

Network Working Group
Request for Comments: 4351
Category: Historic

G. Hellstrom
Omnitor AB
P. Jones
Cisco Systems, Inc.
January 2006

Real-Time Transport Protocol (RTP) Payload for
Text Conversation Interleaved in an Audio Stream

Status of This Memo

This memo defines a Historic Document for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This memo describes how to carry real-time text conversation session contents in RTP packets. Text conversation session contents are specified in ITU-T Recommendation T.140.

One payload format is described for transmitting audio and text data within a single RTP session.

This RTP payload description recommends a method to include redundant text from already transmitted packets in order to reduce the risk of text loss caused by packet loss.

Table of Contents

1. Introduction	3
2. Conventions Used in This Document	4
3. Usage of RTP	4
3.1. Motivations and Rationale	4
3.2. Payload Format for Transmission of audio/t140c Data	4
3.3. The "T140block"	5
3.4. Synchronization of Text with Other Media	5
3.5. Synchronization Considerations for the audio/t140c Format ..	5
3.6. RTP Packet Header	6
4. Protection against Loss of Data	7
4.1. Payload Format When Using Redundancy	7
4.2. Using Redundancy with the audio/t140c Format	8
5. Recommended Procedure	8
5.1. Recommended Basic Procedure	8
5.2. Transmission before and after "Idle Periods"	9
5.3. Detection of Lost Text Packets	9
5.4. Compensation for Packets Out of Order	10
6. Parameter for Character Transmission Rate	10
7. Examples	11
7.1. RTP Packetization Examples for the audio/t140c Format	11
7.2. SDP Examples	12
8. Security Considerations	13
8.1. Confidentiality	13
8.2. Integrity	13
8.3. Source Authentication	13
9. Congestion Considerations	14
10. IANA Considerations	15
10.1. Registration of MIME Media Type audio/t140c	15
10.2. SDP Mapping of MIME Parameters	16
10.3. Offer/Answer Consideration	17
11. Acknowledgements	17
12. Normative References	17
13. Informative References	18

1. Introduction

This document defines a payload type for carrying text conversation session contents in RTP [2] packets. Text conversation session contents are specified in ITU-T Recommendation T.140 [1]. Text conversation is used alone or in connection to other conversational facilities, such as video and voice, to form multimedia conversation services. Text in multimedia conversation sessions is sent character-by-character as soon as it is available, or with a small delay for buffering.

The text is intended to be entered by human users from a keyboard, handwriting recognition, voice recognition, or any other input method. The rate of character entry is usually at a level of a few characters per second or less. In general, only one or a few new characters are expected to be transmitted with each packet. Small blocks of text may be prepared by the user and pasted into the user interface for transmission during the conversation, occasionally causing packets to carry more payload.

T.140 specifies that text and other T.140 elements must be transmitted in ISO 10646-1[5] code with UTF-8 [6] transformation. That makes it easy to implement internationally useful applications and to handle the text in modern information technology environments. The payload of an RTP packet following this specification consists of text encoded according to T.140 without any additional framing. A common case will be a single ISO 10646 character, UTF-8 encoded.

T.140 requires the transport channel to provide characters without duplication and in original order. Text conversation users expect that text will be delivered with no or a low level of lost information.

Therefore a mechanism based on RTP is specified here. It gives text arrival in correct order, without duplication, and with detection and indication of loss. It also includes an optional possibility to repeat data for redundancy to lower the risk of loss. Since packet overhead is usually much larger than the T.140 contents, the increase in bandwidth with the use of redundancy is minimal.

By using RTP for text transmission in a multimedia conversation application, uniform handling of text and other media can be achieved in, as examples, conferencing systems, firewalls, and network translation devices. This, in turn, eases the design and increases the possibility for prompt and proper media delivery.

This document introduces a method of transporting text interleaved with voice within the same RTP session.

2. Conventions Used in This Document

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 RFC 2119 [4].

