Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 17, 2013

C. Perkins
University of Glasgow
M. Westerlund
Ericsson

July 16, 2012

Why RTP Does Not Mandate a Single Security Mechanism
draft-ietf-avt-srtp-not-mandatory-09.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 January 17, 2013.

Table of Contents

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 RTCP control 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.

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 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:

- Point-to-point voice telephony (fixed and wireless networks)
- Point-to-point voice and video conferencing
- Centralised group video conferencing with a multipoint conference unit (MCU)
- Any Source Multicast video conferencing (light-weight sessions; Mbone conferencing)
- Point-to-point streaming audio and/or video
- Source-specific multicast (SSM) streaming to large group (IPTV and 3GPP Multimedia Broadcast Multicast Service (MBMS) [MBMS])

- Replicated unicast streaming to a group
- Interconnecting components in music production studios and video editing suites
- Interconnecting components of distributed simulation systems
- Streaming real-time sensor data (e.g., e-VLBI radio astronomy)

As can be seen, these scenarios vary from point-to-point to 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 are inter-domain, with untrusted (and 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 both low overhead, 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 in use, however, and alternatives are more appropriate for some scenarios, and necessary in some cases where SRTP is not suitable. At present, there is no media security protocol that is appropriate for all the environments where RTP is used. Multiple RTP media security protocols can be expected to remain in wide use for the forseeable future.

The range of available RTP security options, and their applicability, are described in [I-D.ietf-avtcore-rtp-security-options].

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, the range of use cases is wide, and there is no single key management protocol that is appropriate for all scenarios. These solutions have been further diversified by the existence of infrastructure elements such as authentication solutions that are tied into the key management. Some of the available keying options for RTP sessions are described in [I-D.ietf-avtcore-rtp-security-options], although this list is not ensured to be exhaustive but include the ones known to the authors at the time of publication.

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 for use by different classes of applications, in 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, and authentication. When new uses for the framework arise, they need to be studied to check if the existing building blocks satisfy the requirements. 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 where IETF has the appropriate knowledge of the class of applications.

6. Security Mechanisms for RTP

RTP is a framework protocol, so the arguments in in Section 5 apply. The security building blocks available for RTP at the time of this writing are described in [I-D.ietf-avtcore-rtp-security-options]. That memo also gives examples of how those security building blocks can be combined to give mandatory to implement security for some RTP application scenarios.

RTP can be extended in different ways. Two important extension points are RTP Payload Formats and RTP Profiles. An RTP Payload Format defines how the output of a new media codec can be used with RTP. It is appropriate for an RTP payload format to discuss specific security implications of using that codec with RTP, but it is not appropriate for an RTP payload format to mandate the use of SRTP, or any other security building blocks, since that payload format might be used in a range of different scenarios.

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. 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. Any new RTP profile ought to discuss if it makes sense to mandate particular security building blocks be used

with implementations of that profile, but without the expectation that all RTP profiles will mandate particular security solutions.

7. Conclusions

RTP is used in a wide range of scenarios, without comon security requirements. Accordingly, a single security solution cannot be mandated for all scenarios. In the absence of such a solution, it is hoped that this memo explains why SRTP is not mandatory as the media security solution for RTP-based systems, and why we can expect multiple key management solutions for systems using RTP.

It is important consider how strong and interoperable security can be offered for every scenario in which RTP applications are used, and for every class of RTP applications. This will require analysis to determine the security requirements, 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. Commonality of security mechanisms is desirable, where appropriate.

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, and Keith Drage 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-00 (work in progress), July 2012.

[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.

[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

Magnus Westerlund
Ericsson
Farogatan 6
Kista SE-164 80
Sweden

Email: magnus.westerlund@ericsson.com