

Workgroup: AVTCore
Internet-Draft:
draft-ietf-avtcore-multi-party-rtt-mix-10
Updates: [RFC 4103](#) (if approved)
Published: 18 November 2020
Intended Status: Standards Track
Expires: 22 May 2021
Authors: G. Hellstrom
Gunnar Hellstrom Accessible Communication
RTP-mixer formatting of multi-party Real-time text

Abstract

Real-time text mixers for multi-party sessions need to identify the source of each transmitted group of text so that the text can be presented by endpoints in suitable grouping with other text from the same source, while new text from other sources is also presented in readable grouping as received interleaved in real-time.

Regional regulatory requirements specify provision of real-time text in multi-party calls. RFC 4103 mixer implementations can use traditional RTP functions for source identification, but the mixer source switching performance is limited when using the default transmission characteristics with redundancy.

Enhancements for RFC 4103 real-time text mixing is provided in this document, suitable for a centralized conference model that enables source identification and source switching. The intended use is for real-time text mixers and multi-party-aware participant endpoints. The specified mechanism build on the standard use of the CSRC list in the RTP packet for source identification. The method makes use of the same "text/t140" and "text/red" formats as for two-party sessions.

A capability exchange is specified so that it can be verified that a participant can handle the multi-party coded real-time text stream. The capability is indicated by use of a media attribute "rtt-mixer".

The document updates RFC 4103[[RFC4103](#)]

A specifications of how a mixer can format text for the case when the endpoint is not multi-party aware is also provided.

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 <https://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 22 May 2021.

Copyright Notice

Copyright (c) 2020 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 (<https://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](#)
 - [1.1. Selected solution and considered alternative](#)
 - [1.2. Nomenclature](#)
 - [1.3. Intended application](#)
- [2. Overview over the two specified solutions](#)
 - [2.1. Negotiated use of the RFC 4103 format for multi-party transmission in a single RTP stream](#)
 - [2.2. Mixing for multi-party unaware endpoints](#)
- [3. Details for the multi-party aware mixing case](#)
 - [3.1. Offer/answer considerations](#)
 - [3.2. Actions depending on capability negotiation result](#)
 - [3.3. Use of fields in the RTP packets](#)
 - [3.4. Initial transmission of a BOM character](#)
 - [3.5. Keep-alive](#)
 - [3.6. Transmission interval](#)
 - [3.7. Only one source per packet](#)
 - [3.8. Do not send received text to the originating source](#)
 - [3.9. Clean incoming text](#)
 - [3.10. Redundancy](#)
 - [3.11. Source switching](#)

- [3.12. Text placement in packets](#)
- [3.13. Empty T140blocks](#)
- [3.14. Creation of the redundancy](#)
- [3.15. Timer offset fields](#)
- [3.16. Other RTP header fields](#)
- [3.17. Pause in transmission](#)
- [3.18. RTCP considerations](#)
- [3.19. Reception of multi-party contents](#)
- [3.20. Performance considerations](#)
- [3.21. Security for session control and media](#)
- [3.22. SDP offer/answer examples](#)
- [3.23. Packet sequence example from a source switch](#)
- [3.24. Maximum character rate "CPS"](#)
- [4. Presentation level considerations](#)
 - [4.1. Presentation by multi-party aware endpoints](#)
 - [4.2. Multi-party mixing for multi-party unaware endpoints](#)
- [5. Relation to Conference Control](#)
 - [5.1. Use with SIP centralized conferencing framework](#)
 - [5.2. Conference control](#)
- [6. Gateway Considerations](#)
 - [6.1. Gateway considerations with Textphones \(e.g. TTYs\).](#)
 - [6.2. Gateway considerations with WebRTC.](#)
- [7. Updates to RFC 4103](#)
- [8. Congestion considerations](#)
- [9. Acknowledgements](#)
- [10. IANA Considerations](#)
 - [10.1. Registration of the "rtt-mixer" sdp media attribute](#)
- [11. Security Considerations](#)
- [12. Change history](#)
 - [12.1. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-10](#)
 - [12.2. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-09](#)
 - [12.3. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-08](#)
 - [12.4. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-07](#)
 - [12.5. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-06](#)
 - [12.6. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-05](#)
 - [12.7. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-04](#)
 - [12.8. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-03](#)
 - [12.9. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-02](#)
 - [12.10. Changes to draft-ietf-avtcore-multi-party-rtt-mix-01](#)

[12.11. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-03 to draft-ietf-avtcore-multi-party-rtt-mix-00](#)
[12.12. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-02 to -03](#)
[12.13. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-01 to -02](#)
[12.14. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-00 to -01](#)
[13. References](#)
[13.1. Normative References](#)
[13.2. Informative References](#)
[Author's Address](#)

1. Introduction

RFC 4103 [[RFC4103](#)] specifies use of RFC 3550 RTP [[RFC3550](#)] for transmission of real-time text (RTT) and the "text/t140" format. It also specifies a redundancy format "text/red" for increased robustness. RFC 4102 [[RFC4102](#)] registers the "text/red" format. Regional regulatory requirements specify provision of real-time text in multi-party calls.

Real-time text is usually provided together with audio and sometimes with video in conversational sessions.

A requirement related to multi-party sessions from the presentation level standard T.140 for real-time text is: "The display of text from the members of the conversation should be arranged so that the text from each participant is clearly readable, and its source and the relative timing of entered text is visualized in the display."

Another requirement is that the mixing procedure must not introduce delays in the text streams that are experienced disturbing the real-time experience of the receiving users.

The redundancy scheme of RFC 4103 [[RFC4103](#)] enables efficient transmission of redundant text in packets together with new text. However the redundancy header format has no source indicators for the redundant transmissions. The redundant parts in a packet must therefore be from the same source as the new text. The recommended transmission is one new and two redundant generations of text (T140blocks) in each packet and the recommended transmission interval for two-party use is 300 ms.

Real-time text mixers for multi-party sessions therefore need to insert the source of each transmitted group of text from a conference participant so that the text can be transmitted interleaved with text groups from different sources in the rate they are created. This enables the text groups to be presented by

endpoints in suitable grouping with other text from the same source. The presentation can then be arranged so that text from different sources can be presented in real-time and easily read while it is possible for a reading user to also perceive approximately when the text was created in real time by the different parties. The transmission and mixing is intended to be done in a general way so that presentation can be arranged in a layout decided by the endpoint.

There are existing implementations of RFC 4103 without the updates from this document. These will not be able to receive and present real-time text mixed for multi-party aware endpoints.

A negotiation mechanism is therefore needed for verification if the parties are able to handle a multi-party coded stream and agreeing on using that method.

A fall-back mixing procedure is also needed for cases when the negotiation result indicates that a receiving endpoint is not capable of handling the mixed format. This method is called the mixing procedure for multi-party unaware endpoints. The fall-back method is naturally not expected to meet all performance requirements placed on the mixing procedure for multi-party aware endpoints.

The document updates RFC 4103[[RFC4103](#)] by introducing an attribute for indicating capability for the multi-party mixing case and rules for source indications and source switching.

1.1. Selected solution and considered alternative

A number of alternatives were considered when searching an efficient and easily implemented multi-party method for real-time text. This section explains a few of them briefly.

One RTP stream per source, sent in the same RTP session with "text/red" format.

