Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 5, 2015

J. Mattsson
Y. Cheng
Ericsson

July 4, 2014

Privacy Ensured Cloud Conferencing - Use Case and Goals
draft-mattsson-avtvore-cloud-conferencing-use-case-00

Abstract

The aim of this document is to describe the use case of privacy ensured cloud conferencing in a pervasive monitoring landscape and to point out goals for a solution mitigating the pervasive monitoring threat [RFC7258].

Virtualized cloud-based conferencing is happening and IETF should take action to make such services viable and trustworthy from a pervasive monitoring perspective.

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 5, 2015.

Copyright Notice

Table of Contents

1. Introduction

This document discusses the possibility to provide real-time conference communication services to enterprises and other organizations that try to ensure the privacy of their communication in a world with pervasive monitoring. This includes being able to purchase conferencing supporting network services, including cloud-based ones that are resistant to content monitoring.

This document starts with a background section discussing the development in the world that affects considerations for ensuring communication privacy. Next goals and non-goals for privacy ensured cloud conferencing are stated, followed by the considerations around using current technology and standards.

Some strategies for secure cloud systems can be found in [I-D.jennings-perpass-secure-rai-cloud].

2. Background

Within the field of real-time conferencing there is an ongoing transformation. A transformation towards more cloud based, virtualized and software based conferencing server implementations. The central conferencing server on dedicated hardware is under heavy competition from virtualized servers. One enabling factor for this is the increased capabilities of the end-points that allow them to decode and process multiple simultaneously received media streams. This in its turn has made the central conferring media nodes to switch from mixing or composing media in the decoded domain to instead perform the much less heavy-weight operation of selection, switching and forwarding of media streams, at least for video. Thus making virtualized cloud-based conferencing services viable.

This transformation to virtualization and cloud-based services increases the threats towards the confidentiality of the content of any conference. The reason is that an attacker interested in surveillance of a conference now also has the possibility to attack the cloud provider and attempt to get access to the actual hardware or virtualization layer as a method of accessing what happens within the conference services server instances.

From the pervasive monitoring debate we know that there are many organizations that are actively performing large scale pervasive monitoring, this includes national agencies, but also criminal organizations may be engaged in such activities. It has been revealed that several large service providers have been compromised, resulting in people questioning the impact and security of sourcing services for enterprise or governments to external providers. IETF has stated that pervasive monitoring is an attack and that it should be mitigated [RFC7258].

The trend of using virtualized cloud-based services (e.g. conferencing) has a number of positive effects on flexibility, CAPEX, ease of use, etc. IETF should take action to make such services viable and trustworthy from a pervasive monitoring perspective. One important part of pervasive monitoring is the passive pervasive monitoring [I-D.trammell-perpass-ppa].

3. Goals and Non-Goals

This section proposed goals and non-goals for the privacy ensured conferencing use case.

3.1. Goals

3.1.1. Ensure End-To-End(s) Confidentiality

The content of the communication and all its media needs to be confidential within the group of entities explicitly invited into the conference. An external monitoring adversary should be unable to deduce the human to human communication that actually occurred from capturing the media packets within a time frame for which the communication remains confidential.

3.1.2. Ensure End-To-End Source Authentication

In a conference system with multiple participants it is vital that the multi-media content presented to any of the participant humans are from the stated participants, and not an adversary that attempts to inject misleading content. Nor should an adversary be able to fool the system into becoming a trusted party in the conference, only explicitly invited parties shall be able to contribute content.

3.1.3. Provide a more efficient service than Full-Mesh

A multi-party conference that has the goals of confidentiality and source authentication can be established as a full MESH, i.e. each participating end-point directly addresses each of the other participants. However, this has a significant issue with the amount of consumed resources in both the uplink and the downlink from each participant. To reduce this issue one wants to consider other topologies, which implies network or centralized server functionalities.

3.1.4. Support Cloud Based Conferencing

To achieve a cost effective and scalable conferencing support the conference central nodes must be possible to run as instances in a cloud based virtualized environment.

From a security stand point this is a significant issue. In a virtual, possibly multi-tenant, cloud environment the ones running the conferencing supporting implementation instance can’t trust that anything maintained at that instance is secure from an adversary. A pervasive monitoring entity may in fact have direct access to the hardware through vulnerabilities in the virtualization layer the instance runs in.

3.2. Non-Goals

3.2.1. Securing the endpoints

The security of a communication session requires that the endpoints are not compromised and that the users are trustworthy. If not, credentials and decrypted content may be shared with third parties. However this is hard to prevent through system design. Thus, it should be assumed that the endpoint is secure and the user is trustwothy, how to achieve this is out of scope.