3. Usage of RTP

The payload format for real-time text transmission with RTP [2] described in this memo is intended for use between Public Switched Telephone Network (PSTN) gateways and is called audio/t140c.

3.1. Motivations and Rationale

The audio/t140c payload specification is intended to allow gateways that are interconnecting two PSTN networks to interleave, through a single RTP session, audio and text data received on the PSTN circuit. This is comparable to the way in which dual-tone multifrequency (DTMF) is extracted and transmitted within an RTP session [14].

The audio/t140c format SHALL NOT be used for applications other than PSTN gateway applications. In such applications, a specific profiling document MAY make it REQUIRED for a specific application. The reason to prefer to use audio/t140c could be for gateway application where the ports are a limited and scarce resource. Applications SHOULD use RFC 4103 [15] for real-time text communication that falls outside the limited scope of this specification.

3.2. Payload Format for Transmission of audio/t140c Data

An audio/t140c conversation RTP payload format consists of a 16-bit "T140block counter" carried in network byte order (see RFC 791 [11] Annex B), followed by one and only one "T140block" (see section 3.3). The fields in the RTP header are set as defined in section 3.6.

The T140block counter MUST be initialized to zero the first time that a packet containing a T140block is transmitted and MUST be incremented by 1 each time that a new block is transmitted. Once the counter reaches the value 0xFFFF, the counter is reset to 0 the next time the counter is incremented. This T140block counter is used to detect lost blocks and to avoid duplication of blocks.

For the purposes of readability, the remainder of this document refers only to the T140block without making explicit reference to the T140block counter. Readers should understand that when using the audio/t140c format, the T140block counter MUST always precede the actual T140block, including redundant data transmissions.

3.3. The "T140block"

T.140 text is UTF-8 coded as specified in T.140 with no extra framing. The T140block contains one or more T.140 code elements as specified in [1]. Most T.140 code elements are single ISO 10646 [5] characters, but some are multiple-character sequences. Each character is UTF-8 encoded [6] into one or more octets. Each block MUST contain an integral number of UTF-8-encoded characters regardless of the number of octets per character. Any composite character sequence (CCS) SHOULD be placed within one block.

3.4. Synchronization of Text with Other Media

Usually, each medium in a session utilizes a separate RTP stream. As such, if synchronization of the text and other media packets is important, the streams MUST be associated when the sessions are established and the streams MUST share the same reference clock (refer to the description of the timestamp field as it relates to synchronization in section 5.1 of RFC 3550). Association of RTP streams can be done through the CNAME field of RTP Control Protocol (RTCP) SDES function. It is dependent on the particular application and is outside the scope of this document.

3.5. Synchronization Considerations for the audio/t140c Format

The audio/t140c packets are generally transmitted as interleaved packets between voice packets or other kinds of audio packets with the intention to create one common audio signal in the receiving equipment to be used for alternating between text and voice. The audio/t140c payload is then used to play out audio signals according to a PSTN textphone coding method (usually a modem).

One should observe the RTP timestamps of the voice, text, or other audio packets in order to reproduce the stream correctly when playing out the audio. Also, note that incoming text from a PSTN circuit might be at a higher bit-rate than can be played out on an egress PSTN circuit. As such, it is possible that, on the egress side, a gateway may not complete the play out of the text packets before it is time to play the next voice packet. Given that this application is primarily for the benefit of users of PSTN textphone devices, it is strongly RECOMMENDED that all received text packets be properly reproduced on the egress gateway before considering any other subsequent audio packets.

If necessary, voice and other audio packets should be discarded in order to properly reproduce the text signals on the PSTN circuit, even if the text packets arrive late.

The PSTN textphone users commonly use turn-taking indicators in the text stream, so it can be expected that as long as text is transmitted, it is valid text and should be given priority over voice.

