

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

M. Westerlund
B. Burman
F. Jansson
Ericsson
July 16, 2012

Multiple Synchronization sources (SSRC) in RTP Session Signaling
draft-westerlund-avtcore-max-ssrc-02

Abstract

RTP has always been a protocol that supports multiple participants each sending their own media streams in an RTP session. Unfortunately, many implementations are designed only for point to point voice over IP with a single source in each end-point. Even client implementations aimed at video conferences have often been built with the assumption around central mixers that only deliver a single media stream per media type. Thus any application that wants to allow for more advanced usage, where multiple media streams are sent and received by an end-point, has an issue with legacy implementations. This document describes the problem and proposes a signalling solution for how to use multiple SSRCs within one RTP session and at the same time handle the legacy issues.

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.

Copyright Notice

Copyright (c) 2012 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	3
1.1.	Background	3
2.	Definitions	4
2.1.	Requirements Language	4
2.2.	Terminology	4
3.	Multiple Streams Issues	5
3.1.	Legacy Behaviors	5
3.2.	Receiver Limitations	7
3.3.	Transmission Declarations	7
4.	Multiple Streams SDP Extension	8
4.1.	Signaling Support for Multiple Streams	8
4.2.	Declarative Use	9
4.3.	Use in Offer/Answer	10
4.4.	Examples	10
5.	IANA Considerations	11
6.	Security Considerations	11
7.	References	12
7.1.	Normative References	12
7.2.	Informative References	12
	Authors' Addresses	13

1. Introduction

This document discusses the issues of non basic usage of RTP [[RFC3550](#)] where there is multiple media sources sent over an RTP session using the SSRC source identifier to distinguish between the sources. This include multiple sources from the same end-point, multiple end-points each having a source, or an application that sends or receive multiple encodings of a particular media source using multiple SSRCs.

1.1. Background

RTP sessions are a concept which most fundamental part is an SSRC space. This space can encompass a number of network nodes and interconnected transport flows between these nodes. Each node may have zero, one or more source identifiers (SSRCs) used to either identify a real media source such as a camera or a microphone, a conceptual source (like the most active speaker selected by an RTP mixer that switches between incoming media streams based on the media stream or additional information), or simply as an identifier for a receiver that provides feedback and reports on reception. There are also RTP nodes, like translators that are manipulating data, transport or session state without making their presence aware to the other session participants.

RTP was designed to support multiple participants in a session from the beginning. This was not restricted to multicast as many believe but also unicast using either multiple transport flows below RTP or a network node that redistributes the RTP packets, either unchanged in the form of a transport translator (relay) or modified in an RTP mixer. There is also the case where a single end-point have multiple media sources, like multiple cameras or microphones.

However, the most common use cases for RTP have been point to point Voice over IP (VoIP) or streaming applications where there have commonly not been more than one media source per end-point. Even in conferencing applications, especially voice only, the conference focus or bridge have provided a single stream being a mix of the other participants to each participant. Thus there has been little need for handling multiple SSRCs in implementations. This has resulted in an installed legacy base that is not fully RTP specification compliant and will have different issues if they receive multiple SSRCs of media, either simultaneously or in sequence. These issues will manifest themselves in various ways, either by software crashes, or simply in limited functionality, like only decoding and playing back the first or latest SSRC received and discarding any other SSRCs.

The signaling solutions around RTP, especially the SDP [[RFC4566](#)] based, have not considered the fundamental issues around an RTP session's theoretical support of up to 4 billion plus sources all sending media. No end-point has infinite processing resources to decode and mix any number of media sources. In addition the memory for storing related state, especially decoder state is limited, and the network bandwidth to receive multiple streams is also limited. Today, the most likely limitations are processing and network bandwidth although for some use cases memory or other limitations may also exist. The issue is that a given end-point will have some limitations in the number of streams it simultaneously can receive, decode and playback. These limitations need to be possible to expose and enabling the session participants to take them into account.

