

Audio Video Transport WG
Internet-Draft
Updates: RFC [2198](#) (if approved)
Expires: July 2, 2007

Q. Xie
J. Schumacher
Motorola
December 29, 2006

Forward-shifted RTP Redundancy Payload Support
draft-xie-avt-forward-shifted-red-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on July 2, 2007.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document defines a simple enhancement to [RFC 2198](#) to support RTP sessions with forward-shifted redundant encodings, i.e., redundant data is sent before the corresponding primary data. Forward-shifted redundancy can be used to conceal losses of a large number of consecutive media frames (e.g., consecutive loss of seconds or even tens of seconds of media).

Table of Contents

1.	Conventions	3
2.	Introduction	3
3.	Allowing Forward-shifted Redundant Data	3
4.	Updated Registration of Media Type "red"	4
5.	Mapping MIME Parameters into SDP	5
6.	Usage in Offer/Answer	6
7.	Backward Compatibility with RFC 2198	6
8.	IANA Considerations	6
9.	Security Considerations	7
10.	Normative References	7
	Authors' Addresses	7
A.	Anti-shadow Loss Concealment Using Forward-shifted Redundancy	8
A.1	Sender Side Operations	8
A.2	Receiver Side Operations	10
A.2.1	Normal Mode Operation	10
A.2.2	Anti-shadow Mode Operation	11
	Intellectual Property and Copyright Statements	13

1. Conventions

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, NOT RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC 2119](#) [2].

2. Introduction

This document defines a simple enhancement to [RFC 2198](#) [3] to support RTP sessions with forward-shifted redundant encodings, i.e., redundant data is sent before the corresponding primary data.

Forward-shifted redundancy can be used to conceal losses of a large number of consecutive media frames (e.g., consecutive loss of seconds of media). Such capability is highly desirable, especially in wireless mobile communication environments where the radio signal to a mobile wireless media receiver can be temporarily blocked by tall buildings, mountains, tunnels, etc. In other words, the receiver enters into a shadow of the radio coverage. No new data will be received when the receiver is in a shadow.

In some extreme cases, the receiver may have to spend seconds or even tens of seconds in a shadow. The traditional backward-shifted redundant encoding scheme (i.e., redundant data must be sent after the primary data), as currently supported by [RFC 2198](#) [3], is known to be ineffective in dealing with such consecutive frame losses.

However, forward-shifted redundancy, in combination with the anti-shadow loss management at the receiver (as described in [Appendix A](#)) can effectively prevent service interruptions when a mobile receiver runs into such a shadow.

3. Allowing Forward-shifted Redundant Data

In [RFC 2198](#), the timestamp offset in the additional header corresponding to a redundant block is defined as a 14 bits unsigned offset of timestamp relative to timestamp given in the RTP header. As stated in [RFC 2198](#):

"The use of an unsigned offset implies that redundant data must be sent after the primary data, and is hence a time to be subtracted from the current timestamp to determine the timestamp of the data for which this block is the redundancy."

This effectively prevents [RFC 2198](#) from being used to support forward-shifted redundant blocks.

In order to support the use of forward-shifted redundant blocks, an optional MIME parameters, "forwardshift", is introduced for indicating the capability and willingness of using forward-shifted redundancy and the base value of timestamp forward-shifting. The base value of "forwardshift" is an integer equal or greater than '0'. In an RTP session which uses forward-shifted redundant encodings, the timestamp of a redundant block in a received RTP packet is determined as follows:

$$\begin{aligned} \text{timestamp of redundant block} &= \text{timestamp in RTP header} \\ &\quad - \text{timestamp offset in additional header} \\ &\quad + \text{forward shift base value} \end{aligned}$$

4. Updated Registration of Media Type "red"

(The definition is based on [RFC 2198](#) [3], added with the optional "forwardshift" parameter, and updated with the new template specified in [6].)

Type name: audio

Subtype names: red

Required parameters: none

Optional parameters:

forwardshift: An unsigned integer can be specified as value.

If this parameter is present with a value of '0', it indicates that the sender of this parameter is capable of and willing to receive forward-shifted redundant data, but there will be no forward shifting when it sends out its own redundant data.

If this parameter is present with a value greater than '0', it indicates that the sender of this parameter is capable of and willing to receive forward-shifted redundant data, and will use forward shifting with a base value as specified when sending out redundant data if forward-shifted redundancy is allowed for the session.

If this parameter is not present, it MUST be assumed that forward-shifted redundancy is not supported and MUST NOT be used in the session. In other words, forward-shifted redundancy can only be used in a session if both ends indicate the capability.

Encoding considerations:

This media type is framed binary data (see [RFC 4288, Section 4.8](#)) and is only defined for transfer of RTP redundant data frames specified in [RFC 2198](#).