From some points of view, use of multiple RTP streams, one for each source, sent in the same RTP session, called the RTP translator model in [[RFC3550](#)], would be efficient, and use exactly the same packet format as [[RFC4103](#)], the same payload type and a simple SDP declaration. However, the RTP implementation in both mixers and endpoints need to support multiple streams in the same RTP session in order to use this mechanism. For best deployment opportunity, it should be possible to upgrade existing endpoint solutions to be multi-party aware with a reasonable effort. There is currently a lack of support for multi-stream RTP in certain implementation technologies. This

fact made this solution not selected for inclusion in this document.

The "text/red" format in RFC 4103 with shorter transmission interval, and indicating source in CSRC.

The "text/red" format with "text/t140" payload in a single RTP stream can be sent with 100 ms packet intervals instead of the regular 300 ms. The source is indicated in the CSRC field. Source switching can then be done every 100 ms while simultaneous transmission occurs. With five participants sending text simultaneously, the switching and transmission performance is good. With more simultaneously sending participants, there will be a noticeable jerkiness in text presentation. The jerkiness will be more expressed the more participants who send text simultaneously. With ten sending participants, the jerkiness will be about one second. Text sent from a source at the end of the period its text is sent by the mixer will have close to zero extra delay. Recent text will be presented with no or low delay. The one second jerkiness will be noticeable and slightly unpleasant, but corresponds in time to what typing humans often cause by hesitation or changing position while typing. A benefit of this method is that no new packet format needs to be introduced and implemented. Since simultaneous typing by more than two parties is very rare, and in most applications also more than three parties in a call is rare, this method can be used successfully with good performance. Recovery of text in case of packet loss is based on analysis of timestamps of received redundancy versus earlier received text. Negotiation is based on a new sdp media attribute "rtt-mixer". This method is selected to be the main one specified in this document.

A new "text" media subtype with up to 15 sources in each packet.

The mechanism makes use of the RTP mixer model specified in RFC3550[[RFC3550](#)]. Text from up to 15 sources can be included in each packet. Packets are normally sent every 300 ms. The mean delay will be 150 ms. The sources are indicated in strict order in the CSRC list of the RTP packets. A new redundancy packet format is specified. This method would result in good performance, but would require standardisation and implementation of new releases in the target technologies that would take more time than desirable to complete. It was therefore not selected to be included in this document.

The presentation planned by the mixer for multi-party unaware endpoints.

It is desirable to have a method that does not require any modifications in existing user devices implementing RFC 4103 for RTT without explicit support of multi-party sessions. This is possible by having the mixer insert a new line and a text

formatted source label before each switch of text source in the stream. Switch of source can only be done in places in the text where it does not disturb the perception of the contents. Text from only one source can be presented in real time at a time. The delay will therefore be varying. The method has also other limitations, but is included in this document as a fallback method. In calls where parties take turns properly by ending their entries with a new line, the limitations will have limited influence on the user experience. While only two parties send text, these two will see the text in real time with no delay. This method is specified as a fallback method in this document.

RTT transport in WebRTC

Transport of real-time text in the WebRTC technology is specified to use the WebRTC data channel in [[I-D.ietf-mmusic-t140-usage-data-channel](#)]. That specification contains a section briefly describing its use in multi-party sessions. The focus of this document is RTP transport. Therefore, even if the WebRTC transport provides good multi-party performance, it is just mentioned in this document in relation to providing gateways with multi-party capabilities between RTP and WebRTC technologies.

1.2. Nomenclature

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](#)].

The terms SDES, CNAME, NAME, SSRC, CSRC, CSRC list, CC, RTCP, RTP-mixer, RTP-translator are explained in [[RFC3550](#)].

The term "T140block" is defined in RFC 4103 [[RFC4103](#)] to contain one or more T.140 code elements.

"TTY" stands for a text telephone type used in North America.

"WebRTC" stands for web based communication specified by W3C and IETF.

"DTLS-SRTP" stands for security specified in RFC 5764 [[RFC5764](#)].

"multi-party aware" stands for an endpoint receiving real-time text from multiple sources through a common conference mixer being able to present the text in real-time separated by source and presented so that a user can get an impression of the approximate relative timing of text from different parties.

"multi-party unaware" stands for an endpoint not itself being able to separate text from different sources when received through a common conference mixer.

1.3. Intended application

The method for multi-party real-time text specified in this document is primarily intended for use in transmission between mixers and endpoints in centralised mixing configurations. It is also applicable between mixers. An often mentioned application is for emergency service calls with real-time text and voice, where a calltaker want to make an attended handover of a call to another agent, and stay observing the session. Multimedia conference sessions with support for participants to contribute in text is another application. Conferences with central support for speech-to-text conversion is yet another mentioned application.

In all these applications, normally only one participant at a time will send long text utterances. In some cases, one other participant will occasionally contribute with a longer comment simultaneously. That may also happen in some rare cases when text is interpreted to text in another language in a conference. Apart from these cases, other participants are only expected to contribute with very brief utterings while others are sending text.

Text is supposed to be human generated, by some text input means, such as typing on a keyboard or using speech-to-text technology. Occasional small cut-and-paste operations may appear even if that is not the initial purpose of real-time text.

The real-time characteristics of real-time text is essential for the participants to be able to contribute to a conversation. If the text is too much delayed from typing a letter to its presentation, then, in some conference situations, the opportunity to comment will be gone and someone else will grab the turn. A delay of more than one second in such situations is an obstacle for good conversation.

2. Overview over the two specified solutions

This section contains a brief introduction of the two methods specified in this document.

2.1. Negotiated use of the RFC 4103 format for multi-party transmission in a single RTP stream

The main purpose of this document is to specify a method for true multi-party real-time text mixing for multi-party aware endpoints. The method use of the current format for real-time text in [\[RFC4103\]](#). It is an update of RFC 4103 by a clarification on one way to use it in the multi-party situation. It is done by completing a negotiation for this kind of multi-party capability and by indicating source in the CSRC element in the RTP packets. Specific considerations are made to be able to recover text after packet loss.

The detailed procedures for the multi-party aware case are specified in [Section 3](#)

Please use [[RFC4103](#)] as reference when reading the specification.

2.2. Mixing for multi-party unaware endpoints

A method is also specified in this document for cases when the endpoint participating in a multi-party call does not itself implement any solution for multi-party handling of real-time text. The solution requires the mixer to insert text dividers and readable labels and only send text from one source at a time until a suitable point appears for source change. This solution is a fallback method with functional limitations that acts on the presentation level.

A party performing as a mixer, which has not negotiated the "rtt-mixer" sdp media attribute, but negotiated a "text/red" or "text/t140" format in a session with a participant SHOULD, if nothing else is specified for the application, format transmitted text to that participant to be suitable to present on a multi-party unaware endpoint as further specified in [Section 4.2](#).

3. Details for the multi-party aware mixing case

3.1. Offer/answer considerations

RFC 4103[[RFC4103](#)] specifies use of RFC 3550 RTP[[RFC3550](#)], and a redundancy format "text/red" for increased robustness of real-time text transmission. This document updates RFC 4103[[RFC4103](#)] by introducing a capability negotiation for handling multi-party real-time text, a way to indicate the source of transmitted text, and rules for efficient timing of the transmissions.

The capability negotiation is based on use of the sdp media attribute "rtt-mixer".

Both parties shall indicate their capability in a session setup or modification, and evaluate the capability of the counterpart.

The syntax is as follows:

"a=rtt-mixer"

3.2. Actions depending on capability negotiation result

A transmitting party SHALL send text according to the multi-party format only when the negotiation for this method was successful and when the CC field in the RTP packet is set to 1. In all other cases, the packets SHALL be populated and interpreted as for a two-party session.

