

Internet Engineering Task Force (IETF)
Request for Comments: 7656
Category: Informational
ISSN: 2070-1721

J. Lennox
Vidyo
K. Gross
AVA
S. Nandakumar
G. Salgueiro
Cisco Systems
B. Burman, Ed.
Ericsson
November 2015

A Taxonomy of Semantics and Mechanisms for
Real-Time Transport Protocol (RTP) Sources

Abstract

The terminology about, and associations among, Real-time Transport Protocol (RTP) sources can be complex and somewhat opaque. This document describes a number of existing and proposed properties and relationships among RTP sources and defines common terminology for discussing protocol entities and their relationships.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7656>.

Copyright Notice

Copyright (c) 2015 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	5
2.	Concepts	5
2.1.	Media Chain	5
2.1.1.	Physical Stimulus	10
2.1.2.	Media Capture	10
2.1.3.	Raw Stream	10
2.1.4.	Media Source	11
2.1.5.	Source Stream	11
2.1.6.	Media Encoder	12
2.1.7.	Encoded Stream	13
2.1.8.	Dependent Stream	13
2.1.9.	Media Packetizer	13
2.1.10.	RTP Stream	14
2.1.11.	RTP-Based Redundancy	14
2.1.12.	Redundancy RTP Stream	15
2.1.13.	RTP-Based Security	15
2.1.14.	Secured RTP Stream	16
2.1.15.	Media Transport	16
2.1.16.	Media Transport Sender	17
2.1.17.	Sent RTP Stream	18
2.1.18.	Network Transport	18
2.1.19.	Transported RTP Stream	18
2.1.20.	Media Transport Receiver	18
2.1.21.	Received Secured RTP Stream	19
2.1.22.	RTP-Based Validation	19
2.1.23.	Received RTP Stream	19
2.1.24.	Received Redundancy RTP Stream	19
2.1.25.	RTP-Based Repair	19
2.1.26.	Repaired RTP Stream	19
2.1.27.	Media Depacketizer	20
2.1.28.	Received Encoded Stream	20
2.1.29.	Media Decoder	20
2.1.30.	Received Source Stream	20
2.1.31.	Media Sink	21
2.1.32.	Received Raw Stream	21
2.1.33.	Media Render	21
2.2.	Communication Entities	22
2.2.1.	Endpoint	23
2.2.2.	RTP Session	23
2.2.3.	Participant	24
2.2.4.	Multimedia Session	24
2.2.5.	Communication Session	25
3.	Concepts of Inter-Relations	25
3.1.	Synchronization Context	26
3.1.1.	RTCP CNAME	26
3.1.2.	Clock Source Signaling	26

3.1.3.	Implicitly via RtcMediaStream	26
3.1.4.	Explicitly via SDP Mechanisms	26
3.2.	Endpoint	27
3.3.	Participant	27
3.4.	RtcMediaStream	27
3.5.	Multi-Channel Audio	28
3.6.	Simulcast	28
3.7.	Layered Multi-Stream	30
3.8.	RTP Stream Duplication	32
3.9.	Redundancy Format	33
3.10.	RTP Retransmission	33
3.11.	Forward Error Correction	35
3.12.	RTP Stream Separation	36
3.13.	Multiple RTP Sessions over one Media Transport	37
4.	Mapping from Existing Terms	37
4.1.	Telepresence Terms	37
4.1.1.	Audio Capture	37
4.1.2.	Capture Device	37
4.1.3.	Capture Encoding	38
4.1.4.	Capture Scene	38
4.1.5.	Endpoint	38
4.1.6.	Individual Encoding	38
4.1.7.	Media Capture	38
4.1.8.	Media Consumer	38
4.1.9.	Media Provider	39
4.1.10.	Stream	39
4.1.11.	Video Capture	39
4.2.	Media Description	39
4.3.	Media Stream	39
4.4.	Multimedia Conference	39
4.5.	Multimedia Session	40
4.6.	Multipoint Control Unit (MCU)	40
4.7.	Multi-Session Transmission (MST)	40
4.8.	Recording Device	41
4.9.	RtcMediaStream	41
4.10.	RtcMediaStreamTrack	41
4.11.	RTP Receiver	41
4.12.	RTP Sender	41
4.13.	RTP Session	41
4.14.	Single-Session Transmission (SST)	41
4.15.	SSRC	42
5.	Security Considerations	42
6.	Informative References	42
	Acknowledgements	45
	Contributors	45
	Authors' Addresses	46

1. Introduction

The existing taxonomy of sources in the Real-time Transport Protocol (RTP) [RFC3550] has previously been regarded as confusing and inconsistent. Consequently, a deep understanding of how the different terms relate to each other becomes a real challenge. Frequently cited examples of this confusion are (1) how different protocols that make use of RTP use the same terms to signify different things and (2) how the complexities addressed at one layer are often glossed over or ignored at another.

This document improves clarity by reviewing the semantics of various aspects of sources in RTP. As an organizing mechanism, it approaches this by describing various ways that RTP sources are transformed on their way between sender and receiver, and how they can be grouped and associated together.

All non-specific references to ControLling mUltiple streams for tElepresence (CLUE) in this document map to [CLUE-FRAME], and all references to Web Real-time Communications (WebRTC) map to [WEBRTC-OVERVIEW].

2. Concepts

This section defines concepts that serve to identify and name various transformations and streams in a given RTP usage. For each concept, alternate definitions and usages that coexist today are listed along with various characteristics that further describe the concept. These concepts are divided into two categories: one is related to the chain of streams and transformations that Media can be subject to, and the other is for entities involved in the communication.

2.1. Media Chain

In the context of this document, media is a sequence of synthetic or Physical Stimuli (Section 2.1.1) -- for example, sound waves, photons, key strokes -- represented in digital form. Synthesized media is typically generated directly in the digital domain.

This section contains the concepts that can be involved in taking media at a sender side and transporting it to a receiver, which may recover a sequence of physical stimuli. This chain of concepts is of two main types: streams and transformations. Streams are time-based sequences of samples of the physical stimulus in various representations, while transformations change the representation of the streams in some way.

The below examples are basic ones, and it is important to keep in mind that this conceptual model enables more complex usages. Some will be further discussed in later sections of this document. In general the following applies to this model:

- o A transformation may have zero or more inputs and one or more outputs.
- o A stream is of some type, such as audio, video, real-time text, etc.
- o A stream has one source transformation and one or more sink transformations (with the exception of physical stimulus (Section 2.1.1) that may lack source or sink transformation).
- o Streams can be forwarded from a transformation output to any number of inputs on other transformations that support that type.
- o If the output of a transformation is sent to multiple transformations, those streams will be identical; it takes a transformation to make them different.
- o There are no formal limitations on how streams are connected to transformations.

It is also important to remember that this is a conceptual model. Thus, real-world implementations may look different and have a different structure.

To provide a basic understanding of the relationships in the chain, we first introduce the concepts for the sender side (Figure 1). This covers physical stimuli until media packets are emitted onto the network.

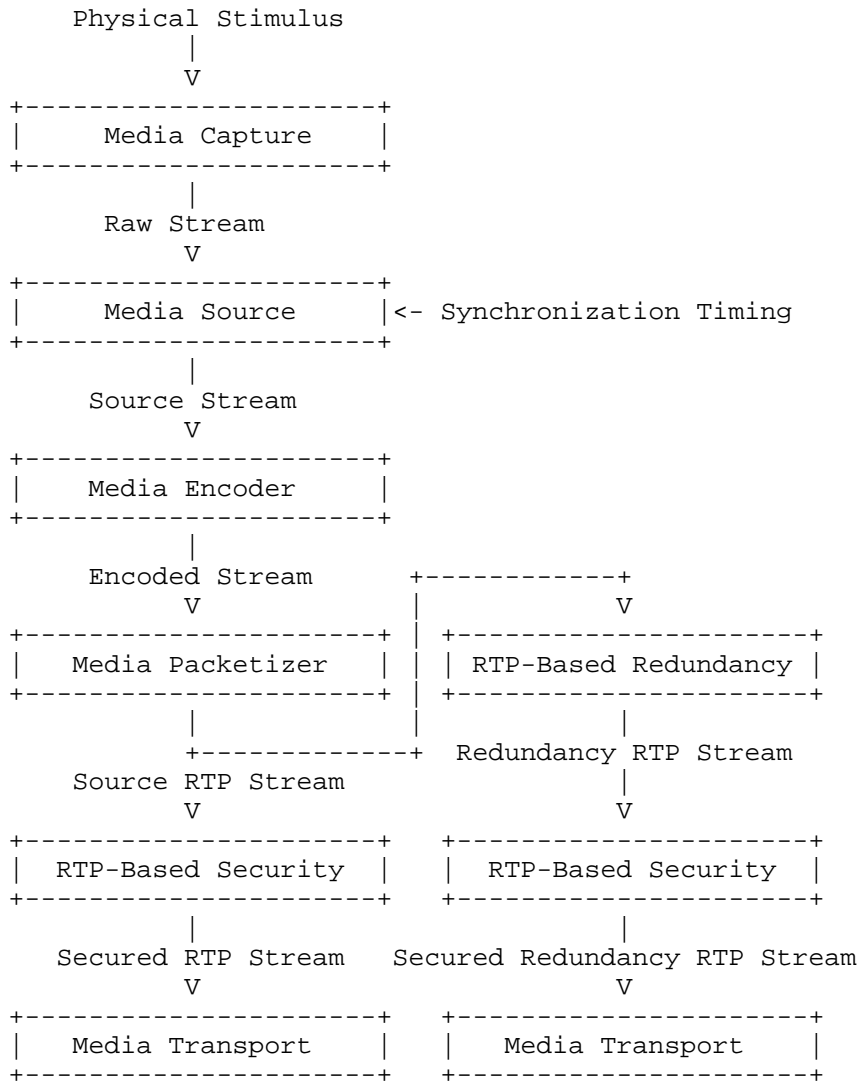


Figure 1: Sender Side Concepts in the Media Chain

In Figure 1, we have included a branched chain to cover the concepts for using redundancy to improve the reliability of the transport. The Media Transport concept is an aggregate that is decomposed in Section 2.1.15.