Security considerations: See [Section 6](#) "Security Considerations" of [RFC 2198](#).

Interoperability considerations:

Existing [RFC 2198](#) implementations will not send "forwardshift" parameter in their SDP and will ignore "forwardshift" parameter they receive. As a result, forward-shifted redundancy will not be used in the session and thus ensures the interoperability between a legacy [RFC 2198](#) endpoint and a forward-shifted redundancy capable endpoint.

Published specification:

RTP redundant data frame format is specified in [RFC 2198](#).

Applications that use this media type:

It is expected that real-time audio/video applications that want protection against losses of a large number of consecutive frames will be interested in using this type.

Additional information: none

Person & email address to contact for further information:

Qiaobing Xie <Qiaobing.Xie@motorola.com>

Intended usage: COMMON

Restrictions on usage:

This media type depends on RTP framing, and hence is only defined for transfer via RTP ([RFC 3550](#) [4]). Transfer within other framing protocols is not defined at this time.

Author:

Qiaobing Xie

Change controller:

IETF Audio/Video Transport working group delegated from the IESG.

5. Mapping MIME Parameters into SDP

The information carried in the MIME media type specification has a specific mapping to fields in the Session Description Protocol (SDP) [5], which is commonly used to describe RTP sessions. When SDP is

used to specify sessions employing the forward-shifted redundant format, the mapping is as follows:

- o The MIME type ("audio") goes in SDP "m=" as the media name.
- o The MIME subtype ("red") goes in SDP "a=rtpmap" as the encoding name.
- o The optional parameter "forwardshift" goes in the SDP "a=fmtp" attribute by copying it directly from the MIME media type string as "forwardshift=value".

Example of usage of indicating forward-shifted (by 5.1 sec) redundancy:

```
m=audio 12345 RTP/AVP 121 0 5
a=rtpmap:121 red/8000/1
a=fmtp:121 0/5 forwardshift=40800
```

Example of usage of indicating only willingness for receiving forward-shifted redundancy (but won't send forward-shifted redundancy):

```
m=audio 12345 RTP/AVP 121 0 5
a=rtpmap:121 red/8000/1
a=fmtp:121 0/5 forwardshift=0
```

6. Usage in Offer/Answer

The optional "forwardshift" SDP parameter specified in this document is declarative, and all reasonable values are expected to be supported. Furthermore, if "forwardshift" parameter is absent from either Offer or Answer, forward-shifted redundancy MUST NOT be used in the session.

7. Backward Compatibility with [RFC 2198](#)

Existing [RFC 2198](#) implementations will not send "forwardshift" parameter in their SDP and will ignore "forwardshift" parameter they receive. As a result, forward-shifted redundancy will not be used in the session and thus ensures the interoperability between a legacy [RFC 2198](#) endpoint and a forward-shifted redundancy capable endpoint.

8. IANA Considerations

The MIME subtype registration, "red" originally defined in [RFC 2198](#) [3], is updated with the additional optional parameter "forwardshift"

(see [Section 4](#)).

9. Security Considerations

See [Section 6](#) "Security Considerations" of [RFC 2198](#) [3].

10. Normative References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", [RFC 3550](#), July 2003.
- [5] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [6] Casner, S., "Media Type Registration of RTP Payload Formats", [draft-ietf-avt-rfc3555bis-05.txt](#) (work in progress), October 2006.

Authors' Addresses

Qiaobing Xie
Motorola, Inc.
1501 W. Shure Drive, 2-F9
Arlington Heights, IL 60004
US

Phone: +1-847-632-3028
Email: Qiaobing.Xie@Motorola.com

Joe Schumacher
Motorola, Inc.
1501 W. Shure Drive, 2-B11
Arlington Heights, IL 60004
US

Phone: +1-847-632-5978

Email: j.schumacher@motorola.com

Appendix A. Anti-shadow Loss Concealment Using Forward-shifted Redundancy

It is not unusual in a wireless mobile communication environment where the radio signal to a mobile wireless media receiver can be temporarily blocked by tall buildings, mountains, tunnels, etc. for a period of time. In other words, the receiver enters into a shadow of the radio coverage. When the receiver is in such a shadow no new data will be received. In some extreme cases, the receiver may have to spend seconds or even tens of seconds in such a shadow.

Without special design considerations to compensate the loss of data due to shadowing, a mobile user may experience an unacceptable level of service interruptions. And traditional redundant encoding schemes (including [RFC 2198](#) and most FEC schemes) are known to be ineffective in dealing with such losses of consecutive frames.

However, the employment of forward-shifted redundancy, in combination with the anti-shadow loss concealment at the receiver, as described here, can effectively prevent service interruptions due to the effect of shadowing.