A party which has negotiated the "rtt-mixer" sdp media attribute MUST populate the CSRC-list and format the packets according to [Section 3](#) if it acts as an rtp-mixer and sends multi-party text.

A party which has negotiated the "rtt-mixer" sdp media attribute MUST interpret the contents of the "CC" field the CSRC-list and the packets according to [Section 3](#) in received rtp packets in the corresponding RTP stream.

A party not performing as a mixer MUST not include the CSRC list.

3.3. Use of fields in the RTP packets

The CC field SHALL show the number of members in the CSRC list, which SHALL be one (1) in transmissions from a mixer involved in a multi-party session, and otherwise 0.

When transmitted from a mixer during a multi-party session, a CSRC list SHALL be included in the packet. The single member in the CSRC-list SHALL contain the SSRC of the source of the T140blocks in the packet. When redundancy is used, the recommended level of redundancy is to use one primary and two redundant generations of T140blocks. In some cases, a primary or redundant T140block is empty, but is still represented by a member in the redundancy header.

From other aspects, the contents of the RTP packets are equal to what is specified in [[RFC4103](#)].

3.4. Initial transmission of a BOM character

As soon as a participant is known to participate in a session with another entity and being available for text reception, a Unicode BOM character SHALL be sent to it by the other entity according to the procedures in this section. If the transmitter is a mixer, then the source of this character SHALL be indicated to be the mixer itself.

Note that the BOM character SHALL be transmitted with the same redundancy procedures as any other text.

3.5. Keep-alive

After that, the transmitter SHALL send keep-alive traffic to the receiver(s) at regular intervals when no other traffic has occurred during that interval, if that is decided for the actual connection. Recommendations for keep-alive can be found in [[RFC6263](#)].

3.6. Transmission interval

A "text/red" transmitter in a mixer SHOULD send packets distributed in time as long as there is something (new or redundant T140blocks)

to transmit. The maximum transmission interval SHOULD then be 330 ms. It is RECOMMENDED to send next packet to a receiver as soon as new text to that receiver is available, as long as the time after the latest sent packet to the same receiver is more than or equal to 100 ms, and also the maximum character rate to the receiver is not exceeded. The intention is to keep the latency low and network load limited while keeping a good protection against text loss in bursty packet loss conditions.

3.7. Only one source per packet

New and redundant text from one source SHALL be transmitted in the same packet if available for transmission at the same time. Text from different sources MUST NOT be transmitted in the same packet.

3.8. Do not send received text to the originating source

Text received to a mixer from a participant SHOULD NOT be included in transmission from the mixer to that participant.

3.9. Clean incoming text

A mixer SHALL handle reception, recovery of packet loss, deletion of superfluous redundancy, marking of possible text loss and deletion of 'BOM' characters from each participant before queueing received text for transmission to receiving participants.

3.10. Redundancy

A transmitting party using redundancy SHALL send redundant repetitions of T140blocks already transmitted in earlier packets.

The number of redundant generations of T140blocks to include in transmitted packets SHALL be deduced from the SDP negotiation. It SHOULD be set to the minimum of the number declared by the two parties negotiating a connection. It is RECOMMENDED to transmit one original and two redundant generations of the T140blocks.

3.11. Source switching

When text from more than one source is available for transmission from a mixer, the mixer SHALL let the sources take turns in having their text transmitted.

The source with the oldest received text in the mixer SHOULD be next in turn to get all its available unsent text transmitted.

3.12. Text placement in packets

The mixer SHOULD compose and transmit an RTP packet to a receiver when one of the following conditions has occurred:

- *100 ms has passed since the latest transmission to that receiver, and there is unsent text available for transmission.

- *New text has arrived and more than 100 ms has passed since latest transmission to that receiver.

- *330 ms has passed since already transmitted text was queued for transmission as redundant text, and more than 100 ms has passed since the latest transmission to that receiver, and the redundant text is still not sent.

At time of transmission, the mixer SHALL populate the RTP packet with all T140blocks queued for transmission originating from the source in turn for transmission as long as this is not in conflict with the allowed number of characters per second ("CPS") or the maximum packet size. In this way, the latency of the latest received text is kept low even in moments of simultaneous transmission from many sources.

The SSRC of the source shall be placed as the only member in the CSRC-list.

Note: The CSRC-list in an RTP packet only includes the participant who's text is included in text blocks. It is not the same as the total list of participants in a conference. With audio and video media, the CSRC-list would often contain all participants who are not muted whereas text participants that don't type are completely silent and thus are not represented in RTP packet CSRC-lists.

3.13. Empty T140blocks

If no unsent T140blocks were available for a source at the time of populating a packet, but T140blocks are available which have not yet been sent the full intended number of redundant transmissions, then the primary T140block for that source is composed of an empty T140block, and populated (without taking up any length) in a packet for transmission. The corresponding SSRC SHALL be placed as usual in its place in the CSRC-list.

The first packet in the session, the first after a source switch and the first after a pause SHALL be populated with the available T140blocks for the source in turn to be sent as primary, and empty T140blocks for the agreed number of redundancy generations.

3.14. Creation of the redundancy

The primary T140block from a source in the latest transmitted packet is saved for populating the first redundant T140block for that source in next transmission of text from that source. The first redundant T140block for that source from the latest transmission is saved for populating the second redundant T140block in next transmission of text from that source.

Usually this is the level of redundancy used. If a higher number of redundancy is negotiated, then the procedure SHALL be maintained until all available redundant levels of T140blocks are placed in the packet. If a receiver has negotiated a lower number of "text/red" generations, then that level shall be the maximum used by the transmitter.

The T140blocks saved for transmission as redundant data are assigned a planned transmission time 330 ms after the current time.

3.15. Timer offset fields

The timestamp offset values are inserted in the redundancy header, with the time offset from the RTP timestamp in the packet when the corresponding T140block was sent as primary.

The timestamp offsets are expressed in the same clock tick units as the RTP timestamp.

The timestamp offset values for empty T140blocks have no relevance but SHOULD be assigned realistic values.

3.16. Other RTP header fields

The number of members in the CSRC list (0 or 1) shall be placed in the "CC" header field. Only mixers place value 1 in the "CC" field. A value of "0" indicates that the source is the transmitting device itself and that the source is indicated by the SSRC field. This value is used by endpoints, and by mixers sending data that it is source of itself.

The current time SHALL be inserted in the timestamp.

The SSRC of the mixer for the RTT session SHALL be inserted in the SSRC field of the RTP header.

The M-bit shall be handled as specified in [[RFC4103](#)].

3.17. Pause in transmission

When there is no new T140block to transmit, and no redundant T140block that has not been retransmitted the intended number of times from any source, the transmission process can stop until either new T140blocks arrive, or a keep-alive method calls for transmission of keep-alive packets.

3.18. RTCP considerations

A mixer SHALL send RTCP reports with SDES, CNAME and NAME information about the sources in the multi-party call. This makes it possible for participants to compose a suitable label for text from each source.

Integrity considerations SHALL be considered when composing these fields.

3.19. Reception of multi-party contents

The "text/red" receiver included in an endpoint with presentation functions will receive RTP packets in the single stream from the mixer, and SHALL distribute the T140blocks for presentation in presentation areas for each source. Other receiver roles, such as gateways or chained mixers are also feasible, and requires consideration if the stream shall just be forwarded, or distributed based on the different sources.

3.19.1. Multi-party vs two-party use

If the "CC" field value of a received packet is 1, it indicates that multi-party transmission is active, and the receiver MUST be prepared to act on the source according to its role. If the CC value is 0, the connection is point-to-point.