In similar ways there is a need for an end-point to express if it intends to produce one or more media streams in an RTP session. Today's SDP signaling support for this is basically the directionality attribute which indicates an end-point intent to send media or not. There is however no way to indicate how many media streams will be sent.

Taking these things together there exist a clear need to enable the usage of multiple simultaneous media streams within an RTP session in a way that allows a system to take legacy implementations into account in addition to negotiate the actual capabilities around the multiple streams in an RTP session.

[2.](#) Definitions

[2.1.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[2.2.](#) Terminology

The following terms and abbreviations are used in this document:

Encoding: A particular encoding is the choice of the media encoder (codec) that has been used to compress the media, the fidelity of that encoding through the choice of sampling, bit-rate and other configuration parameters.

Different encodings: An encoding is different when some parameter that characterize the encoding of a particular media source has been changed. Such changes can be one or more of the following parameters; codec, codec configuration, bit-rate, sampling.

3. Multiple Streams Issues

This section attempts to go a bit more in depth around the different issues when using multiple media streams in an RTP session to make it clear that although in theory multi-stream applications should already be possible to use, there are good reasons to create extensions for signaling. In addition, the RTP specification could benefit from clarifications on how certain mechanisms should be working when an RTP session contains more than two SSRCs.

3.1. Legacy Behaviors

It is a common assumption among many applications using RTP that they do not have a need to support more than one incoming and one outgoing media stream per RTP session. For a number of applications this assumption has been correct. For VoIP and Streaming applications it has been easiest to ensure that a given end-point only receives and/or sends a single stream. However, all end-points should support a source changing SSRC value during a session, e.g due to SSRC value collision between participants in a conference and the requirement to always use unique SSRC values.

Some RTP extension mechanisms require the RTP stacks to handle additional SSRCs, like SSRC multiplexed RTP retransmission described in [[RFC4588](#)]. However, that still has only required handling a single media decoding chain.

There are however applications that clearly can benefit from receiving and using multiple media streams simultaneously. A very basic case would be T.140 conversational text, where the text characters are transmitted as a real-time media stream as you type. When used in a multi-party chat scenario, an end-point can receive input from multiple sending end-points where the T.140 RTP Payload Format [[RFC4103](#)] text media is both low bandwidth and where there is no obvious method to algorithmically distinguish between multiple sources of text, making simple multiplex and identification of separate sources through an identifier (SSRC) a good choice.

An RTP session that contains an end-point with more than two SSRCs actively sending media streams puts some requirements on the receiving client, which is not necessarily fulfilled by a legacy client:

1. The receiving client needs to handle receiving more than one stream simultaneously rather than replacing the already existing stream with the new one.
2. Be capable of decoding multiple streams simultaneously.
3. Be capable of rendering multiple streams simultaneously.

An application using multiple streams may be very similar to existing one media stream applications at signaling level. To avoid connecting two different implementations, one that is built to support multiple streams and one that is not, it is important that the capabilities are signaled. This also enables a particular application to separate legacy clients that do not signal it explicitly from the ones that do, and take appropriate measures for the legacy application. Due to that there exist both legacy applications that only handle a single stream and applications that handle multiple streams, a single authoritative interpretation cannot be defined by this specification.

There exist many models for legacy behaviors when it comes to handling multiple SSRCs. This briefly describes some of the more common ones:

Single SSRC per Direction: As previously discussed there exist a large number of applications where each end-point has only a single SSRC and the number of end-points in the RTP session are only two. This class of applications needs to have a legacy fallback behavior that provide only a single SSRC, or issues will likely occur.

Multiple SSRCs per Direction Single Media Stream: There exist some applications that uses multiple concurrent SSRCs but in the end only consumes a single media stream. This include the streaming applications using RTP retransmission with SSRC multiplexing, but also applications switching between sources although only consuming one at a time. When an offer or answer without the signalling extension is received, the end-point can potentially determine this by inspecting payload types in use and correctly determine the legacy behavior. In some cases it must be determined through application context or SDP external signaling indications.

