security options and building blocks (the RTP/SAVP profile [RFC3711] is an example of this type of RTP Profile). In other cases, though, an RTP profile is applicable to such a wide range of applications that it would not make sense for that profile to mandate particular security building blocks be used (the RTP/AVPF profile [RFC4585] is an example of this type of RTP Profile, since it provides building blocks that can be used in different styles of application). A new RTP Profile specification needs to discuss whether, or not, it makes sense to mandate particular security building blocks that need to be used with all implementations of that profile; however, there is no expectation that all RTP Profiles will mandate particular security solutions. RTP Profiles that do not specify an interoperable usage for a particular class of RTP applications are neither secure in themselves, nor something to which [RFC3365] applies; any future RTP Profiles in this category need to explicitly state this with justification, and include a reference to this memo.

7. Conclusions

The RTP framework is used in a wide range of different scenarios, with no common security requirements. Accordingly, neither SRTP [RFC3711], nor any other single media security solution or keying mechanism, can be mandated for all uses of RTP. In the absence of a single common security solution, it is important to consider what mechanisms can be used to provide strong and interoperable security for each different scenario where RTP applications are used. This will require analysis of each class of application to determine the security requirements for the scenarios in which they are to be used, followed by the selection of a mandatory to implement security building blocks for that class of application, including the desired RTP traffic protection and key-management. A non-exhaustive list of the RTP security options available at the time of this writing is outlined in [I-D.ietf-avtcore-rtp-security-options]. It is expected that each class of application will be supported by a memo describing what security options are mandatory to implement for that usage scenario.

8. Security Considerations

This entire memo is about security.

9. IANA Considerations

None.

10. Acknowledgements

Thanks to Ralph Blom, Hannes Tschofenig, Dan York, Alfred Hoenes, Martin Ellis, Ali Begen, Keith Drage, Ray van Brandenburg, Stephen Farrell, Sean Turner, and John Mattsson for their feedback.

11. Informative References

- [I-D.ietf-avtcore-rtp-security-options]
Westerlund, M. and C. Perkins, "Options for Securing RTP Sessions", draft-ietf-avtcore-rtp-security-options-07 (work in progress), October 2013.
- [I-D.ietf-mmusic-rfc2326bis]
Schulzrinne, H., Rao, A., Lanphier, R., Westerlund, M., and M. Stiemerling, "Real Time Streaming Protocol 2.0 (RTSP)", draft-ietf-mmusic-rfc2326bis-38 (work in progress), October 2013.
- [I-D.ietf-rtcweb-security-arch]
Rescorla, E., "WebRTC Security Architecture", draft-ietf-rtcweb-security-arch-07 (work in progress), July 2013.
- [ISMACrypt2]
Internet Streaming Media Alliance (ISMA), , "ISMA Encryption and Authentication, Version 2.0 release version", November 2007.
- [MBMS]
3GPP, , "Multimedia Broadcast/Multicast Service (MBMS); Protocols and codecs TS 26.346", .
- [RFC0768]
Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.
- [RFC0793]
Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC3365]
Schiller, J., "Strong Security Requirements for Internet Engineering Task Force Standard Protocols", BCP 61, RFC 3365, August 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.
- [RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, March 2006.
- [RFC4585] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July 2006.
- [RFC4614] Duke, M., Braden, R., Eddy, W., and E. Blanton, "A Roadmap for Transmission Control Protocol (TCP) Specification Documents", RFC 4614, September 2006.
- [RFC5479] Wing, D., Fries, S., Tschofenig, H., and F. Audet, "Requirements and Analysis of Media Security Management Protocols", RFC 5479, April 2009.

Authors' Addresses

Colin Perkins
University of Glasgow
School of Computing Science
Glasgow G12 8QQ
UK

Email: csp@csperkins.org
URI: <http://csperkins.org/>

Magnus Westerlund
Ericsson
Farogatan 6
Kista SE-164 80
Sweden

Email: magnus.westerlund@ericsson.com