A.1 Sender Side Operations

For anti-shadow loss management, the RTP sender simply adds a forward-shifted redundant stream (called anti-shadow or AS stream) while transmitting the primary media stream. The amount of forward-shifting, which should remain constant for the duration of the session, will determine the maximal length of shadows that can be completely concealed at the receiver, as explained below.

Except for the fact that it is forward-shifted relative to the primary stream (i.e., the redundant data is sent ahead of the corresponding primary data), the design decision and trade-offs on the quality, encoding, bandwidth overhead, etc. of the redundant stream is not different from the traditional RTP payload redundant scheme.

The following diagram illustrates a segment of the transmission

sequence of a forward-shifted redundant RTP session, in which the AS stream is forward-shifted by 155 frames. If, for simplicity here, we assume the value of timestamp is incremented by 1 between two consecutive frames, this forward-shifted redundancy can then be indicated with:

forwardshift=155

and the setting of timestamp offset to 0 in all the additional headers. This can mean a 3.1 second of forward shifting if each frame represents 20 ms of original media,

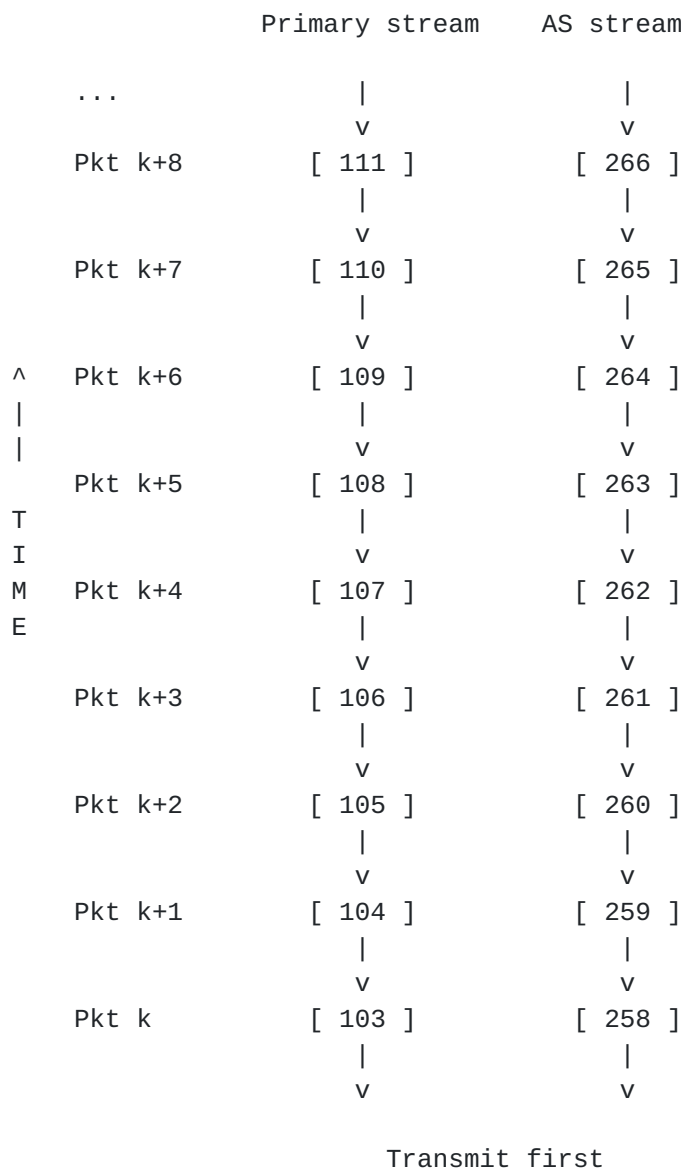


Figure 1. An example of forward-shifted redundant RTP packet transmission.

A.2 Receiver Side Operations

The anti-shadow receiver is illustrated in the following diagram.