In Figure 2, we review a receiver media chain matching the sender side, to look at the inverse transformations and their attempts to recover identical streams as in the sender chain, subject to what may be lossy compression and imperfect media transport. Note that the streams out of a reverse transformation, like the Source Stream out of the Media Decoder, are in many cases not the same as the corresponding ones on the sender side; thus, they are prefixed with a "received" to denote a potentially modified version. The reason for not being the same lies in the transformations that can be of irreversible type. For example, lossy source coding in the Media Encoder prevents the source stream out of the media decoder from being the same as the one fed into the media encoder. Other reasons include packet loss in the media transport transformation that even RTP-based Repair, if used, fails to repair.

2.1.1. Physical Stimulus

The physical stimulus is a physical event in the analog domain that can be sampled and converted to digital form by an appropriate sensor or transducer. This includes sound waves making up audio, photons in a light field, or other excitations or interactions with sensors, like keystrokes on a keyboard.

2.1.2. Media Capture

Media Capture is the process of transforming the analog physical stimulus (Section 2.1.1) into digital media using an appropriate sensor or transducer. The media capture performs a digital sampling of the physical stimulus, usually periodically, and outputs this in some representation as a Raw Stream (Section 2.1.3). This data is considered "media", because it includes data that is periodically sampled or made up of a set of timed asynchronous events. The media capture is normally instantiated in some type of device, i.e., media capture device. Examples of different types of media capturing devices are digital cameras, microphones connected to A/D converters, or keyboards.

Characteristics:

- o A media capture is identified either by hardware/manufacture ID or via a session-scoped device identifier as mandated by the application usage.
- o A media capture can generate an Encoded Stream (Section 2.1.7) if the capture device supports such a configuration.
- o The nature of the media capture may impose constraints on the clock handling in some of the subsequent steps. For example, many audio or video capture devices are not completely free in selecting the sample rate.

2.1.3. Raw Stream

A raw stream is the time progressing stream of digitally sampled information, usually periodically sampled and provided by a media capture (Section 2.1.2). A raw stream can also contain synthesized media that may not require any explicit media capture, since it is already in an appropriate digital form.

2.1.4. Media Source

A Media Source is the logical source of a time progressing digital media stream synchronized to a reference clock. This stream is called a source stream (Section 2.1.5). This transformation takes one or more raw streams (Section 2.1.3) and provides a source stream as output. The output is synchronized with a reference clock (Section 3.1), which can be as simple as a system local wall clock or as complex as an NTP synchronized clock.

The output can be of different types. One type is directly associated with a particular media capture's raw stream. Others are more conceptual sources, like an audio mix of multiple source streams (Figure 3). Mixing multiple streams typically requires that the input streams are possible to relate in time, meaning that they have to be source streams (Section 2.1.5) rather than raw streams. In Figure 3, the generated source stream is a mix of the three input source streams.

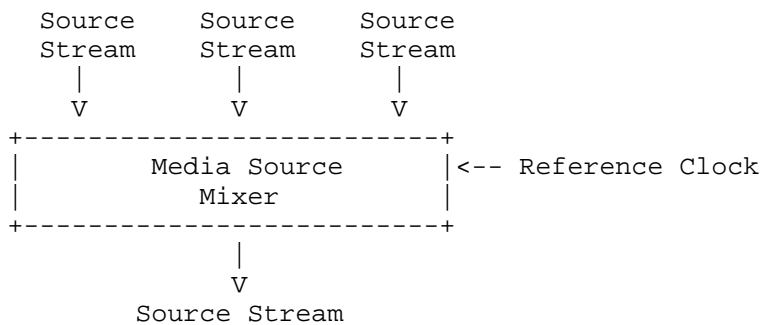


Figure 3: Conceptual Media Source in the form of an Audio Mixer

Another possible example of a conceptual media source is a video surveillance switch, where the input is multiple source streams from different cameras, and the output is one of those source streams based on some selection criteria, such as round robin or some video activity measure.

2.1.5. Source Stream

A source stream is a stream of digital samples that has been synchronized with a reference clock and comes from a particular media source (Section 2.1.4).

2.1.6. Media Encoder

A media encoder is a transform that is responsible for encoding the media data from a source stream (Section 2.1.5) into another representation, usually more compact, that is output as an encoded stream (Section 2.1.7).

The media encoder step commonly includes pre-encoding transformations, such as scaling, resampling, etc. The media encoder can have a significant number of configuration options that affects the properties of the encoded stream. This includes properties such as codec, bitrate, start points for decoding, resolution, bandwidth, or other fidelity affecting properties.

Scalable media encoders need special attention as they produce multiple outputs that are potentially of different types. As shown in Figure 4, a scalable media encoder takes one input source stream and encodes it into multiple output streams of two different types: at least one encoded stream that is independently decodable and one or more Dependent Streams (Section 2.1.8). Decoding requires at least one encoded stream and zero or more dependent streams. A dependent stream's dependency is one of the grouping relations this document discusses further in Section 3.7.

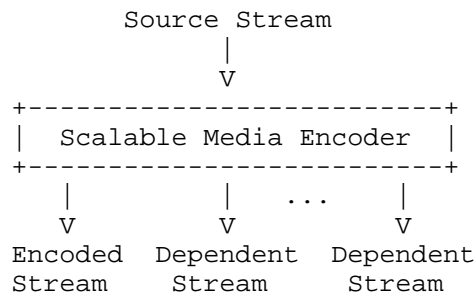


Figure 4: Scalable Media Encoder Input and Outputs

There are also other variants of encoders, like so-called Multiple Description Coding (MDC). Such media encoders produce multiple independent and thus individually decodable encoded streams. However, (logically) combining multiple of these encoded streams into a single Received Source Stream during decoding leads to an improvement in perceptual reproduced quality when compared to decoding a single encoded stream.

Creating multiple encoded streams from the same source stream, where the encoded streams are neither in a scalable nor in an MDC

relationship is commonly utilized in simulcast [SDP-SIMULCAST] environments.

2.1.7. Encoded Stream

A stream of time synchronized encoded media that can be independently decoded.

Due to temporal dependencies, an encoded stream may have limitations in where decoding can be started. These entry points, for example, Intra frames from a video encoder, may require identification and their generation may be event based or configured to occur periodically.

2.1.8. Dependent Stream

A stream of time synchronized encoded media fragments that are dependent on one or more encoded streams (Section 2.1.7) and zero or more dependent streams to be possible to decode.

Each dependent stream has a set of dependencies. These dependencies must be understood by the parties in a Multimedia Session (Section 2.2.4) that intend to use a dependent stream.

2.1.9. Media Packetizer

The transformation of taking one or more encoded (Section 2.1.7) or dependent streams (Section 2.1.8) and putting their content into one or more sequences of packets, normally RTP Packets, and output Source RTP Streams (Section 2.1.10). This step includes both generating RTP Payloads as well as RTP packets. The Media Packetizer then selects which synchronization source(s) (SSRC) [RFC3550] and RTP Sessions (Section 2.2.2) to use.

The media packetizer can combine multiple encoded or dependent streams into one or more RTP Streams:

- o The media packetizer can use multiple inputs when producing a single RTP stream. One such example is Single RTP stream on a Single media Transport (SRST) packetization when using Scalable Video Coding (SVC) (Section 3.7).
- o The media packetizer can also produce multiple RTP streams, for example, when encoded and/or dependent streams are distributed over multiple RTP streams. One example of this is Multiple RTP streams on Multiple media Transports (MRMT) packetization when using SVC (Section 3.7).

2.1.10. RTP Stream

An RTP stream is a stream of RTP packets containing media data, source or redundant. The RTP stream is identified by an SSRC belonging to a particular RTP Session. The RTP session is identified as discussed in Section 2.2.2.

A source RTP stream is an RTP stream directly related to an encoded stream (Section 2.1.7), targeted for transport over RTP without any additional RTP-based Redundancy (Section 2.1.11) applied.

Characteristics:

- o Each RTP stream is identified by an SSRC [RFC3550] that is carried in every RTP and RTP Control Protocol (RTCP) packet header. The SSRC is unique in a specific RTP session context.
- o At any given point in time, an RTP stream can have one and only one SSRC, but SSRCs for a given RTP stream can change over time. SSRC collision and clock rate change [RFC7160] are examples of valid reasons to change SSRC for an RTP stream. In those cases, the RTP stream itself is not changed in any significant way, only the identifying SSRC number.
- o Each SSRC defines a unique RTP sequence numbering and timing space.
- o Several RTP streams, each with their own SSRC, may represent a single media source.
- o Several RTP streams, each with their own SSRC, can be carried in a single RTP session.