3.19.2. Level of redundancy

The used level of redundancy generations SHALL be evaluated from the received packet contents. The number of generations (including the primary) is equal to the number of members in the redundancy header.

3.19.3. Empty T140blocks

Empty T140blocks are included as fillers for unused primary or redundancy levels in the packets. They just do not provide any contents and do not contribute to the received streams.

3.19.4. Detection and indication of possible text loss

The RTP sequence numbers of the received packets SHALL be monitored for gaps and packets out of order. If a sequence number gap appears and still exists after some defined short time for jitter resolution, the packets in the gap SHALL be regarded lost.

If it is known that only one source is active in the RTP session, then it is likely that a gap equal to or larger than the agreed number of redundancy generations causes text loss. In that case a t140block SHALL be created with a marker for possible text loss [[T140ad1](#)] and assigned to the source and inserted in the reception buffer for that source.

If it is known that more than one source is active in the RTP session, then it is not possible in general to evaluate if text was lost when packets were lost. With two active sources and the recommended number of redundancy generations (3), it can take a gap of five lost packets until any text may be lost, but text loss can also appear if three non-consecutive packets are lost when they contained consecutive data from the same source. A simple method to decide when there is risk for resulting text loss is to evaluate if three or more packets were lost within one second. Then a t140block SHALL be created with a marker for possible text loss [[T140ad1](#)] and assigned to the SSRC of the transmitter as a general input from the mixer.

Implementations MAY apply more refined methods for more reliable detection of if text was lost or not. Any refined method SHOULD rather falsely mark possible loss when there was no loss instead of not marking possible loss when there was loss.

3.19.5. Extracting text and handling recovery

When applying the following procedures, the effects MUST be considered of possible timestamp wrap around and the RTP session possibly changing SSRC.

When a packet is received in an RTP session using the packetization for multi-party aware endpoints, its T140blocks SHALL be extracted in the following way. The description is adapted to the default redundancy case using the original and two redundant generations.

When applying the following procedures, the effects MUST be considered of possible timestamp wrap around and the RTP session possibly changing SSRC.

The source SHALL be extracted from the CSRC-list if available, otherwise from the SSRC.

If the received packet is the first packet received from the source, then all T140blocks in the packet SHALL be retrieved and assigned to a receive buffer for the source beginning with the second generation redundancy, continuing with the first generation redundancy and finally the primary.

Note that the normal case is that in the first packet, only the primary data has contents. The redundant data has contents in the first received packet from a source only after initial packet loss.

If the packet is not the first packet from a source, then if the second generation redundant data is available, its timestamp SHALL be created by subtracting its timestamp offset from the RTP timestamp. If the resulting timestamp is later than the latest retrieved data from the same source, then the redundant data SHALL be retrieved and appended to the receive buffer. The process SHALL be continued in the same way for the first generation redundant data. After that, the primary data SHALL be retrieved from the packet and appended to the receive buffer for the source.

3.19.6. Delete 'BOM'

Unicode character 'BOM' is used as a start indication and sometimes used as a filler or keep alive by transmission implementations. These SHALL be deleted after extraction from received packets.

3.20. Performance considerations

This solution has good performance for up to five participants simultaneously sending text. At higher numbers of participants simultaneously sending text, a jerkiness is visible in the presentation of text. With ten participants simultaneously transmitting text, the jerkiness is about one second. Even so, the transmission of text catches up, so there is no resulting total delay introduced. The solution is therefore suitable for emergency service use, relay service use, and small or well-managed larger multimedia conferences. Only in large unmanaged conferences with a high number of participants there may on very rare occasions appear situations when many participants happen to send text simultaneously, resulting in unpleasantly long switching times. It should be noted that it is only the number of users sending text within the same moment that causes jerkiness, not the total number of users with RTT capability.

3.21. Security for session control and media

Security SHOULD be applied on both session control and media. In applications where legacy endpoints without security may exist, a negotiation SHOULD be performed to decide if security by encryption will be applied. If no other security solution is mandated for the

application, then RFC 8643 OSRTP[[RFC8643](#)] SHOULD be applied to negotiate SRTP media security with DTLS. Most SDP examples below are for simplicity expressed without the security additions. The principles (but not all details) for applying DTLS-SRTP security is shown in a couple of the following examples.

3.22. SDP offer/answer examples

This sections shows some examples of SDP for session negotiation of the real-time text media in SIP sessions. Audio is usually provided in the same session, and sometimes also video. The examples only show the part of importance for the real-time text media.

Offer example for "text/red" format and multi-party support:

```
m=text 11000 RTP/AVP 100 98
a=rtpmap:98 t140/1000
a=rtpmap:100 red/1000
a=fmtp:100 98/98/98
a=rtt-mixer
```

Answer example from a multi-party capable device

```
m=text 14000 RTP/AVP 100 98
a=rtpmap:98 t140/1000
a=rtpmap:100 red/1000
a=fmtp:100 98/98/98
a=rtt-mixer
```

Offer example for "text/red" format including multi-party and security:

```
a=fingerprint: (fingerprint1)
m=text 11000 RTP/AVP 100 98
a=rtpmap:98 t140/1000
a=rtpmap:100 red/1000
a=fmtp:100 98/98/98
a=rtt-mixer
```

The "fingerprint" is sufficient to offer DTLS-SRTP, with the media line still indicating RTP/AVP.

Note: For brevity, the entire value of the SDP fingerprint attribute is not shown in this and the following example.

```

Answer example from a multi-party capable device with security
a=fingerprint: (fingerprint2)
m=text 16000 RTP/AVP 100 98
a=rtpmap:98 t140/1000
a=rtpmap:100 red/1000
a=fmtp:100 98/98/98
a=rtt-mixer

```

With the "fingerprint" the device acknowledges use of SRTP/DTLS.

Answer example from a multi-party unaware device that also does not support security:

```

m=text 12000 RTP/AVP 100 98
a=rtpmap:98 t140/1000
a=rtpmap:100 red/1000
a=fmtp:100 98/98/98

```

3.23. Packet sequence example from a source switch

This example shows a symbolic flow of packets from a mixer including loss and recovery. The sequence includes a source switch. A and B are sources of RTT. P indicates primary data. R1 is first redundant generation data and R2 is second redundant generation data. A1, B1, A2 etc are text chunks (T140blocks) received from the respective sources. X indicates dropped packet between the mixer and a receiver. The session is assumed to use original and two redundant generations of RTT.

```

|-----|
|Seq no 101|
|CC=1      |
|CSRC list A|
|R2: A1    |
|R1: A2    |
|P:  A3    |
|-----|

```

Assuming that earlier packets (with text A1 and A2) were received in sequence, text A3 is received from packet 101 and assigned to reception area A. The mixer is now assumed to have received text from source B and will send that text 100 ms after packet 101. Transmission of A2 and A3 as redundancy is planned for 200 ms after packet 101.

```
|-----|
|Seq no 102|
|CC=1      |
|CSRC list B|
|R2 Empty  |
|R1: Empty  |
|P: B1     |
|-----|
```

B1 is retrieved from this packet. Redundant transmission of B1 is planned 200 ms after packet 102.

```
X-----|
X Seq no 103|
X CC=1      |
X CSRC list A|
X R2: A2    |
X R1: A3    |
X P: Empty  |
X-----|
```

Packet 103 is assumed to be dropped in network problems. It contains redundancy for A. Sending A3 as second level redundancy is planned for 100 ms after packet 104.

```
X-----|
X Seq no 104|
X CC=1      |
X CSRC list B|
X R2: Empty  |
X R1: B1     |
X P2: B2     |
X-----|
```

Packet 104 contains text from B, assumed dropped in network problems. The mixer has A3 redundancy to send and plans it after 100 ms.

```
X-----|
X Seq no 105|
X CC=1      |
X CSRC list A|
X R2: A3     |
X R1: Empty  |
X P: Empty  |
X-----|
```

Packet 105 is assumed to be dropped in network problems B1 and B2 still needs to be transmitted as redundancy. This is planned 320 ms after packet 105.

```
|-----|
|Seq no 106|
|CC=1      |
|CSRC list B|
| R2: B1   |
| R1: B2   |
| P:  Empty|
|-----|
```

Packet 106 is received 320 ms after packet 105. The latest received sequence number was 102. 103,104,105 were lost. Three packets were thus lost during less than one second. The simple rule for detection of text loss then results in that an indicator for possible loss (U'FFFD) [[T140ad1](#)], should be inserted generally for the mixer.