Note that the usual RTP semantics apply with regards to switching payload formats within an RTP session. A sender MAY switch between "audio/t140c" and some other format within an RTP session, but MUST NOT send overlapping data using two different audio formats within an RTP session. This does not prohibit an implementation from being split into two logical parts to send overlapping data, each part using a different synchronization source (SSRC) and sending its own RTP and RTCP (such an endpoint will appear to others in the session as two participants with different SSRCs, but the same RTCP SDES CNAME). Further details around using multiple payloads in an RTP session can be found in RFC 3550 [2].

3.6. RTP Packet Header

Each RTP packet starts with a fixed RTP header. The following fields of the RTP fixed header are specified for T.140 text streams:

Payload Type (PT): The assignment of an RTP payload type is specific to the RTP profile under which this payload format is used. For profiles that use dynamic payload type number assignment, this payload format can be identified by the MIME type "audio/t140c" (see section 10). If redundancy is used per RFC 2198, another payload type number needs to be provided for the redundancy format. The MIME type for identifying RFC 2198 is available in RFC 3555 [17].

Sequence number: The definition of sequence numbers is available in RFC 3550 [2]. Character loss is detected through the T140block counter when using the audio/t140c payload format.

Timestamp: The RTP Timestamp encodes the approximate instance of entry of the primary text in the packet. For audio/t140c, the clock frequency MAY be set to any value, and SHOULD be set to the same value as for any audio packets in the same RTP stream in order to avoid RTP timestamp rate switching. The value SHOULD be set by out of band mechanisms. Sequential packets MUST NOT use the same timestamp. Since packets do not represent any constant duration, the timestamp cannot be used to directly infer packet loss.

M-bit: The M-bit MUST be included. The first packet in a session, and the first packet after an idle period, SHOULD be distinguished by setting the marker bit in the RTP data header to one. The

marker bit in all other packets MUST be set to zero. The reception of the marker bit MAY be used for refined methods for detection of loss.

4. Protection against Loss of Data

Consideration must be devoted to keeping loss of text caused by packet loss within acceptable limits. (See ITU-T F.703 [16].)

The default method that MUST be used when no other method is explicitly selected is redundancy in accordance with RFC 2198 [3]. When this method is used, the original text and two redundant generations SHOULD be transmitted if the application or end-to-end conditions do not call for other levels of redundancy to be used.

Other protection methods MAY be used. Forward Error Correction mechanisms as per RFC 2733 [8] or any other mechanism with the purpose of increasing the reliability of text transmission MAY be used as an alternative or complement to redundancy. Text data MAY be sent without additional protection if end-to-end network conditions allow the text quality requirements specified in ITU-T F.703 [16] to be met in all anticipated load conditions.

4.1. Payload Format When Using Redundancy

When using the format with redundant data, the transmitter may select a number of T140block generations to retransmit in each packet. A higher number introduces better protection against loss of text but marginally increases the data rate.

The RTP header is followed by one or more redundant data block headers, one for each redundant data block to be included. Each of these headers provides the timestamp offset and length of the corresponding data block plus a payload type number indicating the payload format audio/t140c.

After the redundant data block headers follows the redundant data fields carrying T140blocks from previous packets, and finally the new (primary) T140block for this packet.

Redundant data that would need a timestamp offset higher than 16383 due to its age at transmission MUST NOT be included in transmitted packets.

4.2. Using Redundancy with the audio/t140c Format

Since sequence numbers are not provided in the redundant header and since the sequence number space is shared by all audio payload types within an RTP session, a sequence number in the form of a T140block counter is added to the T140block for transmission. This allows the redundant T140block data corresponding to missing primary data to be retrieved and used properly into the stream of received T140block data when using the audio/t140c payload format.

All non-empty redundant data blocks MUST contain the same data as a T140block previously transmitted as primary data, and be identified with a T140block counter equating to the original T140block counter for that T140block.

The T140block counters preceding the text in the T140block enables the ordering by the receiver. If there is a gap in the T140block counter value of received audio/t140c packets, and if there are redundant T140blocks with T140block counters matching those that are missing, the redundant T140blocks may be substituted for the missing T140blocks.