2.1.11. RTP-Based Redundancy

RTP-based redundancy is defined here as a transformation that generates redundant or repair packets sent out as a Redundancy RTP Stream (Section 2.1.12) to mitigate Network Transport (Section 2.1.18) impairments, like packet loss and delay. Note that this excludes the type of redundancy that most suitable media encoders (Section 2.1.6) may add to the media format of the encoded stream (Section 2.1.7) that makes it cope better with RTP packet losses.

The RTP-based redundancy exists in many flavors: they may generate independent repair streams that are used in addition to the source stream (like RTP Retransmission (Section 3.10) and some special types of Forward Error Correction (FEC) (Section 3.11), like RTP stream

duplication (Section 3.8)); they may generate a new source stream by combining redundancy information with source information (using XOR FEC as a redundancy payload (Section 3.9)); or they may completely replace the source information with only redundancy packets.

2.1.12. Redundancy RTP Stream

A redundancy RTP stream is an RTP stream (Section 2.1.10) that contains no original source data, only redundant data, which may either be used as standalone or be combined with one or more Received RTP Streams (Section 2.1.23) to produce Repaired RTP Streams (Section 2.1.26).

2.1.13. RTP-Based Security

The optional RTP-based Security transformation applies security services such as authentication, integrity protection, and confidentiality to an input RTP stream, like what is specified in "The Secure Real-time Transport Protocol (SRTP)" [RFC3711], producing a Secured RTP Stream (Section 2.1.14). Either an RTP stream (Section 2.1.10) or a redundancy RTP stream (Section 2.1.12) can be used as input to this transformation.

In SRTP and the related Secure RTCP (SRTCP), all of the above-mentioned security services are optional, except for integrity protection of SRTCP, which is mandatory. Also confidentiality (encryption) is effectively optional in SRTP, since it is possible to use a NULL encryption algorithm. As described in [RFC7201], the strength of SRTP data origin authentication depends on the cryptographic transform and key management used. For example, in group communication, where it is sometimes possible to authenticate group membership but not the actual RTP stream sender.

RTP-based security and RTP-based redundancy can be combined in a few different ways. One way is depicted in Figure 1, where an RTP stream and its corresponding redundancy RTP stream are protected by separate RTP-based security transforms. In other cases, like when a Media Translator is adding FEC in Section 3.2.1.3 of [RTP-TOPOLOGIES], a middlebox can apply RTP-based redundancy to an already secured RTP stream instead of a source RTP stream. One example of that is depicted in Figure 5 below.

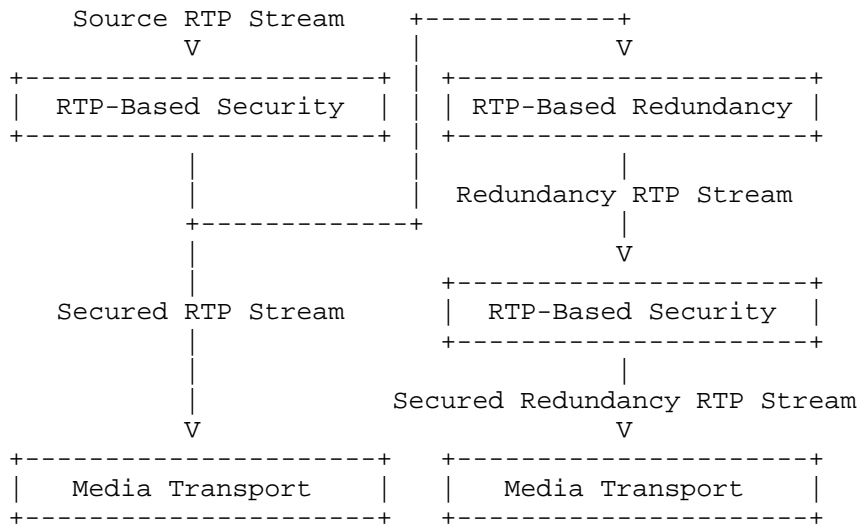


Figure 5: Adding Redundancy to a Secured RTP Stream

In this case, the redundancy RTP stream may already have been secured for confidentiality (encrypted) by the first RTP-based security, and it may therefore not be necessary to apply additional confidentiality protection in the second RTP-based security. To avoid attacks and negative impact on RTP-based Repair (Section 2.1.25) and the resulting repaired RTP stream (Section 2.1.26), it is, however, still necessary to have this second RTP-based security apply both authentication and integrity protection to the redundancy RTP stream.

2.1.14. Secured RTP Stream

A secured RTP stream is a source or redundancy RTP stream that is protected through RTP-based security (Section 2.1.13) by one or more of the confidentiality, integrity, or authentication security services.

2.1.15. Media Transport

A media transport defines the transformation that the RTP streams (Section 2.1.10) are subjected to by the end-to-end transport from one RTP Sender (Section 4.12) to one specific RTP Receiver (Section 4.11) (an RTP session (Section 2.2.2) may contain multiple RTP receivers per sender). Each media transport is defined by a transport association that is normally identified by a 5-tuple (source address, source port, destination address, destination port, transport protocol), but a proposal exists for sending multiple transport associations on a single 5-tuple [TRANSPORT-MULTIPLEX].

Characteristics:

- o Media transport transmits RTP streams of RTP packets from a source transport address to a destination transport address.
- o Each media transport contains only a single RTP session.
- o A single RTP session can span multiple media transports.

The media transport concept sometimes needs to be decomposed into more steps to enable discussion of what a sender emits that gets transformed by the network before it is received by the receiver. Thus, we provide also this media transport decomposition (Figure 6).

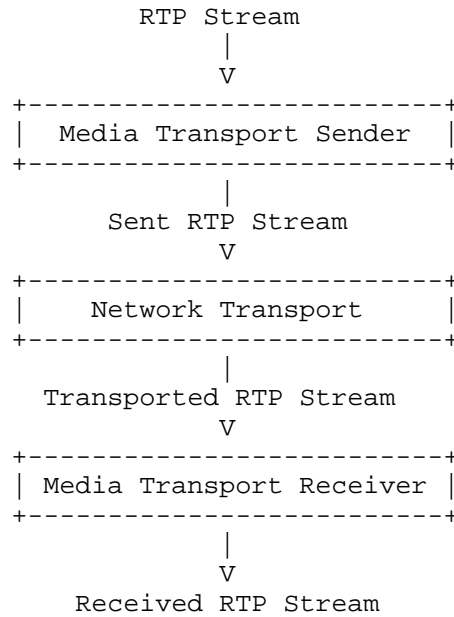


Figure 6: Decomposition of Media Transport

2.1.16. Media Transport Sender

The first transformation within the media transport (Section 2.1.15) is the Media Transport Sender. The sending Endpoint (Section 2.2.1) takes an RTP stream and emits the packets onto the network using the transport association established for this media transport, thereby creating a Sent RTP Stream (Section 2.1.17). In the process, it transforms the RTP stream in several ways. First, it generates the necessary protocol headers for the transport association, for example, IP and UDP headers, thus forming IP/UDP/RTP packets. In

addition, the media transport sender may queue, intentionally pace, or otherwise affect how the packets are emitted onto the network, thereby potentially introducing delay and delay variations [RFC5481] that characterize the sent RTP stream.

2.1.17. Sent RTP Stream

The sent RTP stream is the RTP stream as entering the first hop of the network path to its destination. The sent RTP stream is identified using network transport addresses, like the 5-tuple (source IP address, source port, destination IP address, destination port, and protocol (UDP)) for IP/UDP.

2.1.18. Network Transport

Network transport is the transformation that subjects the sent RTP stream (Section 2.1.17) to traveling from the source to the destination through the network. This transformation can result in loss of some packets, delay, and delay variation on a per-packet basis, packet duplication, and packet header or data corruption. This transformation produces a Transported RTP Stream (Section 2.1.19) at the exit of the network path.

2.1.19. Transported RTP Stream

The transported RTP stream is the RTP stream that is emitted out of the network path at the destination, subjected to the network transport's transformation (Section 2.1.18).

2.1.20. Media Transport Receiver

The Media Transport Receiver is the receiver endpoint's (Section 2.2.1) transformation of the transported RTP stream (Section 2.1.19) by its reception process, which results in the received RTP stream (Section 2.1.23). This transformation includes transport checksums being verified. Sensible system designs typically either discard packets with mismatching checksums or pass them on while somehow marking them in the resulting received RTP stream so to alert subsequent transformations about the possible corrupt state. In this context, it is worth noting that there is typically some probability for corrupt packets to pass through undetected (with a seemingly correct checksum). Other transformations can compensate for delay variations in receiving a packet on the network interface and providing it to the application (de-jitter buffer).

2.1.21. Received Secured RTP Stream

This is the secured RTP stream (Section 2.1.14) resulting from the media transport (Section 2.1.15) aggregate transformation.