The second level redundancy in packet 106 is B1 and has timestamp offset 620 ms. The timestamp of packet 106 minus 620 is the timestamp of packet 101 which was received. So B1 does not need to be retrieved. The first level redundancy in packet 106 is 420. The timestamp of packet 106 minus 420 is later than the latest received packet with source A. Therefore B2 is retrieved and assigned to the input buffer for source B. No primary is available in packet 106

After this sequence, A3 and B1 and B2 have been received. In this case no text was lost.

3.24. Maximum character rate "CPS"

The default maximum rate of reception of "text/t140" real-time text is in RFC 4103 [[RFC4103](#)] specified to be 30 characters per second. The value MAY be modified in the CPS parameter of the FMTP attribute in the media section for the "text/t140" media. A mixer combining real-time text from a number of sources may occasionally have a higher combined flow of text coming from the sources. Endpoints SHOULD therefore specify a suitable higher value for the CPS parameter, corresponding to its real reception capability. A value for "CPS" of 90 is the default for the "text/t140" stream in the "text/red" format when multi-party real-time text is negotiated. See RFC 4103 [[RFC4103](#)] for the format and use of the CPS parameter. The same rules apply for the multi-party case except for the default value.

4. Presentation level considerations

ITU-T T.140 [[T140](#)] provides the presentation level requirements for the RFC 4103 [[RFC4103](#)] transport. T.140 [[T140](#)] has functions for erasure and other formatting functions and has the following general statement for the presentation:

"The display of text from the members of the conversation should be arranged so that the text from each participant is clearly readable, and its source and the relative timing of entered text is visualized in the display. Mechanisms for looking back in the contents from the current session should be provided. The text should be displayed as soon as it is received."

Strict application of T.140 [T140] is of essence for the interoperability of real-time text implementations and to fulfill the intention that the session participants have the same information of the text contents of the conversation without necessarily having the exact same layout of the conversation.

T.140 [T140] specifies a set of presentation control codes to include in the stream. Some of them are optional. Implementations MUST be able to ignore optional control codes that they do not support.

There is no strict "message" concept in real-time text. Line Separator SHALL be used as a separator allowing a part of received text to be grouped in presentation. The characters "CRLF" may be used by other implementations as replacement for Line Separator. The "CRLF" combination SHALL be erased by just one erasing action, just as the Line Separator. Presentation functions are allowed to group text for presentation in smaller groups than the line separators imply and present such groups with source indication together with text groups from other sources (see the following presentation examples). Erasure has no specific limit by any delimiter in the text stream.

4.1. Presentation by multi-party aware endpoints

A multi-party aware receiving party, presenting real-time text MUST separate text from different sources and present them in separate presentation fields. The receiving party MAY separate presentation of parts of text from a source in readable groups based on other criteria than line separator and merge these groups in the presentation area when it benefits the user to most easily find and read text from the different participants. The criteria MAY e.g. be a received comma, full stop, or other phrase delimiters, or a long pause.

When text is received from multiple original sources, the presentation SHOULD provide a view where text is added in multiple presentation fields.

If the presentation presents text from different sources in one common area, the presenting endpoint SHOULD insert text from the local user ended at suitable points merged with received text to

indicate the relative timing for when the text groups were completed. In this presentation mode, the receiving endpoint SHALL present the source of the different groups of text.

A view of a three-party RTT call in chat style is shown in this example .

	^
[Alice] Hi, Alice here.	-
[Bob] Bob as well.	
[Eve] Hi, this is Eve, calling from Paris.	
I thought you should be here.	
[Alice] I am coming on Thursday, my	
performance is not until Friday morning.	
[Bob] And I on Wednesday evening.	
[Alice] Can we meet on Thursday evening?	
[Eve] Yes, definitely. How about 7pm.	
at the entrance of the restaurant	
Le Lion Blanc?	
[Eve] we can have dinner and then take a walk	-
	v
<Eve-typing> But I need to be back to	^
the hotel by 11 because I need	-
<Bob-typing> I wou	-
	v
of course, I underst	

Figure 3: Example of a three-party RTT call presented in chat style seen at participant 'Alice's endpoint.

Other presentation styles than the chat style may be arranged.

This figure shows how a coordinated column view MAY be presented.

Bob	Eve	Alice
		I will arrive by TGV.
My flight is to Orly		Convenient to the main
	Hi all, can we plan	station.
	for the seminar?	
Eve, will you do		
your presentation on		
Friday?	Yes, Friday at 10.	
Fine, wo		We need to meet befo

Figure 4: An example of a coordinated column-view of a three-party session with entries ordered vertically in approximate time-order.

4.2. Multi-party mixing for multi-party unaware endpoints

When the mixer has indicated multi-party capability by the "rtt-mixer" sdp attribute in an SDP negotiation, but the multi-party capability negotiation fails with an endpoint, then the agreed "text/red" or "text/t140" format SHALL be used and the mixer SHOULD compose a best-effort presentation of multi-party real-time text in one stream intended to be presented by an endpoint with no multi-party awareness.

This presentation format has functional limitations and SHOULD be used only to enable participation in multi-party calls by legacy deployed endpoints implementing only RFC 4103 without any multi-party extensions specified in this document.

The principles and procedures below do not specify any new protocol elements. They are instead composed from the information in ITU-T T.140 [T140] and an ambition to provide a best effort presentation on an endpoint which has functions only for two-party calls.

The mixer mixing for multi-party unaware endpoints SHALL compose a simulated limited multi-party RTT view suitable for presentation in one presentation area. The mixer SHALL group text in suitable groups and prepare for presentation of them by inserting a new line between them if the transmitted text did not already end with a new line. A presentable label SHOULD be composed and sent for the source initially in the session and after each source switch. With this procedure the time for source switching is depending on the actions of the users. In order to expedite source switch, a user can for example end its turn with a new line.

4.2.1. Actions by the mixer at reception from the call participants

When text is received by the mixer from the different participants, the mixer SHALL recover text from redundancy if any packets are lost. The mark for lost text [[T140ad1](#)] SHOULD be inserted in the stream if unrecoverable loss appears. Any Unicode "BOM" characters, possibly used for keep-alive shall be deleted. The time of creation of text (retrieved from the RTP timestamp) SHALL be stored together with the received text from each source in queues for transmission to the recipients.

4.2.2. Actions by the mixer for transmission to the recipients

The following procedure SHOULD be applied for each recipient of multi-part text from the mixer.

The text for transmission SHOULD be formatted by the mixer for each receiving user for presentation in one single presentation area. Text received from a participant SHOULD NOT be included in transmission to that participant. When there is text available for transmission from the mixer to a receiving party from more than one participant, the mixer SHOULD switch between transmission of text from the different sources at suitable points in the transmitted stream.