The value of the length field in the redundant header indicates the length of the concatenated T140block counter and the T140block.

5. Recommended Procedure

This section contains RECOMMENDED procedures for usage of the payload format. Based on the information in the received packets, the receiver can:

- reorder text received out of order.
- mark where text is missing because of packet loss.
- compensate for lost packets by using redundant data.

5.1. Recommended Basic Procedure

Packets are transmitted when there is valid T.140 data to transmit.

T.140 specifies that T.140 data MAY be buffered for transmission with a maximum buffering time of 500 ms. A buffering time of 300 ms is RECOMMENDED when the application or end-to-end network conditions are not known to require another value.

If no new data is available for a longer period than the buffering time, the transmission process is in an idle period.

When new text is available for transmission after an idle period, it is RECOMMENDED to send it as soon as possible. After this transmission, it is RECOMMENDED to buffer T.140 data in buffering time intervals until next idle period. This is done in order to keep the maximum bit-rate usage for text at a reasonable level. The buffering time MUST be selected so that text users will perceive a real-time text flow.

5.2. Transmission before and after "Idle Periods"

When valid T.140 data has been sent and no new T.140 data is available for transmission after the selected buffering time, an empty T140block SHOULD be transmitted. This situation is regarded to be the beginning of an idle period. The procedure is recommended in order to more rapidly detect potentially missing text before an idle period or when the audio stream switches from the transmission of audio/t140c to some other form of audio.

An empty T140block contains no data, neither T.140 data nor a T140block counter.

When redundancy is used, transmission continues with a packet at every transmission timer expiration and insertion of an empty T.140block as primary, until the last non-empty T140block has been transmitted as primary and as redundant data with all intended generations of redundancy. The last packet before an idle period will contain only one non-empty T140block as redundant data, and the empty primary T140block.

When using the audio/t140c payload format, empty T140blocks sent as primary data SHOULD NOT be included as redundant T140blocks, as it would simply be a waste of bandwidth to send them and it would introduce a risk of false detection of loss.

After an idle period, the transmitter SHOULD set the M-bit to one in the first packet with new text.

5.3. Detection of Lost Text Packets

Receivers detect the loss of an audio/t140c packet by observing the value of the T140block counter in a subsequent audio/t140c packet.

Missing data SHOULD be marked by insertion of a missing text marker in the received stream for each missing T140block, as specified in ITU-T T.140 Addendum 1 [1].

Procedures based on detection of the packet with the M-bit set to one MAY be used to reduce the risk for introducing false markers of loss.

False detection will also be avoided when using audio/t140c by observing the value of the T140block counter value.

If two successive packets have the same number of redundant generations, it SHOULD be treated as the general redundancy level for the session. Change of the general redundancy level SHOULD only be done after an idle period.

5.4. Compensation for Packets Out of Order

For protection against packets arriving out of order, the following procedure MAY be implemented in the receiver. If analysis of a received packet reveals a gap in the sequence and no redundant data is available to fill that gap, the received packet SHOULD be kept in a buffer to allow time for the missing packet(s) to arrive. It is RECOMMENDED that the waiting time be limited to 1 second.

If a packet with a T140block belonging to the gap arrives before the waiting time expires, this T140block is inserted into the gap and then consecutive T140blocks from the leading edge of the gap may be consumed. Any T140block that does not arrive before the time limit expires should be treated as lost and a missing text marker inserted (see section 5.3).

6. Parameter for Character Transmission Rate

In some cases, it is necessary to limit the rate at which characters are transmitted. For example, when a PSTN gateway is interworking between an IP device and a PSTN textphone, it may be necessary to limit the character rate from the IP device in order to avoid throwing away characters in case of buffer overflow at the PSTN gateway.