Multi-Party Each Using Multiple SSRCs: Multicast applications commonly have multiple SSRCs, if for no other reason than the existence of multiple end-points in the same RTP session. Similar considerations exist in multi-party applications that uses central nodes. There may be no indications in the SDP regarding how the

application will handle multiple concurrent SSRCS and the expectancy to decode these. Thus also this legacy behavior must commonly be determined from the application context and external signaling.

3.2. Receiver Limitations

An RTP end-point that intends to process the media in an RTP session needs to have sufficient resources to receive and process all the incoming streams. It is extremely likely that no receiver is capable to handle the theoretical upper limit of more than 4 billion media sources in an RTP session. Instead, one or more properties will limit the end-points' capabilities to handle simultaneous media streams. These properties are for example memory, processing, network bandwidth, memory bandwidth, or rendering estate to mention a few possible limitations. Those limits in number of simultaneous streams may also differ depending on what codec (payload type) that is used.

We have also considered the issue of how many simultaneous non-active sources an end-point can handle. We cannot see that inactive media sending SSRCS result in significant resource consumption and there should thus be no need to limit them.

A potential issue that needs to be acknowledged is where a limited set of simultaneously active sources varies within a larger set of session members. As each media decoding chain may contain state, it is important that a receiver can flush a decoding state for an inactive source and if that source becomes active again it does not assume that this previous state exists. Thus, we see need for a signaling solution that allows a receiver to indicate its upper limit in terms of capability to handle simultaneous media streams. We see little need for an upper limitation of RTP session members. Applications will need to account for its own capability to use different codecs simultaneously when choosing general and payload specific limits.

3.3. Transmission Declarations

In an RTP based system where an end-point may either be legacy or has an explicit upper limit in the number of simultaneous streams, one will encounter situations where the end-point can not receive and process all simultaneous active streams in the session. Instead, the sending end-points or central nodes, like RTP mixers, will have to provide the end-point with a selected set of streams based on various metrics, such as most active, most interesting, or user selected. In addition, the central node may combine multiple media streams using mixing or composition into a new media stream to enable an end-point

to get sufficient source coverage in the session, despite existing limitations.

For such a system to be able to correctly determine the need for central processing, the capabilities needed for such a central processing node, and the potential need for an end-point to do sender side limitations, it is necessary for an end-point to declare how many simultaneous streams it may send. Thus, enabling negotiation of the number of streams an end-point sends. The limit in number of simultaneous streams may also differ depending on what codec (payload type) that is used.

4. Multiple Streams SDP Extension

This section describes an extension of the media-level SDP attributes to support signaling of the end point's multiple stream capabilities.

4.1. Signaling Support for Multiple Streams

A solution to the issues described in the previous section needs to:

- o Enable signaling between the RTP sender and receiver how many simultaneous RTP streams that can be handled, including a possibility to specify a different number of streams depending on what codec (payload type) that is used.
- o Be able to handle the case where the number of RTP streams that can be sent from a client do not match the number of streams that can be received by the same client.

It is also a requirement that a multiple streams capable RTP sender MUST be able to adapt the number of sent streams to the RTP receiver capability.

For this purpose and for use in SDP, two new media-level SDP attributes are defined, max-send-ssrc and max-recv-ssrc, which can be used independently to establish a limit to the number of simultaneously active SSRCS for the send and receive directions, respectively. Active SSRCS are the ones counted as senders according to [\[RFC3550\]](#), i.e. they have sent RTP packets during the last two regular RTCP reporting intervals. Alternatively, if the end-points supports PAUSE and RESUME signaling [\[I-D.westerlund-avtext-rtp-stream-pause\]](#), any stream that the media sender sent a PAUSED indication for can be regarded as non-Active and can immediately be replaced by another SSRC, considering the limits established by this specification, including possible changes in the applicable limit required by a change of codec.