3.2.2. Concealing that communication occurs

A non-goal is to attempt to prevent a pervasive monitoring adversary from knowing that the communication session has occurred. The reason for excluding this as a goal is that first of all it is extremely difficult to achieve, as a pervasive monitoring adversary can be expected to be able to have knowledge of all IP flows that enter or exit local ISPs, across links that straddle nation borders or internet exchange points. Thus the flows required to achieve the communication session need to be highly difficult to correlate between different legs of the communication. At this stage this is deemed too difficult to attempt and will need to be subject for future studies. Existing attempts include The Onion Router (TOR), which has been claimed to be possible to monitor, at least partially, by an adversary with sufficient reach.

3.2.3. Individual Media Source Authentication

Although the participants in the conference are authenticated, it should probably not be a goal to provide source authentication of the media at the individual user level, instead being satisfied with being able to authenticate media as coming from an invited conference participant or not.

There exist solutions that can provide individual media source authentication, e.g. TESLA. However they impact the performance or security properties they provide. Thus, further studies are required to determine impact and resulting security properties if desired to have individual source authentication.

3.2.4. Preventing invited user to access content

As an invited user will be provided with the content protection keys, an invited user can unless active measures are taken against this, decrypt content from the time periods prior and post the user is officially part of the conference.

If this is a concern the solution could be extended to re-key the content protection keys every time a user joins or leaves the conference so each particular set of conference participants uses a unique key. However, this also changes the trust level required on the conference rooster handling at any point and how to keep that accurate and secured. Further the re-keying operations and their timely completion become an obstacle in system design.

4. Problems with Current Technology

If SRTP is used end-to-end, a multiparty conference where the middlebox/server duplicates packets and forwards the complete encrypted packets from a client to multiple participants, the RTP handling is problematic.

The RTP mixer will be forced to behave like a video switching MCU in RFC 5117. SRTP prevents the mixer from performing any type of RTP or RTCP rewrite. However, to keep the bitrate in check its switching decision will result in stopping RTP streams from reaching the client. This results in RTP sequences with large gaps in them. These gaps hide packet losses at the edges of the gaps, resulting in that the receiver has issues in determining if loss near switching point is intentional or not. This can cause repair attempts, buffering issues, and triggering bit-rate adaptation. In addition the congestion control mechanism has significant difficulties to act correctly in such an environment.

Further the above topology requires the RTP stacks to be capable of handling multiple remote peers, including for adaptation of congestion control. This has previously been limited to any-source multicast and transport relay topologies, not RTP mixer ones.

To enable this system’s properties, enable RTP mixing, while not letting the mixer get access to content, it’s required to specify a two level security mechanism. In any multi-vendor environments this will require a specification as it will affect the cipher operations and the data transmitted between participants. The application could handle automatic key-management in the group of authorized participants. One approach on how to do this is [I-D.cheng-srtp-cloud]

To support video switching/relaying knowing from which points in the video streams a receiving endpoint will be able to decode is important. Thus markers for where switching points in the media stream are will be required.

To enable the middle boxes to take local decisions on this, each sender will need to include some speaker activity indication;

preferably including some type of ranking of how likely this is to contain speech. However, this activity indication also needs to leak as little information as possible about the actual content of the speech.

5. Conclusions

Virtualized cloud-based conferencing is happening and IETF should take action to make such services viable and trustworthy from a pervasive monitoring perspective.

From the goals and discussion above it is clear that to provide effective cloud based conferencing while protecting from pervasive monitoring, two layers of security is needed. This is not supported by SRTP.

There is currently active work on secure objects in the form of JSON objects in the IETF WG JOSE. But this is not applicable to RTP based real-time media.

6. Security Considerations

The whole document is about making cloud-based conferencing viable and trustworthy from a pervasive monitoring perspective.

7. Acknowledgements

The authors would like to thank Magnus Westerlund for providing much of the input to this document as well as much of the text.

8. References

[I-D.cheng-srtp-cloud]
Cheng, Y., Mattsson, J., and Naslund, M., "Secure Real-time Transport Protocol (SRTP) for Cloud Services", draft-cheng-srtp-cloud-00 (work in progress), July 2014.

[I-D.jennings-perpass-secure-rai-cloud]
Jennings, C. and S. Nandakumar, "Trustable Cloud Systems - Strategies and Recommendations", draft-jennings-perpass-secure-rai-cloud-01 (work in progress), January 2014.

[I-D.trammell-perpass-ppa]
Trammell, B., Borkmann, D., and C. Huitema, "A Threat Model for Pervasive Passive Surveillance", draft-trammell-perpass-ppa-01 (work in progress), November 2013.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", BCP 188, RFC 7258, May 2014.

Authors' Addresses

John Mattsson
Ericsson AB
SE-164 80 Stockholm
Sweden

Phone: +46 10 71 43 501
Email: john.mattsson@ericsson.com

Yi Cheng
Ericsson
SE-164 80 Stockholm
Sweden

Phone: +46 10 71 17 589
Email: yi.cheng@ericsson.com