To control the character transmission rate, the MIME parameter "cps" in the "fmt" attribute [7] is defined (see section 10). It is used in Session Description Protocol (SDP) with the following syntax:

```
a=fmt:<format> cps=<integer>
```

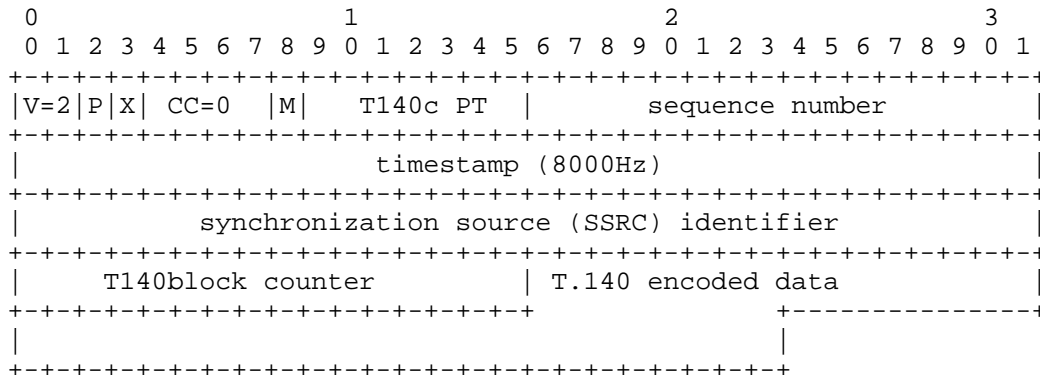
The <format> field is populated with the payload type that is used for text. The <integer> field contains an integer representing the maximum number of characters that may be received per second. The value shall be used as a mean value over any 10-second interval. The default value is 30.

In receipt of this parameter, devices MUST adhere to the request by transmitting characters at a rate at or below the specified <integer> value. Examples of use in SDP are found in section 7.2.