When switching source, the mixer SHOULD insert a line separator if the already transmitted text did not end with a new line (line separator or CRLF). A label SHOULD be composed from information in the CNAME and NAME fields in RTCP reports from the participant to have its text transmitted, or from other session information for that user. The label SHOULD be delimited by suitable characters (e.g. '[' ']') and transmitted. The CSRC SHOULD indicate the selected source. Then text from that selected participant SHOULD be transmitted until a new suitable point for switching source is reached.

Integrity considerations SHALL be taken when composing the label.

Seeking a suitable point for switching source SHOULD be done when there is older text waiting for transmission from any party than the age of the last transmitted text. Suitable points for switching are:

- *A completed phrase ended by comma

- *A completed sentence

- *A new line (line separator or CRLF)

- *A long pause (e.g. > 10 seconds) in received text from the currently transmitted source

*If text from one participant has been transmitted with text from other sources waiting for transmission for a long time (e.g. > 1 minute) and none of the other suitable points for switching has occurred, a source switch MAY be forced by the mixer at next word delimiter, and also if even a word delimiter does not occur within a time (e.g. 15 seconds) after the scan for word delimiter started.

When switching source, the source which has the oldest text in queue SHOULD be selected to be transmitted. A character display count SHOULD be maintained for the currently transmitted source, starting at zero after the label is transmitted for the currently transmitted source.

The status SHOULD be maintained for the latest control code for Select Graphic Rendition (SGR) from each source. If there is an SGR code stored as the status for the current source before the source switch is done, a reset of SGR shall be sent by the sequence SGR 0 [009B 0000 006D] after the new line and before the new label during a source switch. See SGR below for an explanation. This transmission does not influence the display count.

If there is an SGR code stored for the new source after the source switch, that SGR code SHOULD be transmitted to the recipient before the label. This transmission does not influence the display count.

4.2.3. Actions on transmission of text

Text from a source sent to the recipient SHOULD increase the display count by one per transmitted character.

4.2.4. Actions on transmission of control codes

The following control codes specified by T.140 require specific actions. They SHOULD cause specific considerations in the mixer. Note that the codes presented here are expressed in UCS-16, while transmission is made in UTF-8 transform of these codes.

BEL 0007 Bell Alert in session, provides for alerting during an active session. The display count SHOULD not be altered.

NEW LINE 2028 Line separator. Check and perform a source switch if appropriate. Increase display count by 1.

CR LF 000D 000A A supported, but not preferred way of requesting a new line. Check and perform a source switch if appropriate. Increase display count by 1.

INT ESC 0061 Interrupt (used to initiate mode negotiation procedure). The display count SHOULD not be altered.

SGR 009B Ps 006D

Select graphic rendition. Ps is rendition parameters specified in ISO 6429. The display count SHOULD not be altered. The SGR code SHOULD be stored for the current source.

SOS 0098 Start of string, used as a general protocol element introducer, followed by a maximum 256 bytes string and the ST. The display count SHOULD not be altered.

ST 009C String terminator, end of SOS string. The display count SHOULD not be altered.

ESC 001B Escape - used in control strings. The display count SHOULD not be altered for the complete escape code.

Byte order mark "BOM" (U+FEFF) "Zero width, no break space", used for synchronization and keep-alive. SHOULD be deleted from incoming streams. Shall be sent first after session establishment to the recipient. The display count shall not be altered.

Missing text mark (U+FFFD) "Replacement character", represented as a question mark in a rhombus, or if that is not feasible, replaced by an apostrophe ', marks place in stream of possible text loss. SHOULD be inserted by the reception procedure in case of unrecoverable loss of packets. The display count SHOULD be increased by one when sent as for any other character.

SGR If a control code for selecting graphic rendition (SGR), other than reset of the graphic rendition (SGR 0) is sent to a recipient, that control code shall also be stored as status for the source in the storage for SGR status. If a reset graphic rendition (SGR 0) originated from a source is sent, then the SGR status storage for that source shall be cleared. The display count shall not be increased.

BS (U+0008) Back Space, intended to erase the last entered character by a source. Erasure by backspace cannot always be performed as the erasing party intended. If an erasing action erases all text up to the end of the leading label after a source switch, then the mixer must not transmit more backspaces. Instead it is RECOMMENDED that a letter "X" is inserted in the text stream for each backspace as an indication of the intent to erase more. A new line is usually coded by a Line Separator, but the character combination "CRLF" MAY be used instead. Erasure of a new line is in both cases done by just one erasing action (Backspace). If the display count has a positive value it is decreased by one when the BS is sent. If the display count is at zero, it is not altered.

4.2.5. Packet transmission

A mixer transmitting to a multi-party unaware terminal SHOULD send primary data only from one source per packet. The SSRC SHOULD be the SSRC of the mixer. The CSRC list SHOULD contain one member and be the SSRC of the source of the primary data.

4.2.6. Functional limitations

When a multi-party unaware endpoint presents a conversation in one display area in a chat style, it inserts source indications for remote text and local user text as they are merged in completed text groups. When an endpoint using this layout receives and presents text mixed for multi-party unaware endpoints, there will be two levels of source indicators for the received text; one generated by the mixer and inserted in a label after each source switch, and another generated by the receiving endpoint and inserted after each switch between local and remote source in the presentation area. This will waste display space and look inconsistent to the reader.

New text can be presented only from one source at a time. Switch of source to be presented takes place at suitable places in the text, such as end of phrase, end of sentence, line separator and inactivity. Therefore the time to switch to present waiting text from other sources may become long and will vary and depend on the actions of the currently presented source.

Erase can only be done up to the latest source switch. If a user tries to erase more text, the erasing actions will be presented as letter X after the label.

Text loss because of network errors may hit the label between entries from different parties, causing risk for misunderstanding from which source a piece of text is.

These facts makes it strongly RECOMMENDED to implement multi-party awareness in RTT endpoints. The use of the mixing method for multi-party-unaware endpoints should be left for use with endpoints which are impossible to upgrade to become multi-party aware.

4.2.7. Example views of presentation on multi-party unaware endpoints

The following pictures are examples of the view on a participant's display for the multi-party-unaware case.

Conference	Alice
	I will arrive by TGV.
[Bob]:My flight is to Orly.	Convenient to the main station.
[Eve]:Hi all, can we plan for the seminar.	
[Bob]:Eve, will you do your presentation on Friday?	
[Eve]:Yes, Friday at 10.	
[Bob]: Fine, wo	We need to meet befo

Figure 5: Alice who has a conference-unaware client is receiving the multi-party real-time text in a single-stream. This figure shows how a coordinated column view MAY be presented on Alice's device.

	^
[Alice] Hi, Alice here.	-
[mix][Bob] Bob as well.	
[Eve] Hi, this is Eve, calling from Paris	
I thought you should be here.	
[Alice] I am coming on Thursday, my	
performance is not until Friday morning.	
[mix][Bob] And I on Wednesday evening.	
[Eve] we can have dinner and then walk	
[Eve] But I need to be back to	
the hotel by 11 because I need	
	-
	v
of course, I underst	

Figure 6: An example of a view of the multi-party unaware presentation in chat style. Alice is the local user.