2.1.22. RTP-Based Validation

RTP-based Validation is the reverse transformation of RTP-based security (Section 2.1.13). If this transformation fails, the result is either not usable and must be discarded or may be usable but cannot be trusted. If the transformation succeeds, the result can be a received RTP stream (Section 2.1.23) or a Received Redundancy RTP Stream (Section 2.1.24), depending on what was input to the corresponding RTP-based security transformation, but it can also be a Received Secured RTP Stream (Section 2.1.21) in case several RTP-based security transformations were applied.

2.1.23. Received RTP Stream

The received RTP stream is the RTP stream (Section 2.1.10) resulting from the media transport's aggregate transformation (Section 2.1.15), i.e., subjected to packet loss, packet corruption, packet duplication, delay, and delay variation from sender to receiver.

2.1.24. Received Redundancy RTP Stream

The received redundancy RTP stream is the redundancy RTP stream (Section 2.1.12) resulting from the media transport's aggregate transformation, i.e., subjected to packet loss, packet corruption, packet duplication, delay, and delay variation from sender to receiver.

2.1.25. RTP-Based Repair

RTP-based repair is a transformation that takes as input zero or more received RTP streams (Section 2.1.23) and one or more received redundancy RTP streams (Section 2.1.24) and produces one or more repaired RTP streams (Section 2.1.26) that are as close to the corresponding sent source RTP streams (Section 2.1.10) as possible, using different RTP-based repair methods, for example, the ones referred to in RTP-based redundancy (Section 2.1.11).

2.1.26. Repaired RTP Stream

A repaired RTP stream is a received RTP stream (Section 2.1.23) for which received redundancy RTP stream (Section 2.1.24) information has been used to try to recover the source RTP stream (Section 2.1.10) as it was before media transport (Section 2.1.15).

2.1.27. Media Depacketizer

A Media Depacketizer takes one or more RTP streams (Section 2.1.10), depacketizes them, and attempts to reconstitute the encoded streams (Section 2.1.7) or dependent streams (Section 2.1.8) present in those RTP streams.

In practical implementations, the media depacketizer and the media decoder may be tightly coupled and share information to improve or optimize the overall decoding and error concealment process. It is, however, not expected that there would be any benefit in defining a taxonomy for those detailed (and likely very implementation-dependent) steps.

2.1.28. Received Encoded Stream

The Received Encoded Stream is the received version of an encoded stream (Section 2.1.7).

2.1.29. Media Decoder

A media decoder is a transformation that is responsible for decoding encoded streams (Section 2.1.7) and any dependent streams (Section 2.1.8) into a source stream (Section 2.1.5).

In practical implementations, the media decoder and the media depacketizer may be tightly coupled and share information to improve or optimize the overall decoding process in various ways. It is, however, not expected that there would be any benefit in defining a taxonomy for those detailed (and likely very implementation-dependent) steps.

A media decoder has to deal with any errors in the encoded streams that resulted from corruption or failure to repair packet losses. Therefore, it commonly is robust to error and losses, and includes concealment methods.

2.1.30. Received Source Stream

The received source stream is the received version of a source stream (Section 2.1.5).

2.1.31. Media Sink

The Media Sink receives a source stream (Section 2.1.5) that contains, usually periodically, sampled media data together with associated synchronization information. Depending on application, this source stream then needs to be transformed into a raw stream (Section 2.1.3) that is conveyed to the Media Render (Section 2.1.33) and synchronized with the output from other media sinks. The media sink may also be connected with a media source (Section 2.1.4) and be used as part of a conceptual media source.

The media sink can further transform the source stream into a representation that is suitable for rendering on the media render as defined by the application or system-wide configuration. This includes sample scaling, level adjustments, etc.

2.1.32. Received Raw Stream

The Received Raw Stream is the received version of a raw stream (Section 2.1.3).

2.1.33. Media Render

A media render takes a raw stream (Section 2.1.3) and converts it into physical stimulus (Section 2.1.1) that a human user can perceive. Examples of such devices are screens and D/A converters connected to amplifiers and loudspeakers.

An endpoint can potentially have multiple media renders for each media type.

The terms used in Figure 7 are further elaborated in the subsections below.

2.2.1. Endpoint

An endpoint is a single addressable entity sending or receiving RTP packets. It may be decomposed into several functional blocks, but as long as it behaves as a single RTP stack entity, it is classified as a single "endpoint".

Characteristics:

- o Endpoints can be identified in several different ways. While RTCP Canonical Names (CNAMEs) [RFC3550] provide a globally unique and stable identification mechanism for the duration of the communication session (see Section 2.2.5), their validity applies exclusively within a Synchronization Context (Section 3.1). Thus, one endpoint can handle multiple CNAMEs, each of which can be shared among a set of endpoints belonging to the same participant (Section 2.2.3). Therefore, mechanisms outside the scope of RTP, such as application-defined mechanisms, must be used to provide endpoint identification when outside this synchronization context.
- o An endpoint can be associated with at most one participant (Section 2.2.3) at any single point in time.
- o In some contexts, an endpoint would typically correspond to a single "host", for example, a computer using a single network interface and being used by a single human user. In other contexts, a single "host" can serve multiple participants, in which case each participant's endpoint may share properties, for example, the IP address part of a transport address.

2.2.2. RTP Session

An RTP session is an association among a group of participants communicating with RTP. It is a group communications channel that can potentially carry a number of RTP streams. Within an RTP session, every participant can find metadata and control information (over RTCP) about all the RTP streams in the RTP session. The bandwidth of the RTCP control channel is shared between all participants within an RTP session.

Characteristics:

- o An RTP session can carry one or more RTP streams.

- o An RTP session shares a single SSRC space as defined in [RFC3550]. That is, the endpoints participating in an RTP session can see an SSRC identifier transmitted by any of the other endpoints. An endpoint can receive an SSRC either as SSRC or as a contributing source (CSRC) in RTP and RTCP packets, as defined by the endpoints' network interconnection topology.
- o An RTP session uses at least two media transports (Section 2.1.15): one for sending and one for receiving. Commonly, the receiving media transport is the reverse direction of the media transport used for sending. An RTP session may use many media transports and these define the session's network interconnection topology.
- o A single media transport always carries a single RTP session.
- o Multiple RTP sessions can be conceptually related, for example, originating from or targeted for the same participant (Section 2.2.3) or endpoint (Section 2.2.1), or by containing RTP streams that are somehow related (Section 3).

2.2.3. Participant

A participant is an entity reachable by a single signaling address and is thus related more to the signaling context than to the media context.

Characteristics:

- o A single signaling-addressable entity, using an application-specific signaling address space, for example, a SIP URI.
- o A participant can participate in several multimedia sessions (Section 2.2.4).
- o A participant can be comprised of several associated endpoints (Section 2.2.1).

2.2.4. Multimedia Session

A multimedia session is an association among a group of participants (Section 2.2.3) engaged in the communication via one or more RTP sessions (Section 2.2.2). It defines logical relationships among media sources (Section 2.1.4) that appear in multiple RTP sessions.

Characteristics:

- o A multimedia session can be composed of several RTP sessions with potentially multiple RTP streams per RTP session.
- o Each participant in a multimedia session can have a multitude of media captures and media rendering devices.
- o A single multimedia session can contain media from one or more synchronization contexts (Section 3.1). An example of that is a multimedia session containing one set of audio and video for communication purposes belonging to one synchronization context, and another set of audio and video for presentation purposes (like playing a video file) with a separate synchronization context that has no strong timing relationship and need not be strictly synchronized with the audio and video used for communication.

2.2.5. Communication Session

A communication session is an association among two or more participants (Section 2.2.3) communicating with each other via one or more multimedia sessions (Section 2.2.4).

Characteristics:

- o Each participant in a communication session is identified via an application-specific signaling address.
- o A communication session is composed of participants that share at least one multimedia session, involving one or more parallel RTP sessions with potentially multiple RTP streams per RTP session.

For example, in a full mesh communication, the communication session consists of a set of separate multimedia sessions between each pair of participants. Another example is a centralized conference, where the communication session consists of a set of multimedia sessions between each participant and the conference handler.

3. Concepts of Inter-Relations

This section uses the concepts from previous sections and looks at different types of relationships among them. These relationships occur at different abstraction levels and for different purposes, but the reason for the needed relationship at a certain step in the media handling chain may exist at another step. For example, the use of simulcast (Section 3.6) implies a need to determine relations at the

RTP stream level, but the underlying reason is that multiple media encoders use the same media source, i.e., to be able to identify a common media source.

3.1. Synchronization Context

A synchronization context defines a requirement for a strong timing relationship between the media sources, typically requiring alignment of clock sources. Such a relationship can be identified in multiple ways as listed below. A single media source can only belong to a single synchronization context, since it is assumed that a single media source can only have a single media clock and requiring alignment to several synchronization contexts (and thus reference clocks) will effectively merge those into a single synchronization context.

3.1.1. RTCP CNAME

[RFC3550] describes inter-media synchronization between RTP sessions based on RTCP CNAME, RTP, and timestamps of a reference clock formatted using the Network Time Protocol (NTP) [RFC5905]. As indicated in [RFC7273], despite using NTP format timestamps, it is not required that the clock be synchronized to an NTP source.

3.1.2. Clock Source Signaling

[RFC7273] provides a mechanism to signal the clock source in the Session Description Protocol (SDP) [RFC4566] both for the reference clock as well as the media clock, thus allowing a synchronization context to be defined beyond the one defined by the usage of CNAME source descriptions.