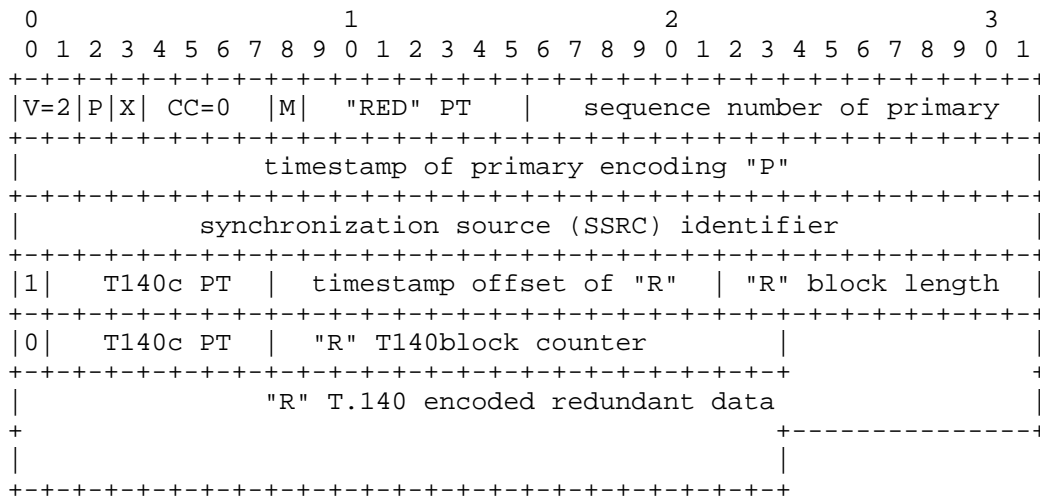
7. Examples

7.1. RTP Packetization Examples for the audio/t140c Format

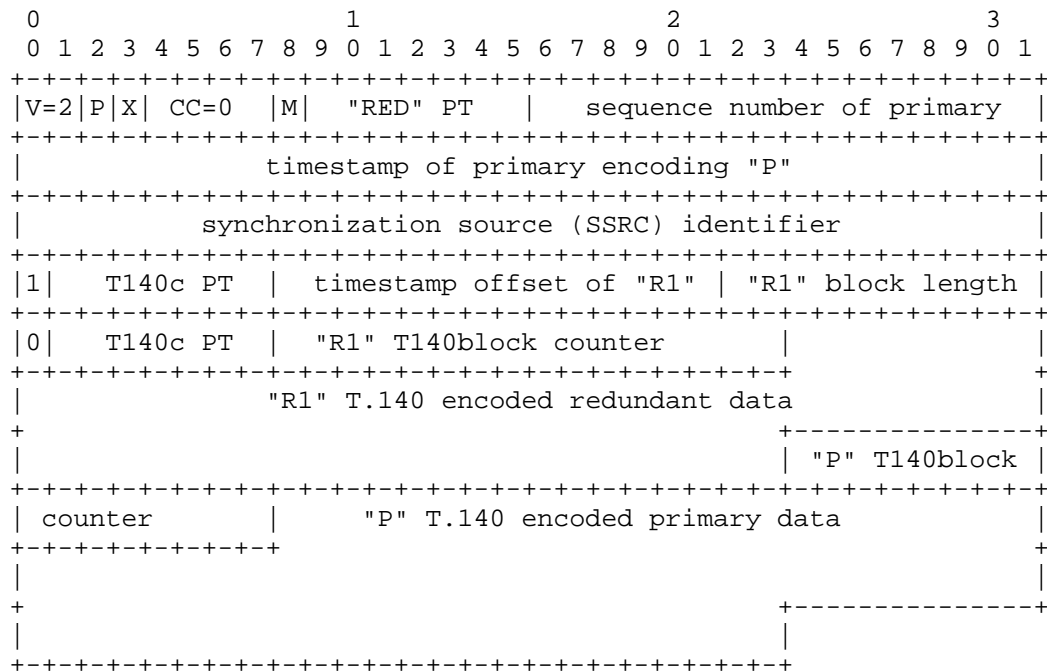
Below is an example of an audio/t140c RTP packet without redundancy.



Below is an example of an RTP packet with one redundant T140block using audio/t140c payload format. The primary data block is empty, which is the case when transmitting a packet for the sole purpose of forcing the redundant data to be transmitted in the absence of any new data. Note that since this is the audio/t140c payload format, the redundant block of T.140 data is immediately preceded with a T140block counter.



As a follow-on to the previous example, the example below shows the next RTP packet in the sequence that does contain a new real T140block when using the audio/t140c payload format. This example has 2 levels of redundancy and one primary data block. Since the previous primary block was empty, no redundant data is included for that block. This is because when using the audio/t140c payload format, any previously transmitted "empty" T140blocks are NOT included as redundant data in subsequent packets.



7.2. SDP Examples

Below is an example of SDP describing RTP text interleaved with G.711 audio packets within the same RTP session from port 7200 and at a maximum text rate of 6 characters per second:

```

m=audio 7200 RTP/AVP 0 98
a=rtpmap:98 t140c/8000
a=fmtp:98 cps=6

```

Below is an example using RFC 2198 to provide the recommended two levels of redundancy to the text packets in an RTP session with interleaving text and G.711 at a text rate no faster than 20 characters per second:

```
m=audio 7200 RTP/AVP 0 98 100
a=rtpmap:98 t140c/8000
a=fmtp:98 cps=20
a=rtpmap:100 red/8000
a=fmtp:100 98/98/98
```

Note: While these examples utilize the RTP/AVP profile, it is not intended to limit the scope of this memo to use with only that profile. Rather, any appropriate profile may be used in conjunction with this memo.

8. Security Considerations

All of the security considerations from section 14 of RFC 3550 [2] apply.