5. Relation to Conference Control

5.1. Use with SIP centralized conferencing framework

The SIP conferencing framework, mainly specified in RFC 4353[[RFC4353](#)], RFC 4579[[RFC4579](#)] and RFC 4575[[RFC4575](#)] is suitable for coordinating sessions including multi-party RTT. The RTT stream between the mixer and a participant is one and the same during the conference. Participants get announced by notifications when participants are joining or leaving, and further user information may be provided. The SSRC of the text to expect from joined users MAY be included in a notification. The notifications MAY be used both for security purposes and for translation to a label for presentation to other users.

5.2. Conference control

In managed conferences, control of the real-time text media SHOULD be provided in the same way as other for media, e.g. for muting and unmuting by the direction attributes in SDP [[RFC4566](#)].

Note that floor control functions may be of value for RTT users as well as for users of other media in a conference.

6. Gateway Considerations

6.1. Gateway considerations with Textphones (e.g. TTYs).

Multi-party RTT sessions may involve gateways of different kinds. Gateways involved in setting up sessions SHALL correctly reflect the multi-party capability or unawareness of the combination of the gateway and the remote endpoint beyond the gateway.

One case that may occur is a gateway to PSTN for communication with textphones (e.g. TTYs). Textphones are limited devices with no multi-party awareness, and it SHOULD therefore be suitable for the gateway to not indicate multi-party awareness for that case. Another solution is that the gateway indicates multi-party capability towards the mixer, and includes the multi-party mixer function for multi-party unaware endpoints itself. This solution makes it possible to make adaptations for the functional limitations of the textphone (TTY).

More information on gateways to textphones (TTYs) is found in RFC 5194[[RFC5194](#)]

6.2. Gateway considerations with WebRTC.

Gateway operation to real-time text in WebRTC may also be required. In WebRTC, RTT is specified in [[I-D.ietf-mmusic-t140-usage-data-channel](#)].

A multi-party bridge may have functionality for communicating by RTT both in RTP streams with RTT and WebRTC t140 data channels. Other configurations may consist of a multi-party bridge with either technology for RTT transport and a separate gateway for conversion of the text communication streams between RTP and t140 data channel.

In WebRTC, it is assumed that for a multi-party session, one t140 data channel is established for each source from a gateway or bridge to each participant. Each participant also has a data channel with two-way connection with the gateway or bridge.

The t140 channel used both ways is for text from the WebRTC user and from the bridge or gateway itself to the WebRTC user. The label parameter of this t140 channel is used as NAME field in RTCP to participants on the RTP side. The other t140 channels are only for text from other participants to the WebRTC user.

When a new participant has entered the session with RTP transport of rtt, a new t140 channel SHOULD be established to WebRTC users with the label parameter composed from the NAME field in RTCP on the RTP side.

When a new participant has entered the multi-party session with RTT transport in a WebRTC t140 data channel, the new participant SHOULD be announced by a notification to RTP users. The label parameter from the WebRTC side SHOULD be used as the NAME RTCP field on the RTP side, or other available session information.

7. Updates to RFC 4103

This document updates RFC 4103 [[RFC4103](#)] by introducing an sdp media attribute "rtt-mixer" for negotiation of multi-party mixing capability with the [[RFC4103](#)] format, and by specifying the rules for packets when multi-party capability is negotiated and in use.

8. Congestion considerations

The congestion considerations and recommended actions from RFC 4103 [[RFC4103](#)] are valid also in multi-party situations.

The first action in case of congestion SHOULD be to temporarily increase the transmission interval up to two seconds.

If the unlikely situation appears that more than 20 participants in a conference send text simultaneously, it will take more than 7 seconds between presentation of text from each of these participants. More time than that can cause confusion in the session. It is therefore RECOMMENDED that the mixer discards such text in excess inserts a general indication of possible text loss [[T140ad1](#)] in the session. If the main text contributor is indicated in any way, the mixer MAY avoid deleting text from that participant.

9. Acknowledgements

James Hamlin for format and performance aspects.

10. IANA Considerations

10.1. Registration of the "rtt-mixer" sdp media attribute

[RFC EDITOR NOTE: Please replace all instances of RFCXXXX with the RFC number of this document.]

IANA is asked to register the new sdp attribute "rtt-mixer".

Contact name: IESG

Contact email: iesg@ietf.org

Attribute name: rtt-mixer

Attribute semantics: See RFCXXXX [Section 3.1](#)

Attribute value: none

Usage level: media

Purpose: Indicate support by mixer and endpoint of multi-party mixing for real-time text transmission, using a common RTP-stream for transmission of text from a number of sources mixed with one source at a time and the source indicated in a single CSRC-list member.

Charset Dependent: no

O/A procedure: See RFCXXXX [Section 3.1](#)

Mux Category: normal

Reference: RFCXXXX

11. Security Considerations

The RTP-mixer model requires the mixer to be allowed to decrypt, pack and encrypt secured text from the conference participants. Therefore the mixer needs to be trusted. This is similar to the situation for central mixers of audio and video.

The requirement to transfer information about the user in RTCP reports in SDES, CNAME and NAME fields, and in conference notifications, for creation of labels may have privacy concerns as already stated in RFC 3550 [[RFC3550](#)], and may be restricted of privacy reasons. The receiving user will then get a more symbolic label for the source.

Participants with malicious intentions may appear and e.g. disturb the multi-party session by a continuous flow of text, or masquerading as text from other participants. Counteractions should be to require secure signaling, media and authentication, and to provide higher level conference functions e.g. for blocking and expelling participants.

12. Change history

12.1. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-10

The packet composition was modified for interleaving packets from different sources.

The packet reception was modified for the new interleaving method.

The packet sequence examples was adjusted for the new interleaving method.

Modifications according to responses to Brian Rosen of 2020-11-03

12.2. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-09

Changed name on the SDP media attribute to "rtt-mixer"

Restructure of section 2 for balance between aware and unaware cases.

Moved conference control to own section.

Improved clarification of recovery and loss in the packet sequence example.

A number of editorial corrections and improvements.

12.3. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-08

Deleted the method requiring a new packet format "text/rex" because of the longer standardization and implementation period it needs.

Focus on use of RFC 4103 text/red format with shorter transmission interval, and source indicated in CSRC.

12.4. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-07

Added a method based on the "text/red" format and single source per packet, negotiated by the "rtt-mixer" sdp attribute.

Added reasoning and recommendation about indication of loss.

The highest number of sources in one packet is 15, not 16. Changed.

Added in information on update to RFC 4103 that RFC 4103 explicitly allows addition of FEC method. The redundancy is a kind of forward error correction..

12.5. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-06

Improved definitions list format.

The format of the media subtype parameters is made to match the requirements.

The mapping of media subtype parameters to sdp is included.

The CPS parameter belongs to the t140 subtype and does not need to be registered here.

12.6. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-05

nomenclature and editorial improvements

"this document" used consistently to refer to this document.

12.7. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-04

'Redundancy header' renamed to 'data header'.

More clarifications added.

Language and figure number corrections.