The syntax for the attributes is in ABNF [[RFC5234](#)]:

```
max-ssrc      = "a="( "max-send-ssrc:" / "max-recv-ssrc:" ) alt-list
alt-list      = alt-set *(WSP alt-set)
alt-set       = "{" alt *("&" alt)) "}"
alt           = pt ":" limit
pt            = ( pt-list / pt-wildcard )
pt-list       = ( pt-value / pt-range ) *(","( pt-value / pt-range ))
pt-value      = 1*3DIGIT
pt-range      = pt-value "-" pt-value
pt-wildcard   = "*"
limit         = 1*8DIGIT
; WSP and DIGIT defined in [RFC5234]
```

A Payload Type (PT)-agnostic upper limit to the total number of simultaneous SSRCs that can be sent or received in this RTP session is signaled with a "*" instead of the PT number. A value of 0 MAY be used as maximum number of SSRC, but it is then RECOMMENDED that this is also reflected using the sendonly or recvonly attribute.

A PT-specific upper limit to the total number of simultaneous SSRCs in the RTP session with that specific PT is signaled with a defined PT (static, or dynamic through rtpmap). Multiple, alternative sets of PT and limit MAY be specified on the same line, where each set indicates the codec-dependent limit when a certain PT is used. A combination of a comma-separated list of PT and a range of PT sharing a single limit MAY be used. Within the alternative set, there MAY be a specification of limits valid when different PT are used simultaneously. Any PT-agnostic specification on a line MUST be interpreted as valid for any PT that was not included within an explicit limit within that alternative set. Multiple lines with max-send-ssrc or max-recv-ssrc attributes specifying a single PT MAY be used, but MUST NOT contain conflicting limits. PT values that are not defined in the media block MUST be ignored.

When max-send-ssrc or max-recv-ssrc are not included in the SDP, it is to be interpreted in the application's context. The default number of allowed SSRCs can vary depending on the type of application.

[4.2.](#) Declarative Use

When used as a declarative media description, the specified limit in max-send-ssrc indicates the maximum number of simultaneous streams of the specified payload types that the configured end-point may send at any single point in time. Similarly, max-recv-ssrc indicates the maximum number of simultaneous streams of the specified payload types

that may be sent to the configured end-point. Payload-agnostic limits MAY be used with or without additional payload-specific limits.

4.3. Use in Offer/Answer

When used in an offer [[RFC3264](#)], the specified limits indicate the agent's intent of sending and/or capability of receiving that number of simultaneous SSRCS. An offerer supporting this specification MUST include both attributes for sendrecv media streams, even if one or both has a value of 1. For sendonly or recvonly m= blocks, the one matching the Offer/Answer agent's role MUST be included when using this extension, and the other directionality MAY be included for informational purpose, if bi-directionality can potentially be used in the future for this m= block. The answerer MUST reverse the directionality of recognized attributes such that max-send-ssrc becomes max-recv-ssrc and vice versa. The answerer SHOULD modify the offered limits in the answer to suit the answering client's capability and intentions. A sender MUST NOT send more simultaneous streams of the specified payload type(s) than the receiver has indicated ability to receive, taking into account also any payload type-agnostic limit.

In case an answer fails to include any of the limitation attributes, the agent is RECOMMENDED to have defined a suitable interpretation for the application context. This may be the choice of being capable of supporting only a single stream (SSRC) in the direction for which attributes are missing. It may also indicate multiple SSRC support depending on the application. In case the offer lacks both max-send-ssrc and max-recv-ssrc, they MUST NOT be included in the answer.