3.1.3. Implicitly via RtcMediaStream

WebRTC defines RtcMediaStream with one or more RtcMediaStreamTracks. All tracks in a RtcMediaStream are intended to be synchronized when rendered, implying that they must be generated such that synchronization is possible.

3.1.4. Explicitly via SDP Mechanisms

The SDP Grouping Framework [RFC5888] defines an "m=" line (Section 4.2) grouping mechanism called Lip Synchronization (with LS identification-tag) for establishing the synchronization requirement across "m=" lines when they map to individual sources.

Source-Specific Media Attributes in SDP [RFC5576] extends the above mechanism when multiple media sources are described by a single "m=" line.

3.2. Endpoint

Some applications require knowledge of what media sources originate from a particular endpoint (Section 2.2.1). This can include such decisions as packet routing between parts of the topology, knowing the endpoint origin of the RTP streams.

In RTP, this identification has been overloaded with the synchronization context (Section 3.1) through the usage of the RTCP source description CNAME (Section 3.1.1). This works for some usages, but in others it breaks down. For example, if an endpoint has two sets of media sources that have different synchronization contexts, like the audio and video of the human participant as well as a set of media sources of audio and video for a shared movie, CNAME would not be an appropriate identification for that endpoint. Therefore, an endpoint may have multiple CNAMEs. The CNAMEs or the media sources themselves can be related to the endpoint.

3.3. Participant

In communication scenarios, information about which media sources originate from which participant (Section 2.2.3) is commonly needed. One reason is, for example, to enable the application to correctly display participant identity information associated with the media sources. This association is handled through signaling to point at a specific multimedia session where the media sources may be explicitly or implicitly tied to a particular endpoint.

Participant information becomes more problematic when there are media sources that are generated through mixing or other conceptual processing of raw streams or source streams that originate from different participants. These types of media sources can thus have a dynamically varying set of origins and participants. RTP contains the concept of CSRC that carries information about the previous step origin of the included media content on the RTP level.

3.4. RtcMediaStream

An RtcMediaStream in WebRTC is an explicit grouping of a set of media sources (RtcMediaStreamTracks) that share a common identifier and a single synchronization context (Section 3.1).

3.5. Multi-Channel Audio

There exist a number of RTP payload formats that can carry multi-channel audio, despite the codec being a single-channel (mono) encoder. Multi-channel audio can be viewed as multiple media sources sharing a common synchronization context. These are independently encoded by a media encoder and the different encoded streams are packetized together in a time-synchronized way into a single source RTP stream, using the used codec's RTP payload format. Examples of codecs that support multi-channel audio are PCMA and PCMU [RFC3551], Adaptive Multi Rate (AMR) [RFC4867], and G.719 [RFC5404].

3.6. Simulcast

A media source represented as multiple independent encoded streams constitutes a simulcast [SDP-SIMULCAST] or Modification Detection Code (MDC) of that media source. Figure 8 shows an example of a media source that is encoded into three separate simulcast streams, that are in turn sent on the same media transport flow. When using simulcast, the RTP streams may be sharing an RTP session and media transport, or be separated on different RTP sessions and media transports, or be any combination of these two. One major reason to use separate media transports is to make use of different quality of service (QoS) for the different source RTP streams. Some considerations on separating related RTP streams are discussed in Section 3.12.

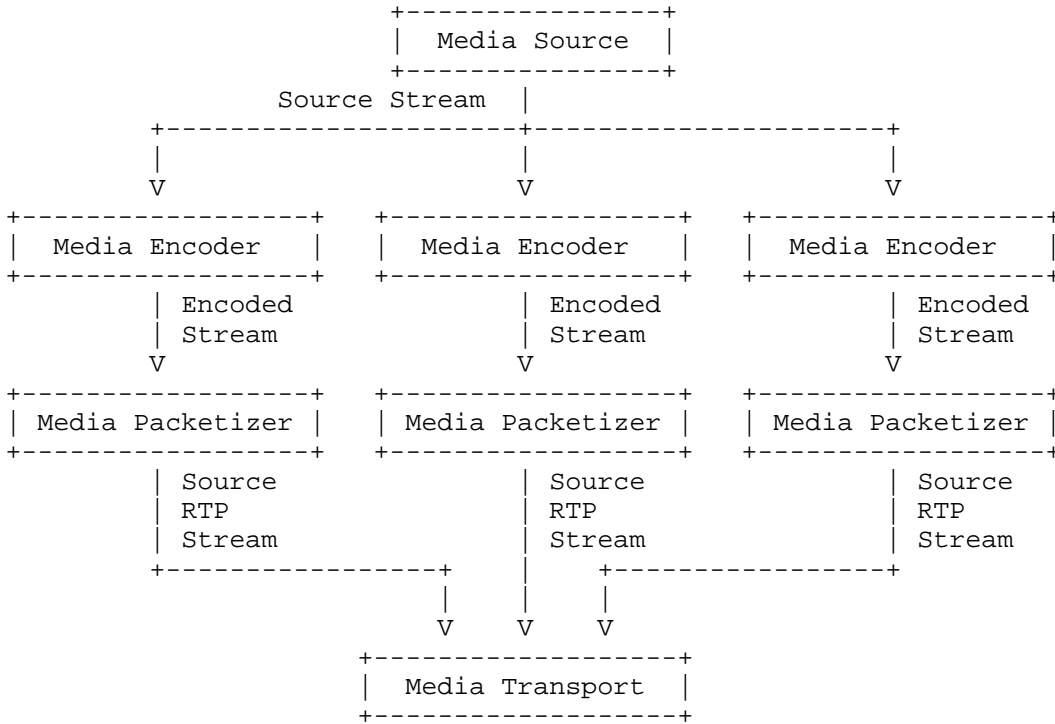


Figure 8: Example of Media Source Simulcast

The simulcast relation between the RTP streams is the common media source. In addition, to be able to identify the common media source, a receiver of the RTP stream may need to know which configuration or encoding goals lay behind the produced encoded stream and its properties. This enables selection of the stream that is most useful in the application at that moment.

3.7. Layered Multi-Stream

Layered Multi-Stream (LMS) is a mechanism by which different portions of a layered or scalable encoding of a source stream are sent using separate RTP streams (sometimes in separate RTP sessions). LMSs are useful for receiver control of layered media.

A media source represented as an encoded stream and multiple dependent streams constitutes a media source that has layered dependencies. Figure 9 represents an example of a media source that is encoded into three dependent layers, where two layers are sent on the same media transport using different RTP streams, i.e., SSRCs, and the third layer is sent on a separate media transport.

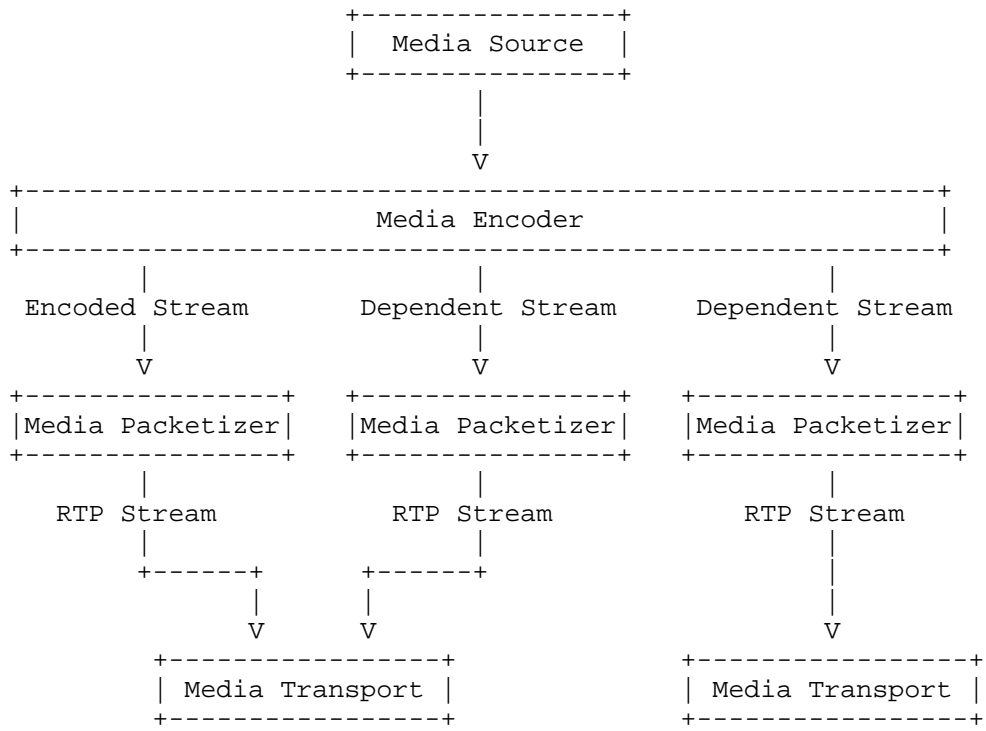


Figure 9: Example of Media Source Layered Dependency

It is sometimes useful to make a distinction between using a single media transport or multiple separate media transports when (in both cases) using multiple RTP streams to carry encoded streams and dependent streams for a media source. Therefore, the following new terminology is defined here:

SRST: Single RTP stream on a Single media Transport

MRST: Multiple RTP streams on a Single media Transport

MRMT: Multiple RTP streams on Multiple media Transports

MRST and MRMT relations need to identify the common media encoder origin for the encoded and dependent streams. When using different RTP sessions (MRMT), a single RTP stream per media encoder, and a single media source in each RTP session, common SSRCs and CNAMEs can be used to identify the common media source. When multiple RTP streams are sent from one media encoder in the same RTP session (MRST), then CNAME is the only currently specified RTP identifier that can be used. In cases where multiple media encoders use multiple media sources sharing synchronization context, and thus have a common CNAME, additional heuristics or identification need to be applied to create the MRST or MRMT relationships between the RTP streams.