12.8. Changes included in draft-ietf-avtc core-multi-party-rtt-mix-03

Mention possible need to mute and raise hands as for other media.
---done ----

Make sure that use in two-party calls is also possible and explained. - may need more wording -

Clarify the RTT is often used together with other media. --done--

Tell that text mixing is N-1. A users own text is not received in the mix. -done-

In 3. correct the interval to: A "text/rex" transmitter SHOULD send packets distributed in time as long as there is something (new or redundant T140blocks) to transmit. The maximum transmission interval SHOULD then be 300 ms. It is RECOMMENDED to send a packet to a receiver as soon as new text to that receiver is available, as long as the time after the latest sent packet to the same receiver is more than 150 ms, and also the maximum character rate to the receiver is not exceeded. The intention is to keep the latency low while keeping a good protection against text loss in bursty packet loss conditions. -done-

In 1.3 say that the format is used both ways. -done-

In 13.1 change presentation area to presentation field so that reader does not think it shall be totally separated. -done-

In Performance and intro, tell the performance in number of simultaneous sending users and introduced delay 16, 150 vs requirements 5 vs 500. -done --

Clarify redundancy level per connection. -done-

Timestamp also for the last data header. To make it possible for all text to have time offset as for transmission from the source. Make that header equal to the others. -done-

Mixer always use the CSRC list, even for its own BOM. -done-

Combine all talk about transmission interval (300 ms vs when text has arrived) in section 3 in one paragraph or close to each other. -done-

Documents the goal of good performance with low delay for 5 simultaneous typers in the introduction. -done-

Describe better that only primary text shall be sent on to receivers. Redundancy and loss must be resolved by the mixer. -done-

12.9. Changes included in draft-ietf-avtcore-multi-party-rtt-mix-02

SDP and better description and visibility of security by OSRTP RFC 8634 needed.

The description of gatewaying to WebRTC extended.

The description of the data header in the packet is improved.

12.10. Changes to draft-ietf-avtcore-multi-party-rtt-mix-01

2,5,6 More efficient format "text/rex" introduced and attribute a=rtt-mix deleted.

3. Brief about use of OSRTP for security included- More needed.

4. Brief motivation for the solution and why not rtp-translator is used added to intro.

7. More limitations for the multi-party unaware mixing method inserted.

8. Updates to RFC 4102 and 4103 more clearly expressed.

9. Gateway to WebRTC started. More needed.

12.11. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-03 to draft-ietf-avtcore-multi-party-rtt-mix-00

Changed file name to draft-ietf-avtcore-multi-party-rtt-mix-00

Replaced CDATA in IANA registration table with better coding.

Converted to xml2rfc version 3.

12.12. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-02 to -03

Changed company and e-mail of the author.

Changed title to "RTP-mixer formatting of multi-party Real-time text" to better match contents.

Check and modification where needed of use of RFC 2119 words SHALL etc.

More about the CC value in sections on transmitters and receivers so that 1-to-1 sessions do not use the mixer format.

Enhanced section on presentation for multi-party-unaware endpoints

A paragraph recommending CPS=150 inserted in the performance section.

12.13. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-01 to -02

In Abstract and 1. Introduction: Introduced wording about regulatory requirements.

In section 5: The transmission interval is decreased to 100 ms when there is text from more than one source to transmit.

In section 11 about SDP negotiation, a SHOULD-requirement is introduced that the mixer should make a mix for multi-party unaware endpoints if the negotiation is not successful. And a reference to a later chapter about it.

The presentation considerations chapter 14 is extended with more information about presentation on multi-party aware endpoints, and a new section on the multi-party unaware mixing with low functionality but SHOULD a be implemented in mixers. Presentation examples are added.

A short chapter 15 on gateway considerations is introduced.

Clarification about the text/t140 format included in chapter 10.

This sentence added to the chapter 10 about use without redundancy. "The text/red format SHOULD be used unless some other protection against packet loss is utilized, for example a reliable network or transport."

Note about deviation from RFC 2198 added in chapter 4.

In chapter 9. "Use with SIP centralized conferencing framework" the following note is inserted: Note: The CSRC-list in an RTP packet only includes participants who's text is included in one or more text blocks. It is not the same as the list of participants in a conference. With audio and video media, the CSRC-list would often contain all participants who are not muted whereas text participants that don't type are completely silent and so don't show up in RTP packet CSRC-lists.

12.14. Changes from draft-hellstrom-avtcore-multi-party-rtt-source-00 to -01

Editorial cleanup.

Changed capability indication from fmtp-parameter to SDP attribute "rtt-mix".

Swapped order of redundancy elements in the example to match reality.

Increased the SDP negotiation section

13. References

13.1. Normative References

[I-D.ietf-mmusic-t140-usage-data-channel]

Holmberg, C. and G. Hellstrom, "T.140 Real-time Text Conversation over WebRTC Data Channels", Work in Progress, Internet-Draft, draft-ietf-mmusic-t140-usage-data-channel-14, 10 April 2020, <<https://tools.ietf.org/html/draft-ietf-mmusic-t140-usage-data-channel-14>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[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, <<https://www.rfc-editor.org/info/rfc3550>>.

[RFC4102] Jones, P., "Registration of the text/red MIME Sub-Type", RFC 4102, DOI 10.17487/RFC4102, June 2005, <<https://www.rfc-editor.org/info/rfc4102>>.

[RFC4103] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", RFC 4103, DOI 10.17487/RFC4103, June 2005, <<https://www.rfc-editor.org/info/rfc4103>>.

[RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, DOI 10.17487/RFC4566, July 2006, <<https://www.rfc-editor.org/info/rfc4566>>.

[RFC5764] McGrew, D. and E. Rescorla, "Datagram Transport Layer Security (DTLS) Extension to Establish Keys for the Secure Real-time Transport Protocol (SRTP)", RFC 5764, DOI 10.17487/RFC5764, May 2010, <<https://www.rfc-editor.org/info/rfc5764>>.

[RFC6263] Marjou, X. and A. Sollaud, "Application Mechanism for Keeping Alive the NAT Mappings Associated with RTP / RTP Control Protocol (RTCP) Flows", RFC 6263, DOI 10.17487/RFC6263, June 2011, <<https://www.rfc-editor.org/info/rfc6263>>.

[RFC8643] Johnston, A., Aboba, B., Hutton, A., Jesske, R., and T. Stach, "An Opportunistic Approach for Secure Real-time Transport Protocol (OSRTP)", RFC 8643, DOI 10.17487/

RFC8643, August 2019, <<https://www.rfc-editor.org/info/rfc8643>>.

[T140] ITU-T, "Recommendation ITU-T T.140 (02/1998), Protocol for multimedia application text conversation", February 1998, <<https://www.itu.int/rec/T-REC-T.140-199802-I/en>>.

[T140ad1] ITU-T, "Recommendation ITU-T.140 Addendum 1 - (02/2000), Protocol for multimedia application text conversation", February 2000, <<https://www.itu.int/rec/T-REC-T.140-200002-I!Add1/en>>.

13.2. Informative References

[RFC4353] Rosenberg, J., "A Framework for Conferencing with the Session Initiation Protocol (SIP)", RFC 4353, DOI 10.17487/RFC4353, February 2006, <<https://www.rfc-editor.org/info/rfc4353>>.

[RFC4575] Rosenberg, J., Schulzrinne, H., and O. Levin, Ed., "A Session Initiation Protocol (SIP) Event Package for Conference State", RFC 4575, DOI 10.17487/RFC4575, August 2006, <<https://www.rfc-editor.org/info/rfc4575>>.

[RFC4579] Johnston, A. and O. Levin, "Session Initiation Protocol (SIP) Call Control - Conferencing for User Agents", BCP 119, RFC 4579, DOI 10.17487/RFC4579, August 2006, <<https://www.rfc-editor.org/info/rfc4579>>.

[RFC5194] van Wijk, A., Ed. and G. Gybels, Ed., "Framework for Real-Time Text over IP Using the Session Initiation Protocol (SIP)", RFC 5194, DOI 10.17487/RFC5194, June 2008, <<https://www.rfc-editor.org/info/rfc5194>>.

Author's Address

Gunnar Hellstrom
Gunnar Hellstrom Accessible Communication
SE-13670 Vendelso
Sweden

Email: gunnar.hellstrom@ghaccess.se