8.1. Confidentiality

Since the intention of the described payload format is to carry text in a text conversation, security measures in the form of encryption are of importance. The amount of data in a text conversation session is low, and therefore any encryption method MAY be selected and applied to T.140 session contents or to the whole RTP packets. Secure Realtime Transport Protocol (SRTP) [13] provides a suitable method for ensuring confidentiality.

8.2. Integrity

It may be desirable to protect the text contents of an RTP stream against manipulation. SRTP [13] provides methods for providing integrity that MAY be applied.

8.3. Source Authentication

Measures to make sure that the source of text is the intended one can be accomplished by a combination of methods.

Text streams are usually used in a multimedia control environment. Security measures for authentication are available and SHOULD be applied in the registration and session establishment procedures, so that the identity of the sender of the text stream is reliably associated with the person or device setting up the session. Once established, SRTP [13] mechanisms MAY be applied to ascertain that the source is maintained the same during the session.

9. Congestion Considerations

The congestion considerations from section 10 of RFC 3550 [2], section 6 of RFC 2198 [3], and any used profile (e.g., the part about congestion in section 2 of RFC 3551 [10]) apply with the following application-specific considerations.

Automated systems MUST NOT use this format to send large amounts of text at a rate significantly above that which a human user could enter.

Even if the network load from users of text conversation is usually very low, for best-effort networks an application MUST monitor the packet loss rate and take appropriate actions to reduce its sending rate if this application sends at higher rate than what TCP would achieve over the same path. The reason is that this application, due to its recommended usage of two or more redundancy levels, is very robust against packet loss. At the same time, due to the low bit-rate of text conversations, if one considers the discussion in RFC 3714 [12], this application will experience very high packet loss rates before it needs to perform any reduction in the sending rate.

If the application needs to reduce its sending rate, it SHOULD NOT reduce the number of redundancy levels below the default amount specified in section 4. Instead, the following actions are RECOMMENDED in order of priority:

- Increase the shortest time between transmissions described in section 5.1 from the recommended 300 ms to 500 ms that is the highest value allowable according to T.140.
- Limit the maximum rate of characters transmitted.
- Increase the shortest time between transmissions to a higher value, not higher than 5 seconds. This will cause unpleasant delays in transmission, beyond what is allowed according to T.140, but text will still be conveyed in the session with some usability.
- Exclude participants from the session.

Please note that if the reduction in bit-rate achieved through the above measures is not sufficient, the only remaining action is to terminate the session.

As guidance, some load figures are provided here as examples based on use of IPv4, including the load from IP, UDP, and RTP headers without compression.

- Experience tells that a common mean character transmission rate during a complete PSTN text telephony session in reality is around 2 characters per second.
- A maximum performance of 20 characters per second is enough even for voice-to-text applications.
- With the (unusually high) load of 20 characters per second, in a language that make use of three-octet UTF-8 characters, two redundant levels, and 300 ms between transmissions, the maximum load of this application is 3500 bits/s.
- When the restrictions mentioned above are applied, limiting transmission to 10 characters per second, using 5 s between transmissions, the maximum load of this application in a language that uses one octet per UTF-8 character is 300 bits/s.

Note also, that this payload can be used in a congested situation as a last resort to maintain some contact when audio and video media need to be stopped. The availability of one low bit-rate stream for text in such adverse situations may be crucial for maintaining some communication in a critical situation.

10. IANA Considerations

This document defines one RTP payload format named "t140" and an associated MIME type "audio/t140c". They have been registered by the IANA.

10.1. Registration of MIME Media Type audio/t140c

MIME media type name: audio

MIME subtype name: t140c

Required parameters:

rate: The RTP timestamp clock rate, which is equal to the sampling rate. This parameter SHOULD have the same value as for any audio codec packets interleaved in the same RTP stream.

Optional parameters:

cps: The maximum number of characters that may be received per second. The default value is 30.

Encoding considerations: T.140 text can be transmitted with RTP as specified in RFC 4351.

Security considerations: See section 8 of RFC 4351.