3.8. RTP Stream Duplication

RTP Stream Duplication [RFC7198], using the same or different media transports, and optionally also delaying the duplicate [RFC7197], offers a simple way to protect media flows from packet loss in some cases (see Figure 10). This is a specific type of redundancy. All but one source RTP stream (Section 2.1.10) are effectively redundancy RTP streams (Section 2.1.12), but since both source and redundant RTP streams are the same, it does not matter which one is which. This can also be seen as a specific type of simulcast (Section 3.6) that transmits the same encoded stream (Section 2.1.7) multiple times.

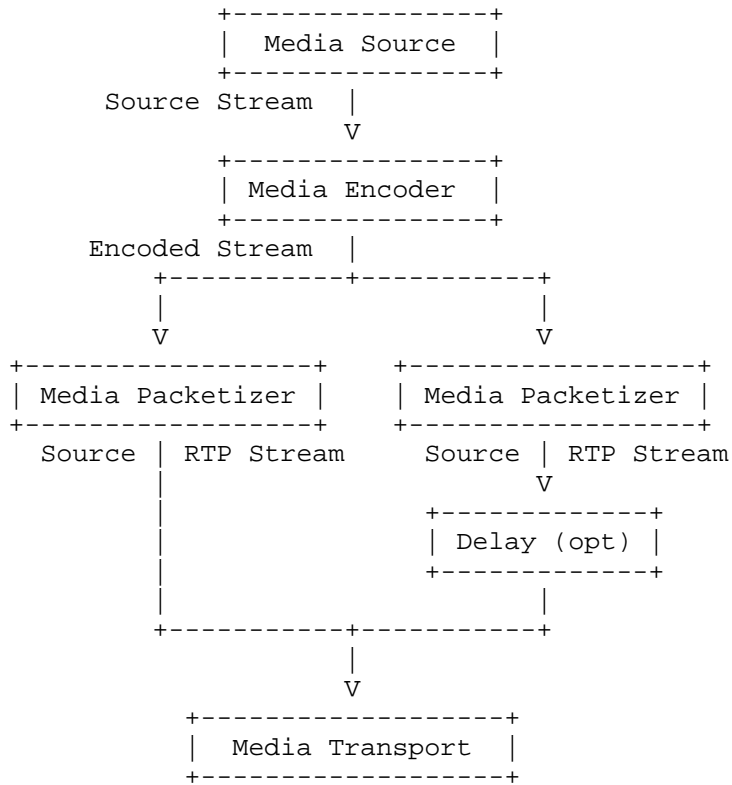


Figure 10: Example of RTP Stream Duplication

3.9. Redundancy Format

"RTP Payload for Redundant Audio Data" [RFC2198] defines a transport for redundant audio data together with primary data in the same RTP payload. The redundant data can be a time-delayed version of the primary or another time-delayed encoded stream using a different media encoder to encode the same media source as the primary, as depicted in Figure 11.

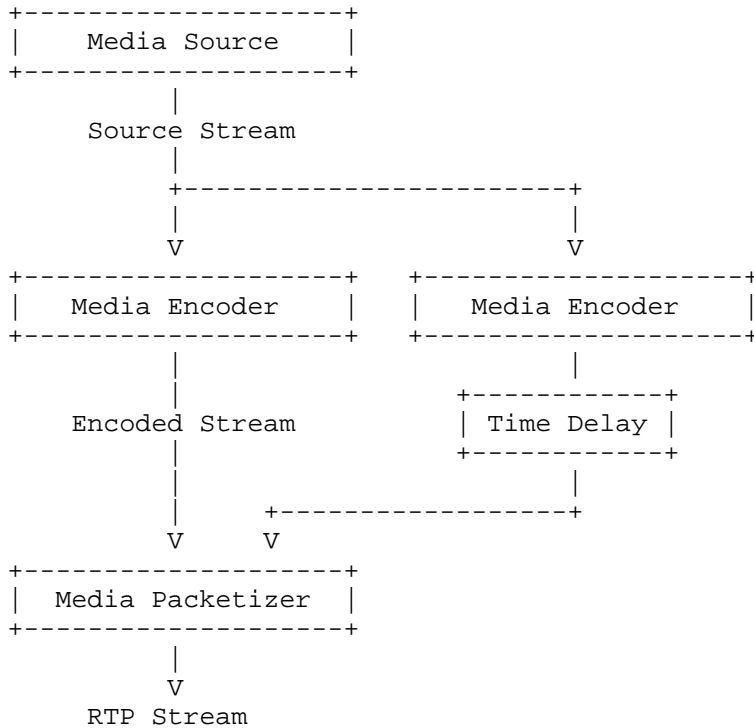


Figure 11: Concept for Usage of Audio Redundancy with Different Media Encoders

The redundancy format is thus providing the necessary meta information to correctly relate different parts of the same encoded stream. The case depicted above (Figure 11) relates the received source stream fragments coming out of different media decoders, to be able to combine them together into a less erroneous source stream.

3.10. RTP Retransmission

Figure 12 shows an example where a media source's source RTP stream is protected by a retransmission (RTX) flow [RFC4588]. In this

example, the source RTP stream and the redundancy RTP stream share the same media transport.

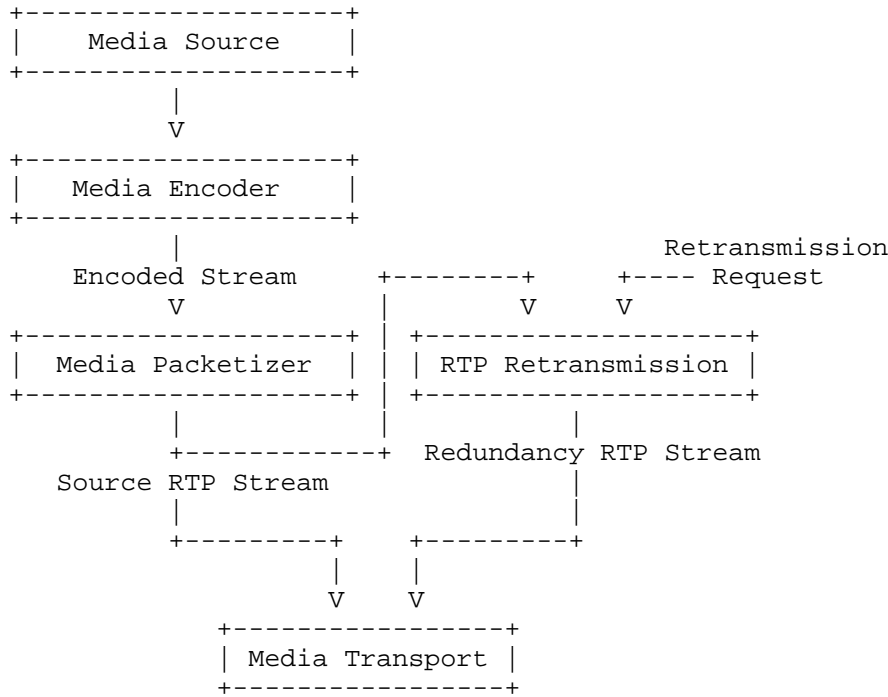


Figure 12: Example of Media Source Retransmission Flows

The RTP retransmission example (Figure 12) illustrates that this mechanism works purely on the source RTP stream. The RTP retransmission transforms buffers from the sent source RTP stream and, upon request, emits a retransmitted packet with an extra payload header as a redundancy RTP stream. The RTP retransmission mechanism [RFC4588] is specified such that there is a one-to-one relation between the source RTP stream and the redundancy RTP stream. Therefore, a redundancy RTP stream needs to be associated with its source RTP stream. This is done based on CNAME selectors and heuristics to match requested packets for a given source RTP stream with the original sequence number in the payload of any new redundancy RTP stream using the RTX payload format. In cases where the redundancy RTP stream is sent in a different RTP session than the source RTP stream, the RTP session relation is signaled by using the SDP media grouping's [RFC5888] Flow Identification (FID identification-tag) semantics.

As FEC encoding exists in various forms, the methods for relating FEC redundancy RTP streams with its source information in source RTP streams are many. The XOR-based RTP FEC payload format [RFC5109] is defined in such a way that a redundancy RTP stream has a one-to-one relation with a source RTP stream. In fact, the RFC requires the redundancy RTP stream to use the same SSRC as the source RTP stream. This requires the use of either a separate RTP session or the redundancy RTP payload format [RFC2198]. The underlying relation requirement for this FEC format and a particular redundancy RTP stream is to know the related source RTP stream, including its SSRC.

3.12. RTP Stream Separation

RTP streams can be separated exclusively based on their SSRCs, at the RTP session level, or at the multimedia session level.

