

RTP Payload for Text Conversation

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

This memo describes how to carry text conversation session contents in RTP 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.

This RTP payload description contains an optional possibility to include redundant text from already transmitted packets in order to reduce the risk of text loss caused by packet loss. The redundancy coding follows RFC 2198.

1. Introduction

This memo defines a payload type for carrying text conversation session contents in RTP 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 text conversation sessions is sent as soon as it is available, or with a small delay for buffering.

The text is supposed 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. Therefore, the expected number of characters to transmit is low. Only one or a few new characters are expected to be transmitted with each packet.

T.140 specifies that text and other T.140 elements MUST be transmitted in ISO 10 646-1 code with UTF-8 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. If lost information can be indicated, the willingness to accept loss is expected to be higher.

Therefore a mechanism based on RTP is specified here. It gives text arrival in correct order, without duplications, and with detection and indication of losses. 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 channel load by the redundancy scheme is minimal.

1.1 Terminology

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]

2. Usage of RTP

When transport of T.140 text session data in RTP is desired, the payload as described in this specification SHOULD be used.

A text conversation RTP packet as specified by this payload format consists of an RTP header as defined in RFC 1889 [2] followed immediately by a block of T.140 data, defined here to be a "T140block". There is no additional header specific to this payload format. 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. This implies that each block MUST contain an integral number of UTF-8 encoded

characters regardless of the number of octets per character. It also implies that any composite character sequence (CCS) SHOULD be placed within one block.

The T140blocks MAY be transmitted redundantly according to the payload format defined in RFC 2198 [3]. In that case, the RTP header is followed by one or more redundant data block headers, the same number of redundant data fields carrying T140blocks from previous packets, and finally the new (primary) T140block for this packet.

2.1 RTP packet header

Each RTP packet starts with a fixed RTP header. The following fields of the RTP fixed header are used 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 which use dynamic payload type number assignment, this payload format is identified by the name "T140" (see section 6). If redundancy is used per RFC 2198, the Payload Type MUST indicate that payload format ("RED").

Sequence number: The Sequence Number MUST be increased by one for each new transmitted packet. It is used for detection of packet loss and packets out of order, and can be used in the process of retrieval of redundant text, reordering of text and marking missing text.

Timestamp: The RTP Timestamp encodes the approximate instance of entry of the primary text in the packet. A clock frequency of 1000 Hz MUST be used. 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 losses.

2.2 Additional headers

There are no additional headers defined specific to this payload format.

When redundant transmission of the data according to RFC 2198 is desired, 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 this payload format ("T140").

2.3 T.140 Text structure

T.140 text is UTF-8 coded as specified in T.140 with no extra framing. 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 increases the data rate.

Since packets are not generated at regular intervals, the timestamp is not sufficient to identify a packet in the presence of loss unless extra information is provided. Since sequence numbers are not provided in the redundant header, some additional rules must be followed to allow the redundant data corresponding to missing primary data to be merged properly into the stream of primary data T140blocks:

- Each redundant data block MUST contain the same data as a T140block previously transmitted as primary data, and be identified with a timestamp offset equating to the original timestamp for that T140block.
- The redundant data MUST be placed in age order with most recent redundant T140block last in the redundancy area.
- All T140blocks from the oldest desired generation up through the generation immediately preceding the new (primary) T140block MUST be included.

These rules allow the sequence numbers for the redundant T140blocks to be inferred by counting backwards from the sequence number in the RTP header. The result will be that all the text in the payload will be contiguous and in order.

3. Recommended procedures

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.

3.1 Recommended basic procedure

Packets are transmitted only when there is valid T.140 data to transmit. The sequence number is used for sequencing of T.140 data.

On reception, the RTP sequence number is compared with the sequence number of the last correctly received packet. If they are consecutive, the (only or primary) T140block is retrieved from the packet.

3.2 Recommended procedure for compensation for lost packets.

For reduction of data loss in case of packet loss, redundant data MAY be included in the packets following to the procedures in RFC 2198. If network conditions are not known, it is RECOMMENDED to use one redundant T140block in each packet. If there is a gap in the RTP sequence numbers, and redundant T140blocks are available in a subsequent packet, the sequence numbers for the redundant T140blocks should be inferred by counting backwards from the sequence number in the RTP header for that packet. If there are redundant T140blocks with sequence numbers matching those that are missing, the redundant T140blocks may be substituted for the missing T140blocks.

Both for the case when redundancy is used and not used, 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].

3.3 Recommended procedure for 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 can be kept in a buffer to allow time for the missing packet(s) to arrive. It is suggested that the waiting time be limited to 0.5 seconds. For the case when redundancy is used the waiting time SHOULD be extended to the number of redundancy generations times the T.140 buffering timer if this product is known to be greater than 0.5 seconds.

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 which does not arrive before the time limit expires should be treated as lost.

3.4 Transmission during "silent periods" when redundancy is used.


When using the redundancy transmission scheme, and there is nothing more to transmit from T.140, the latest T140block has a risk of getting old before it is transmitted as redundant data. The result is less useful protection against packet loss at the end of a text input sequence. For cases where this should be avoided, a zero-length primary T140block MAY be transmitted with the redundant data.

Any zero-length T140blocks that are sent as primary data MUST be included as redundant T140blocks on subsequent packets just as normal text T140blocks would be so that sequence number inference for the redundant T140blocks will be correct, as explained in section 2.3.

Redundancy for the last T140block SHOULD NOT be implemented by repeatedly transmitting the same packet (with the same sequence number) because this will cause the packet loss count, as reported in RTCP, to decrement.

4. Examples

This is an example of a T140 RTP packet without redundancy.



This is an example of an RTP packet with one redundant T140block.




Figure: Examples of RTP text packets.

5. Security Considerations

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. When redundant data is included, the same security considerations as for RFC 2198 apply.

6. MIME Media Type Registrations

This document defines a new RTP payload name and associated MIME type, T140 (text/t140).

6.1 Registration of MIME media type text/t140

MIME media type name: text

MIME subtype name: t140

Required parameters: None

Optional parameters: None

Encoding considerations: T140 text can be transmitted with RTP as specified in RFC 2793.

Security considerations: None

Interoperability considerations: None

Published specification: ITU-T T.140 Recommendation.
RFC 2793.

Applications which use this media type:

Text communication terminals and text conferencing tools.

Additional information: None

Magic number(s): None

File extension(s): None

Macintosh File Type Code(s): None

Person & email address to contact for further information:

Gunnar Hellstrom

e-mail: gunnar.hellstrom@omnitor.se

Intended usage: COMMON

Author	/ Change controller:
Gunnar Hellstrom	IETF avt WG
gunnar.hellstrom@omnitor.se	c/o Steve Casner casner@cisco.com

7. Author's Address

Gunnar Hellstrom
Omnitor AB
Alsnogatan 7, 4 tr
SE-116 41 Stockholm
Sweden

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

8. Acknowledgements

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

9. 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", RFC 1889, January 1996.
- [3] Perkins, C., Kouvelas, I., Hardman, V., Handley, M. and J. Bolot, "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", RFC 2279, January 1998.

10. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