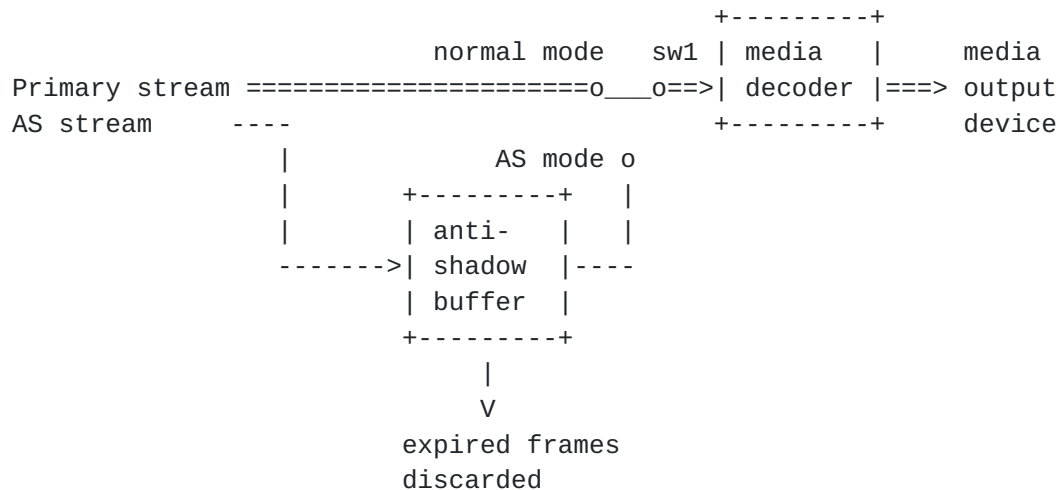


Figure 2. Anti-shadow RTP receiver.

The anti-shadow receiver operates between two modes - "normal mode" and "AS mode". When the receiver is not in a shadow (it can easily tell that if it is still receiving new data), the receiver operates in the normal mode. Otherwise, it operates in the AS mode.

A.2.1 Normal Mode Operation

In the normal mode, after receiving a new RTP packet that contains the primary data and forward-shifted AS data, the receiver passes the primary data directly to the appropriate media decoder for play-out (a de-jittering buffer may be used before the play-out, but for simplicity we assume none is used here), while the received AS data is stored in an anti-shadow buffer.

Moreover, data stored in the anti-shadow buffer will be continuously checked to determine whether it has expired. If a redundant data in the anti-shadow buffer is found to have a timestamp older (i.e., smaller) than that of the last primary frame passed to the media decoder, it will be considered expired and be purged from the anti-shadowing buffer.

The following example illustrates the operation of the anti-shadow buffer in normal mode. We use the same assumption as in Figure 1, and assume that the initial timestamp value is 103 when the session starts.

Time (in ms)	Timestamp being played out	Timestamp of media in AS buffer	Note
t < 0		--	(buffer empty)
...			
t=0	103	258	(hold 1 AS frame)
t=20	104	258-259	(hold 2 AS frames)
t=40	105	258-260	(hold 3 AS frames)
...			
t=3080	257	258-412	(full, hold 154 AS frames)
t=3100	258	259-413	(full, frame 258 purged)
t=3120	259	260-414	(full, frame 259 purged)
...			
t=6240	415	416-570	(always holds 3.08 sec worth of redundant data)
...			

Figure 3. Example of anti-shadow buffer operation in normal mode.

At the beginning of the session ($t=0$), the anti-shadow buffer will be empty. When the first primary frame is received, the play-out will start immediately, and the first received AS frame is stored in the anti-shadow buffer. And with the arriving of more forward-shifted redundant frames, the anti-shadow buffer will gradually be filled up.

For the example shown in Figure 1, after 3.08 seconds (the amount of the forward-shifting minus one frame) from the start of the session, the anti-shadow buffer will be full, holding exactly 3.08 seconds worth of redundant data, with the oldest frame corresponding to $t=3.1$ sec and youngest frame corresponding to $t=6.18$ sec.

And it is not difficult to see that in normal mode because of the continuous purge of expired frames and the addition of new frames, the anti-shadowing buffer will always be full holding the next forward-shift amount of redundant frames.

[A.2.2](#) Anti-shadow Mode Operation

When the receiver enters a shadow (or any other conditions that prevent the receiver from getting new media data), the receiver switches to the anti-shadow mode, in which it simply continues the play-out from the forward-shifted redundant data stored in the anti-shadow buffer.

For the example in Figure 3, if the receiver enters a shadow at $t=3120$, it can continue the play-out by using the forward-shifted redundant frames ($ts=260-414$) from the anti-shadow buffer. As far as the receiver can move out of the shadow by $t=6240$, there will be no service interruption.

When the shadow condition ends (meaning new data starts to arrive again), the receiver immediately switches back to normal mode of operation, playing out from newly arrived primary frames. And at the same time, the arrival of new AS frames will start to re-fill the anti-shadow buffer.

However, if the duration of the shadow is longer than the amount of forward-shifting, the receiver will run out of media frames from its anti-shadow buffer. At that point, service interruption will occur.

Anti-shadow loss concealment described above can be readily applied to the streaming of pre-recorded media. Because of the need of generating the forward-shifted anti-shadow redundant stream, to apply anti-shadow loss concealment to the streaming of live media will require the insertion of a delay equal to or greater than the amount of forward-shifting at the source of media.

Intellectual Property Statement

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.

Disclaimer of Validity

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.

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.

Acknowledgment

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