When the RTP streams that have a relationship are all sent in the same RTP session and are uniquely identified based on their SSRC only, it is termed an "SSRC-only-based separation". Such streams can be related via RTCP CNAME to identify that the streams belong to the same endpoint. SSRC-based approaches [RFC5576], when used, can explicitly relate various such RTP streams.

On the other hand, when RTP streams that are related are sent in the context of different RTP sessions to achieve separation, it is known as "RTP session-based separation". This is commonly used when the different RTP streams are intended for different media transports.

Several mechanisms that use RTP session-based separation rely on it as a grouping mechanism expressing the relationship. The solutions have been based on using the same SSRC value in the different RTP sessions to implicitly indicate their relation. That way, no explicit RTP level mechanism has been needed; only signaling level relations have been established using semantics from the media-line grouping framework [RFC5888]. Examples of this are RTP retransmission [RFC4588], SVC Multi-Session Transmission [RFC6190], and XOR-based FEC [RFC5109]. RTCP CNAME explicitly relates RTP streams across different RTP sessions, as explained in the previous section. Such a relationship can be used to perform inter-media synchronization.

RTP streams that are related and need to be associated can be part of different multimedia sessions, rather than just different RTP sessions within the same multimedia session context. This puts further demand on the scope of the mechanism(s) and its handling of identifiers used for expressing the relationships.

3.13. Multiple RTP Sessions over one Media Transport

[TRANSPORT-MULTIPLEX] describes a mechanism that allows several RTP sessions to be carried over a single underlying media transport. The main reasons for doing this are related to the impact of using one or more media transports (using a common network path or potentially having different ones). The fewer media transports used, the less need for NAT/firewall traversal resources and smaller number of flow-based QoS.

However, multiple RTP sessions over one media transport imply that a single media transport 5-tuple is not sufficient to express in which RTP session context a particular RTP stream exists. Complexities in the relationship between media transports and RTP sessions already exist as one RTP session contains multiple media transports, e.g., even a Peer-to-Peer RTP Session with RTP/RTCP Multiplexing requires two media transports, one in each direction. The relationship between media transports and RTP sessions as well as additional levels of identifiers needs to be considered in both signaling design and when defining terminology.

4. Mapping from Existing Terms

This section describes a selected set of terms from some relevant RFCs and Internet-Drafts (at the time of writing), using the concepts from previous sections.

4.1. Telepresence Terms

The terms in this subsection are used in the context of CLUE [CLUE-FRAME]. Note that some terms listed in this subsection use the same names as terms defined elsewhere in this document. Unless explicitly stated (as "RTP Taxonomy") and in this subsection, they are to be read as references to the CLUE-specific term within this subsection.

4.1.1. Audio Capture

Defined in CLUE as a Media Capture (Section 4.1.7) for audio. Describes an audio media source (Section 2.1.4).

4.1.2. Capture Device

Defined in CLUE as a device that converts physical input into an electrical signal. Identifies a physical entity performing an RTP Taxonomy media capture (Section 2.1.2) transformation.

4.1.3. Capture Encoding

Defined in CLUE as a specific Encoding (Section 4.1.6) of a Media Capture (Section 4.1.7). Describes an encoded stream (Section 2.1.7) related to CLUE-specific semantic information.

4.1.4. Capture Scene

Defined in CLUE as a structure representing a spatial region captured by one or more Capture Devices (Section 4.1.2), each capturing media representing a portion of the region. Describes a set of spatially related media sources (Section 2.1.4).

4.1.5. Endpoint

Defined in CLUE as a CLUE-capable device that is the logical point of final termination through receiving, decoding, and rendering and/or initiation through capturing, encoding, and sending of media Streams (Section 4.1.10). CLUE further defines it to consist of one or more physical devices with source and sink media streams, and exactly one participant [RFC4353]. Describes exactly one participant (Section 2.2.3) and one or more RTP Taxonomy endpoints (Section 2.2.1).

4.1.6. Individual Encoding

Defined in CLUE as a set of parameters representing a way to encode a Media Capture (Section 4.1.7) to become a Capture Encoding (Section 4.1.3). Describes the configuration information needed to perform a media encoder (Section 2.1.6) transformation.

4.1.7. Media Capture

Defined in CLUE as a source of media, such as from one or more Capture Devices (Section 4.1.2) or constructed from other media Streams (Section 4.1.10). Describes either an RTP Taxonomy media capture (Section 2.1.2) or a media source (Section 2.1.4), depending on in which context the term is used.

4.1.8. Media Consumer

Defined in CLUE as a CLUE-capable device that intends to receive Capture Encodings (Section 4.1.3). Describes the media receiving part of an RTP Taxonomy endpoint (Section 2.2.1).

4.1.9. Media Provider

Defined in CLUE as a CLUE-capable device that intends to send Capture Encodings (Section 4.1.3). Describes the media sending part of an RTP Taxonomy endpoint (Section 2.2.1).

4.1.10. Stream

Defined in CLUE as a Capture Encoding (Section 4.1.3) sent from a Media Provider (Section 4.1.9) to a Media Consumer (Section 4.1.8) via RTP. Describes an RTP stream (Section 2.1.10).

4.1.11. Video Capture

Defined in CLUE as a Media Capture (Section 4.1.7) for video. Describes a video media source (Section 2.1.4).

4.2. Media Description

A single Session Description Protocol (SDP) [RFC4566] Media Description (or media block; an "m=" line and all subsequent lines until the next "m=" line or the end of the SDP) describes part of the necessary configuration and identification information needed for a media encoder transformation, as well as the necessary configuration and identification information for the media decoder to be able to correctly interpret a received RTP stream.

A media description typically relates to a single media source. This is, for example, an explicit restriction in WebRTC. However, nothing prevents that the same media description (and same RTP session) is reused for multiple media sources [RTP-MULTI-STREAM]. It can thus describe properties of one or more RTP streams, and can also describe properties valid for an entire RTP session (via [RFC5576] mechanisms, for example).

4.3. Media Stream

RTP [RFC3550] uses media stream, audio stream, video stream, and a stream of (RTP) packets interchangeably, which are all RTP streams.

4.4. Multimedia Conference

A Multimedia Conference is a communication session (Section 2.2.5) between two or more participants (Section 2.2.3), along with the software they are using to communicate.

4.5. Multimedia Session

SDP [RFC4566] defines a multimedia session as a set of multimedia senders and receivers and the data streams flowing from senders to receivers, which would correspond to a set of endpoints and the RTP streams that flow between them. In this document, multimedia session (Section 2.2.4) also assumes those endpoints belong to a set of participants that are engaged in communication via a set of related RTP streams.

RTP [RFC3550] defines a multimedia session as a set of concurrent RTP sessions among a common group of participants. For example, a video conference may contain an audio RTP session and a video RTP session. This would correspond to a group of participants (each using one or more endpoints) sharing a set of concurrent RTP sessions. In this document, multimedia session also defines those RTP sessions to have some relation and be part of a communication among the participants.

4.6. Multipoint Control Unit (MCU)

This term is commonly used to describe the central node in any type of star topology [RTP-TOPOLOGIES] conference. It describes a device that includes one participant (Section 2.2.3) (usually corresponding to a so-called conference focus) and one or more related endpoints (Section 2.2.1) (sometimes one or more per conference participant).

4.7. Multi-Session Transmission (MST)

One of two transmission modes defined in H.264-based SVC [RFC6190], the other mode being a Single-Session Transmission (SST) (Section 4.14). In Multi-Session Transmission (MST), the SVC media encoder sends encoded streams and dependent streams distributed across two or more RTP streams in one or more RTP sessions. The term "MST" is ambiguous in RFC 6190, especially since the name indicates the use of multiple "sessions", while MST-type packetization is in fact required whenever two or more RTP streams are used for the encoded and dependent streams, regardless if those are sent in one or more RTP sessions. Corresponds either to MRST or MRMT (Section 3.7) stream relations defined in this document. The SVC RTP payload RFC [RFC6190] is not particularly explicit about how the common media encoder (Section 2.1.6) relation between encoded streams (Section 2.1.7) and dependent streams (Section 2.1.8) is to be implemented.

4.8. Recording Device

WebRTC specifications use this term to refer to locally available entities performing a media capture (Section 2.1.2) transformation.

4.9. RtcMediaStream

A WebRTC RtcMediaStream is a set of media sources (Section 2.1.4) sharing the same synchronization context (Section 3.1).

4.10. RtcMediaStreamTrack

A WebRTC RtcMediaStreamTrack is a media source (Section 2.1.4).

4.11. RTP Receiver

RTP [RFC3550] uses this term, which can be seen as the RTP protocol part of a media depacketizer (Section 2.1.27).

4.12. RTP Sender

RTP [RFC3550] uses this term, which can be seen as the RTP protocol part of a media packetizer (Section 2.1.9).

4.13. RTP Session

Within the context of SDP, a single "m=" line can map to a single RTP session (Section 2.2.2), or multiple "m=" lines can map to a single RTP session. The latter is enabled via multiplexing schemes such as BUNDLE [SDP-BUNDLE], for example, which allows mapping of multiple "m=" lines to a single RTP session.