4.4. Examples

The SDP examples below are not complete. Only relevant parts have been included, for brevity and readability.

```
m=video 49200 RTP/AVP 99
a=rtpmap:99 H264/90000
a=max-send-ssrc:{*:2}
a=max-recv-ssrc:{*:4}
```

An offer with a stated intention of sending 2 simultaneous SSRCS and a capability to receive 4 simultaneous SSRCS.

```
m=video 50324 RTP/AVP 96 97
a=rtpmap:96 H264/90000
a=rtpmap:97 H263-2000/90000
a=max-recv-ssrc:{96:2&97:3} {96:1&97:4} {97:5}
```



```
a=max-send-ssrc:{* 1}
```

An offer to receive at most 5 SSRCs, at most 2 of which using payload type 96 and the rest using payload type 97. The "max-send-ssrc" is used to explicitly indicate the value of 1.

```
m=video 50324 RTP/AVP 96 97 98
a=rtpmap:96 H264/900000
a=rtpmap:97 H263-2000/900000
a=rtpmap:98 H263/900000
a=max-recv-ssrc:{96:2&97:3} {96-97:2&98:1} {96,98:2&97:1}
a=max-recv-ssrc:{96,98:1&97:3} {96:1&97-98:2} {96-97:1&98:3}
a=max-recv-ssrc:{96:1&98:4} {97:3&98:2} {97:2&98:3} {97:1&98:4}
a=max-recv-ssrc:{98:5}
a=max-send-ssrc:{* 1}
```

An offer to receive at most 5 SSRCs, at most 2 of which using payload type 96, at most 3 of which using payload type 97, and at most 5 using payload type 98. Since there are many combinations, more than one "max-recv-ssrc" line is used, which is OK since no specifications on those lines are in conflict.

5. IANA Considerations

This document registers two media level SDP attributes.

6. Security Considerations

The SDP attributes defined in this document "a=max-recv-ssrc" and "a=max-send-ssrc" signals capabilities of the end-point. Thus they are vulnerable to attacks. The primary security concerns would be with third parties that modifies the values of the attributes or inserts the attributes in a signalling context. Thus changing the peers view of the others peers capabilities and proposals. A modification reducing either of send or receive values will degrade the service, potentially preventing the service all together. Increasing the value or inserting the attribute with a value different from 1 have the potential of being even more effective. It can result in that an end-point that only supports a single stream will be sent multiple streams. First of all, this potentially exposes software flaws regarding handling of multiple streams, thus causing crashes, less severe it can cause media degradation as the receiving entity flaps between media streams, or plays only a single one, where the other side assumes both will be played. In addition, negotiation of several streams has transport impact, potentially increasing the bit-rate consumed towards the end-point, and in

addition forcing an adaptation response over a limited path, thus degrading the media stream the end-point may play out.

To prevent third party manipulation of the SDP, it should be source authenticated and integrity protected. The solution suitable for this depends on the signalling protocol being used. For SIP S/MIME [[RFC3261](#)] is the ideal, and hop by hop TLS provides at least some protection, although not perfect. For SDPs retrieved using RTSP DESCRIBE [[RFC2326](#)], TLS would be the RECOMMENDED solution.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.

7.2. Informative References

- [I-D.westerlund-avtext-rtp-stream-pause]
Akram, A., Burman, B., Grondal, D., and M. Westerlund,
"RTP Media Stream Pause and Resume",
[draft-westerlund-avtext-rtp-stream-pause-02](#) (work in progress), July 2012.
- [RFC2326] Schulzrinne, H., Rao, A., and R. Lanphier, "Real Time Streaming Protocol (RTSP)", [RFC 2326](#), April 1998.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", [RFC 3261](#), June 2002.
- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC4103] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", [RFC 4103](#), June 2005.

- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", [RFC 4588](#), July 2006.

Authors' Addresses

Magnus Westerlund
Ericsson
Farogatan 6
SE-164 80 Kista
Sweden

Phone: +46 10 714 82 87
Email: magnus.westerlund@ericsson.com

Bo Burman
Ericsson
Farogatan 6
SE-164 80 Kista
Sweden

Phone: +46 10 714 13 11
Email: bo.burman@ericsson.com

Fredrik Jansson
Ericsson
Farogatan 6
Kista, SE-164 80
Sweden

Phone: +46 10 719 00 00
Fax:
Email: fredrik.k.jansson@ericsson.com
URI:

