

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 6, 2009

C. Perkins
University of Glasgow
M. Westerlund
Ericsson
November 2, 2008

Why RTP Does Not Mandate a Single Security Mechanism
draft-ietf-avt-srtp-not-mandatory-01.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 6, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This memo discusses the problem of securing real-time multimedia sessions, and explains why the Real-time Transport Protocol (RTP) does not mandate a single media security mechanism.

Table of Contents

1. Introduction 3
2. RTP Applications and Deployment Scenarios 3
3. Implications for RTP Media Security 4
4. Implications for Key Management 5
5. On the Requirement for Strong Security in IETF protocols . . . 6
6. Conclusions 7
7. Security Considerations 7
8. IANA Considerations 7
9. Informative References 7
Authors' Addresses 9
Intellectual Property and Copyright Statements 10

1. Introduction

The Real-time Transport Protocol (RTP) [1] is widely used for voice over IP, Internet television, video conferencing, and various other real-time and streaming media applications. Despite this, the base RTP specification provides very limited options for media security, and defines no standard key exchange mechanism. Rather, a number of extensions are defined to provide confidentiality and authentication of media streams, and to exchange security keys. This memo outlines why it is appropriate that multiple extension mechanisms are defined, rather than mandating a single media security and keying mechanism.

This memo provides information for the community; it does not specify a standard of any kind.

The structure of this memo is as follows: we begin, in Section 2 by describing the scenarios in which RTP is deployed. Following this, Section 3 outlines the implications of this range of scenarios for media confidentiality and authentication, and Section 4 outlines the implications for key exchange. Section 5 outlines how the RTP framework meets the requirement of BCP 61. Section 6 then concludes and gives some recommendations. Finally, Section 7 outlines the security considerations, and Section 8 outlines IANA considerations.

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 (fixed and wireless networks)
- o Point-to-point video conferencing
- o Centralised group video conferencing with a multipoint conference unit (MCU)
- o Any Source Multicast video conferencing (light-weight sessions; Mbone conferencing)
- o Point-to-point streaming audio and/or video
- o Single Source Multicast streaming to large group (IPTV and MBMS [2])
- o Replicated unicast streaming to a group

- o Interconnecting components in music production studios and video editing suites
- o Interconnecting components of distributed simulation systems
- o Streaming real-time sensor data

As can be seen, these scenarios vary from point-to-point to very large multicast groups, from interactive to non-interactive, and from low bandwidth (kilobits per second) to very high bandwidth (multiple gigabits per second). While most of these applications run over UDP, some use TCP or DCCP 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. Implications for RTP Media Security

The wide range of application scenarios where RTP is used has led to the development of multiple solutions for media security, considering different requirements. Perhaps the most widely applicable of these solutions is the Secure RTP (SRTP) framework [3]. This is an application-level media security solution, encrypting the media payload data (but not the RTP headers) to provide some degree of confidentiality, and providing optional source origin authentication. It was carefully designed to be both low overhead, and to support the group communication 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. For example, many client-server streaming media applications run over a single TCP connection, multiplexing media data with control information on that connection (for example, on an RTSP connection). The natural way to provide media security for such client-server media applications is to use TLS to protect the TCP connection, sending the RTP media data over the TLS connection. Using the SRTP framework in addition to TLS is unnecessary, and would result in double encryption of the media, and SRTP cannot be used instead of TLS since it is RTP-specific, and so cannot protect the control traffic.

Other RTP use cases work over networks which provide security at the network layer, using IPsec. For example, certain 3GPP networks need IPsec security associations for other purposes, and can reuse those to secure the RTP session [4]. SRTP is, again, unnecessary in such environments, and its use would only introduce overhead for no gain.

For some applications it is sufficient to protect the RTP payload data while leaving RTP, transport, and network layer headers unprotected. An example of this is RTP broadcast over DVB-H [5], where one mode of operation uses ISMACryp to protect the media data only.

Finally, the link layer may be secure, and it may be known that the RTP media data is constrained to that single link (for example, when operating in a studio environment, with physical link security). An environment like this is inherently constrained, but might avoid the need for application, transport, or network layer media security.

All these are application scenarios where RTP has seen commercial deployment. Other use case also exist, with additional requirements. There is no media security protocol that is appropriate for all these environments. Accordingly, multiple RTP media security protocols can be expected to remain in wide use.

4. Implications for Key Management

With such a diverse range of use case come a range of different protocols for RTP session establishment. Mechanisms used to provide security keying for these different session establishment protocols can basically be put into two categories: inband and out-of-band in relation to the session establishment mechanism. The requirements for these solutions are highly varying. Thus a wide range of solutions have been developed in this space:

- o The most common use case for RTP is probably point-to-point voice calls or centralised group conferences, negotiated using SIP with the SDP offer/answer model, operating on a trusted infrastructure. In such environments, SDP security descriptions [6] or the MIKEY [7] protocol are appropriate keying mechanisms, piggybacked onto the SDP exchange. The infrastructure may be secured by protecting the SIP exchange using SIPS or S/MIME, for example.
- o Point-to-point RTP sessions may be negotiated using SIP with the offer/answer model, but operating over a network with untrusted infrastructure. In such environments, the key management protocol is run on the media path, bypassing the untrusted infrastructure. Protocols such as DTLS [8] or ZRTP [9] are useful here.
- o For point-to-point client-server streaming of RTP over RTSP [10], a TLS association is appropriate to manage keying material, in much the same manner as would be used to secure an HTTP session.

- o A session description may be sent by email, secured using X.500 or PGP, or retrieved from a web page, using HTTP with TLS.
- o A session description may be distributed to a multicast group using SAP or FLUTE secured with S/MIME.
- o A session description may be distributed using OMA's DRM key management [11] with pointer for point to point streaming setup with RTSP in 3GPP [12].
- o In the 3GPP MBMS system, HTTP and MIKEY are used for key management [13].

A more detailed survey of requirements for media security management protocols can be found in [14]. As can be seen, the range of use cases is wide, and there is no single 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.

5. On the Requirement for Strong Security in IETF protocols

BCP 61 [15] puts a requirement on IETF protocols to provide strong, mandatory to implement, security solutions. This is actually quite a difficult requirement for any type of framework protocol, like RTP, since one can never know all the deployment scenarios, and if they are covered by the security solution. It would clearly be desirable if a single media security solution and a single key management solution could be developed, satisfying the range of use cases for RTP. The authors are not aware of any such solution, however, and it is not clear that any single solution can be developed.

For a framework protocol it appears that the only sensible solution to the requirement of BCP 61 is to develop or use security building blocks, like SRTP, SDES, MIKEY, DTLS, or IPsec, to provide the basic security services of authorization, data integrity protection and data confidentiality protection. When new usages of the RTP framework arise, one needs to analyze the situation, to determine if the existing building blocks satisfy the requirements. If not, it is necessary to develop new security building blocks.

When it comes to fulfilling the "MUST Implement" strong security for a specific application, it will fall on that application to actually consider what building blocks it is required to support. To maximize interoperability it is desirable if certain applications, or classes of application with similar requirements, agree on what data security mechanisms and key-management should be used. If such agreement is

not possible, there will be increased cost, either in the lack of interoperability, or in the need to implement more solutions. Unfortunately this situation, if not handled reasonably well, can result in a failure to satisfy the requirement of providing the users with an option of turning on strong security when desired.

6. Conclusions

As discussed earlier it appears that a single solution can't be designed to meet the diverse requirements. 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 for any RTP-based application to consider how it meets the security requirements. This will require some analysis to determine these requirements, followed by a selection of a mandatory to implement solution, or in exceptional scenarios several solutions, including the desired RTP traffic protection and key-management. SRTP is a preferred solution for the protection of the RTP traffic in those use cases where it is applicable. It is out of scope for this memo to recommend a preferred key management solution.

7. Security Considerations

This entire memo is about security.

8. IANA Considerations

No IANA actions are required.

9. Informative References

- [1] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [2] 3GPP, "Multimedia Broadcast/Multicast Service (MBMS); Protocols and codecs TS 26.346".
- [3] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.

- [4] 3GPP, "IP network layer security", 3GPP TS 33.210, September 2008.
- [5] ETSI, "Digital Video Broadcasting (DVB); IP Datacast over DVB-H: Service Purchase and Protection", ETSI TS 102 474, November 2007.
- [6] Andreasen, F., Baugher, M., and D. Wing, "Session Description Protocol (SDP) Security Descriptions for Media Streams", RFC 4568, July 2006.
- [7] Arkko, J., Lindholm, F., Naslund, M., Norrman, K., and E. Carrara, "Key Management Extensions for Session Description Protocol (SDP) and Real Time Streaming Protocol (RTSP)", RFC 4567, July 2006.
- [8] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for Secure Real-time Transport Protocol (SRTP)", draft-ietf-avt-dtls-srtp-06 (work in progress), October 2008.
- [9] Zimmermann, P., Johnston, A., and J. Callas, "ZRTP: Media Path Key Agreement for Secure RTP", draft-zimmermann-avt-zrtp-10 (work in progress), October 2008.
- [10] Schulzrinne, H., Rao, A., Lanphier, R., Westerlund, M., and M. Stiemerling, "Real Time Streaming Protocol 2.0 (RTSP)", draft-ietf-mmusic-rfc2326bis-18 (work in progress), May 2008.
- [11] Open Mobile Alliance, "DRM Specification 2.0".
- [12] 3GPP, "Transparent end-to-end Packet-switched Streaming Service (PSS); Protocols and codecs TS 26.234".
- [13] 3GPP, "Security of Multimedia Broadcast/Multicast Service (MBMS) TS 33.246".
- [14] Wing, D., Fries, S., Tschofenig, H., and F. Audet, "Requirements and Analysis of Media Security Management Protocols", draft-ietf-sip-media-security-requirements-08 (work in progress), October 2008.
- [15] Schiller, J., "Strong Security Requirements for Internet Engineering Task Force Standard Protocols", BCP 61, RFC 3365, August 2002.

Authors' Addresses

Colin Perkins
University of Glasgow
Department of Computing Science
Sir Alwyn Williams Building
Lilybank Gardens
Glasgow G12 8QQ
UK

Email: csp@csperkins.org

Magnus Westerlund
Ericsson
Torshamgatan 23
Stockholm SE-164 80
Sweden

Email: magnus.westerlund@ericsson.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