4.14. Single-Session Transmission (SST)

One of two transmission modes defined in H.264-based SVC [RFC6190], the other mode being MST (Section 4.7). In SST, the SVC media encoder sends encoded streams (Section 2.1.7) and dependent streams (Section 2.1.8) combined into a single RTP stream (Section 2.1.10) in a single RTP session (Section 2.2.2), using the SVC RTP payload format. The term "SST" is ambiguous in RFC 6190, in that it sometimes refers to the use of a single RTP stream, like in sections relating to packetization, and sometimes appears to refer to use of a single RTP session, like in the context of discussing SDP. Closely corresponds to SRST (Section 3.7) defined in this document.

4.15. SSRC

RTP [RFC3550] defines this as "the source of a stream of RTP packets", which indicates that an SSRC is not only a unique identifier for the encoded stream (Section 2.1.7) carried in those packets but is also effectively used as a term to denote a media packetizer (Section 2.1.9). In [RFC3550], it is stated that "a synchronization source may change its data format, e.g., audio encoding, over time". The related encoded stream data format in an RTP stream (Section 2.1.10) is identified by the RTP payload type. Changing the data format for an encoded stream effectively also changes what media encoder (Section 2.1.6) is used for the encoded stream. No ambiguity is introduced to SSRC as an encoded stream identifier by allowing RTP payload type changes, as long as only a single RTP payload type is valid for any given RTP Timestamp. This is aligned with and further described by Section 5.2 of [RFC3550].

5. Security Considerations

The purpose of this document is to make clarifications and reduce the confusion prevalent in RTP taxonomy because of inconsistent usage by multiple technologies and protocols making use of the RTP protocol. It does not introduce any new security considerations beyond those already well documented in the RTP protocol [RFC3550] and each of the many respective specifications of the various protocols making use of it.

Having a well-defined common terminology and understanding of the complexities of the RTP architecture will help lead us to better standards, avoiding security problems.

6. Informative References

[CLUE-FRAME]

Duckworth, M., Pepperell, A., and S. Wenger, "Framework for Telepresence Multi-Streams", Work in Progress, draft-ietf-clue-framework-22, April 2015.

[RFC2198] Perkins, C., Kouvelas, I., Hodson, O., Hardman, V., Handley, M., Bolot, J., Vega-Garcia, A., and S. Fosse-Parisis, "RTP Payload for Redundant Audio Data", RFC 2198, DOI 10.17487/RFC2198, September 1997, <<http://www.rfc-editor.org/info/rfc2198>>.

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, DOI 10.17487/RFC3550, July 2003, <<http://www.rfc-editor.org/info/rfc3550>>.

- [RFC3551] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, DOI 10.17487/RFC3551, July 2003, <<http://www.rfc-editor.org/info/rfc3551>>.
- [RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, DOI 10.17487/RFC3711, March 2004, <<http://www.rfc-editor.org/info/rfc3711>>.
- [RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, DOI 10.17487/RFC4353, February 2006, <<http://www.rfc-editor.org/info/rfc4353>>.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <<http://www.rfc-editor.org/info/rfc4566>>.
- [RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R. Hakenberg, "RTP Retransmission Payload Format", RFC 4588, DOI 10.17487/RFC4588, July 2006, <<http://www.rfc-editor.org/info/rfc4588>>.
- [RFC4867] Sjöberg, J., Westerlund, M., Lankaniemi, A., and Q. Xie, "RTP Payload Format and File Storage Format for the Adaptive Multi-Rate (AMR) and Adaptive Multi-Rate Wideband (AMR-WB) Audio Codecs", RFC 4867, DOI 10.17487/RFC4867, April 2007, <<http://www.rfc-editor.org/info/rfc4867>>.
- [RFC5109] Li, A., Ed., "RTP Payload Format for Generic Forward Error Correction", RFC 5109, DOI 10.17487/RFC5109, December 2007, <<http://www.rfc-editor.org/info/rfc5109>>.
- [RFC5404] Westerlund, M. and I. Johansson, "RTP Payload Format for G.719", RFC 5404, DOI 10.17487/RFC5404, January 2009, <<http://www.rfc-editor.org/info/rfc5404>>.
- [RFC5481] Morton, A. and B. Claise, "Packet Delay Variation Applicability Statement", RFC 5481, DOI 10.17487/RFC5481, March 2009, <<http://www.rfc-editor.org/info/rfc5481>>.
- [RFC5576] Lennox, J., Ott, J., and T. Schierl, "Source-Specific Media Attributes in the Session Description Protocol (SDP)", RFC 5576, DOI 10.17487/RFC5576, June 2009, <<http://www.rfc-editor.org/info/rfc5576>>.

- [RFC5888] Camarillo, G. and H. Schulzrinne, "The Session Description Protocol (SDP) Grouping Framework", RFC 5888, DOI 10.17487/RFC5888, June 2010, <<http://www.rfc-editor.org/info/rfc5888>>.
- [RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <<http://www.rfc-editor.org/info/rfc5905>>.
- [RFC6190] Wenger, S., Wang, Y., Schierl, T., and A. Eleftheriadis, "RTP Payload Format for Scalable Video Coding", RFC 6190, DOI 10.17487/RFC6190, May 2011, <<http://www.rfc-editor.org/info/rfc6190>>.
- [RFC7160] Petit-Huguenin, M. and G. Zorn, Ed., "Support for Multiple Clock Rates in an RTP Session", RFC 7160, DOI 10.17487/RFC7160, April 2014, <<http://www.rfc-editor.org/info/rfc7160>>.
- [RFC7197] Begen, A., Cai, Y., and H. Ou, "Duplication Delay Attribute in the Session Description Protocol", RFC 7197, DOI 10.17487/RFC7197, April 2014, <<http://www.rfc-editor.org/info/rfc7197>>.
- [RFC7198] Begen, A. and C. Perkins, "Duplicating RTP Streams", RFC 7198, DOI 10.17487/RFC7198, April 2014, <<http://www.rfc-editor.org/info/rfc7198>>.
- [RFC7201] Westerlund, M. and C. Perkins, "Options for Securing RTP Sessions", RFC 7201, DOI 10.17487/RFC7201, April 2014, <<http://www.rfc-editor.org/info/rfc7201>>.
- [RFC7273] Williams, A., Gross, K., van Brandenburg, R., and H. Stokking, "RTP Clock Source Signalling", RFC 7273, DOI 10.17487/RFC7273, June 2014, <<http://www.rfc-editor.org/info/rfc7273>>.
- [RTP-MULTI-STREAM]
Lennox, J., Westerlund, M., Wu, W., and C. Perkins, "Sending Multiple Media Streams in a Single RTP Session", Work in Progress, draft-ietf-avtcore-rtp-multi-stream-08, July 2015.
- [RTP-TOPOLOGIES]
Westerlund, M. and S. Wenger, "RTP Topologies", Work in Progress, draft-ietf-avtcore-rtp-topologies-update-10, July 2015.

[SDP-BUNDLE]

Holmberg, C., Alvestrand, H., and C. Jennings,
"Negotiating Media Multiplexing Using the Session
Description Protocol (SDP)", Work in Progress,
draft-ietf-mmusic-sdp-bundle-negotiation-23, July 2015.

[SDP-SIMULCAST]

Burman, B., Westerlund, M., Nandakumar, S., and M. Zanaty,
"Using Simulcast in SDP and RTP Sessions", Work in
Progress, draft-ietf-mmusic-sdp-simulcast-01, July 2015.

[TRANSPORT-MULTIPLEX]

Westerlund, M. and C. Perkins, "Multiplexing Multiple RTP
Sessions onto a Single Lower-Layer Transport", Work in
Progress, draft-westerlund-avtcore-transport-multiplexing-
07, October 2013.

[WEBRTC-OVERVIEW]

Alvestrand, H., "Overview: Real Time Protocols for
Browser-based Applications", Work in Progress,
draft-ietf-rtcweb-overview-14, June 2015.

Acknowledgements

This document has many concepts borrowed from several documents such as WebRTC [WEBRTC-OVERVIEW], CLUE [CLUE-FRAME], and Multiplexing Architecture [TRANSPORT-MULTIPLEX]. The authors would like to thank all the authors of each of those documents.

The authors would also like to acknowledge the insights, guidance, and contributions of Magnus Westerlund, Roni Even, Paul Kyzivat, Colin Perkins, Keith Drage, Harald Alvestrand, Alex Eleftheriadis, Mo Zanaty, Stephan Wenger, and Bernard Aboba.

Contributors

Magnus Westerlund has contributed the concept model for the media chain using transformations and streams model, including rewriting pre-existing concepts into this model and adding missing concepts. The first proposal for updating the relationships and the topologies based on this concept was also performed by Magnus.

Authors' Addresses

Jonathan Lennox
Vidyo, Inc.
433 Hackensack Avenue
Seventh Floor
Hackensack, NJ 07601
United States

Email: jonathan@vidyo.com

Kevin Gross
AVA Networks, LLC
Boulder, CO
United States

Email: kevin.gross@avanw.com

Suhas Nandakumar
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
United States

Email: snandaku@cisco.com

Gonzalo Salgueiro
Cisco Systems
7200-12 Kit Creek Road
Research Triangle Park, NC 27709
United States

Email: gsalguei@cisco.com

Bo Burman (editor)
Ericsson
Kistavagen 25
SE-16480 Stockholm
Sweden

Email: bo.burman@ericsson.com