Interoperability considerations: None

Published specification: ITU-T T.140 Recommendation.
RFC 4351.

Applications which use this media type:

Text communication systems and text conferencing tools that transmit text associated with audio and within the same RTP session as the audio, such as PSTN gateways that transmit audio and text signals between two PSTN textphone users over an IP network.

Additional information: This type is only defined for transfer via RTP.

Magic number(s): None
File extension(s): None
Macintosh File Type Code(s): None

Person & email address to contact for further information:

Paul E. Jones
E-mail: paulej@packetizer.com

Intended usage: COMMON

Author	/ Change controller:
Paul E. Jones	IETF avt WG delegated from the IESG
paulej@packetizer.com	

10.2. SDP Mapping of MIME Parameters

The information carried in the MIME media type specification has a specific mapping to fields in the Session Description Protocol (SDP) [7], which is commonly used to describe RTP sessions. When SDP is used to specify sessions employing the audio/t140c format, the mapping is as follows:

- The MIME type ("audio") goes in SDP "m=" as the media name.
- The MIME subtype (payload format name) goes in SDP "a=rtpmap" as the encoding name. For audio/t140c, the clock rate MAY be set to any value, and SHOULD be set to the same value as for any audio packets in the same RTP stream.
- The parameter "cps" goes in SDP "a=fmtp" attribute.

- When the payload type is used with redundancy according to RFC 2198, the level of redundancy is shown by the number of elements in the slash-separated payload type list in the "fmt" parameter of the redundancy declaration as defined in RFC 2198 [3].

10.3. Offer/Answer Consideration

In order to achieve interoperability within the framework of the offer/answer model [9], the following consideration should be made:

- The "cps" parameter is declarative. Both sides may provide a value, which is independent of the other side.

11. Acknowledgements

The authors want to thank Stephen Casner, Magnus Westerlund, and Colin Perkins for valuable support with reviews and advice on creation of this document; Mickey Nasiri at Ericsson Mobile Communication for providing the development environment; Michele Mizarro for verification of the usability of the payload format for its intended purpose; and Andreas Piirimets for editing support.

12. Normative References

- [1] ITU-T Recommendation T.140 (1998) - Text conversation protocol for multimedia application, with amendment 1, (2000).
- [2] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [3] 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, September 1997.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [5] ISO/IEC 10646-1: (1993), Universal Multiple Octet Coded Character Set.
- [6] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [7] Handley, M. and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

- [8] Rosenberg, J. and H. Schulzrinne, "An RTP Payload Format for Generic Forward Error Correction", RFC 2733, December 1999.
- [9] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [10] Schulzrinne, H. and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, RFC 3551, July 2003.
- [11] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

13. Informative References

- [12] Floyd, S. and J. Kempf, "IAB Concerns Regarding Congestion Control for Voice Traffic in the Internet", RFC 3714, March 2004.
- [13] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K. Norrman, "The Secure Real-time Transport Protocol (SRTP)", RFC 3711, March 2004.
- [14] Schulzrinne, H. and S. Petrack, "RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals", RFC 2833, May 2000.
- [15] Hellstrom, G. and P. Jones, "RTP Payload for Text Conversation", RFC 4103, June 2005.
- [16] ITU-T Recommendation F.703, Multimedia Conversational Services, Nov 2000.
- [17] Casner, S. and P. Hoschka, "MIME Type Registration of RTP Payload Formats", RFC 3555, July 2003.

Authors' Addresses

Gunnar Hellstrom
Omnitor AB
Renathvagen 2
SE-121 37 Johanneshov
Sweden

Phone: +46 708 204 288 / +46 8 556 002 03
Fax: +46 8 556 002 06
EMail: gunnar.hellstrom@omnitor.se

Paul E. Jones
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
USA

Phone: +1 919 392 6948
EMail: paulej@packetizer.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

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

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

Intellectual Property

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

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

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

Acknowledgement

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