Duplicating RTP Streams
draft-begen-avtcore-rtp-duplication-01

Abstract

   Packet loss is undesirable for real-time multimedia sessions, but can
   occur due to congestion, or other unplanned network outages. This is
   especially true for IP multicast networks, where packet loss patterns
   can vary greatly between receivers. One technique that can be used
   to recover from packet loss without incurring unbounded delay for all
   the receivers is to duplicate the packets and send them in separate
   redundant streams. This document explains how Real-time Transport
   Protocol (RTP) streams can be duplicated without breaking RTP media
   streams, or RTP Control Protocol (RTCP) rules.

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 http://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 September 11, 2012.

Table of Contents

1. **Introduction**

   The Real-time Transport Protocol (RTP) [RFC3550] is widely used today
   for delivering IPTV traffic, and other real-time multimedia sessions.
   Many of these applications support very large numbers of receivers,
   and rely on intra-domain UDP/IP multicast for efficient distribution
   of traffic within the network.

   While this combination has proved successful, there does exist a
   weakness. As [RFC2354] noted, packet loss is not avoidable, even in
   a carefully managed network. This loss might be due to congestion,
   it might also be a result of an unplanned outage caused by a flapping
   link, link or interface failure, a software bug, or a maintenance
   person accidentally cutting the wrong fiber. Since UDP/IP flows do
   not provide any means for detecting loss and retransmitting packets,
   it leaves up to the RTP layer and the applications to detect, and
   recover from, packet loss.

   One technique to recover from packet loss without incurring unbounded
   delay for all the receivers is to duplicate the packets and send them
   in separate redundant streams. Variations on this idea have been
   implemented and deployed today [IC2011]. However, duplication of RTP
   streams without breaking the RTP and RTCP functionality has not been
   documented properly. This document explains how duplication can be
   achieved for RTP streams.

   Stream duplication offers a simple way to protect media flows from
   packet loss. It has a comparatively high bandwidth overhead, since
   everything is sent twice, but with a low processor overhead. It is
   also very predictable in its overheads. Alternative approaches may
   be suitable in some cases, for example retransmission-based recovery
   [RFC4588] or forward error correction [RFC5109].

2. **Terminology and Requirements Notation**

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and
   "OPTIONAL" in this document are to be interpreted as described in
   [RFC2119].

3. **Dual Streaming Use Cases**

   Dual streaming refers to a technique that involves transmitting two
   redundant RTP streams of the same content, with each stream capable
   of supporting the playback when there is no packet loss. Therefore,
   adding an additional RTP stream provides a protection against packet

loss. The level of protection depends on how the packets are sent
and transmitted inside the network.

It is important to note that dual streaming can easily be extended to
support cases when more than two streams are desired. However, using
three or more streams is rare in practise, due to the high overhead
that it incurs.

3.1. Temporal Redundancy

From a routing perspective, two streams are considered identical if
the following two IP header fields are the same, since they will be
both routed over the same path:

o IP Source Address

o IP Destination Address

Two routing-plane identical RTP streams might carry the same payload,
but can use different Synchronization Sources (SSRC) to differentiate
the RTP packets belonging to each stream. In the context of dual RTP
streaming, we assume that the source duplicates the RTP packets and
sends them in separate RTP streams, each with a unique SSRC. All the
redundant streams are transmitted in the same RTP session.

For example, one main and one redundant RTP stream can be sent to the
same IP destination address and UDP destination port with a certain
delay between them [I-D.begen-mmusic-temporal-interleaving]. The
streams carry the same payload in their respective RTP packets with
identical sequence numbers. This allows receivers (or other nodes
responsible for gap filling and duplicate suppression) to identify
and suppress the duplicate packets, and subsequently produce a
hopefully loss-free and duplication-free output stream. This process
is called stream merging.

3.2. Spatial Redundancy

An RTP source might be associated with multiple network interfaces,
allowing it to send two redundant streams from two separate source
addresses. Such streams can be routed over diverse or identical
paths depending on the routing algorithm used inside the network. At
the receiving end, the node responsible for duplicate suppression can
look into various RTP header fields, for example SSRC and sequence
number, to identify and suppress the duplicate packets.

If source-specific multicast (SSM) transport is used to carry such
redundant streams, there will be a separate SSM session for each
redundant stream since the streams are sourced from different

interfaces (i.e., IP addresses). Thus, the receiving host has to
join each SSM session separately.

Alternatively, an RTP source might send the redundant streams to
separate IP destination addresses.

3.3. Dual Streaming over a Single Path or Multiple Paths

Having described the characteristics of the streams, one can reach
the following conclusions:

1. When two routing-plane identical streams are used, the two
    streams will have identical IP headers. This makes it
    impractical to forward the packets onto different paths. In
    order to minimize packet loss, the packets belonging to one
    stream are often interleaved with packets belonging to the other,
    and with a delay, so that if there is a packet loss, such a delay
    would allow the same packet from the other stream to reach the
    receiver because the chances that the same packet is lost in
    transit again is often small. This is what is also known as
    Time-shifted Redundancy, Temporal Redundancy or simply Delayed
    Duplication [I-D.begen-mmusic-temporal-interleaving] [IC2011].
    This approach can be used with both types of dual streaming,
    described in Section 3.1 and Section 3.2.

2. If the two streams have different IP headers, an additional
    opportunity arises in that one is able to build a network, with
    physically diverse paths, to deliver the two streams concurrently
    to the intended receivers. This reduces the delay when packet
    loss occurs and needs to be recovered. Additionally, it also
    further reduces chances for packet loss. An unrecoverable loss
    happens only when two network failures happen in such a way that
    the same packet is affected on both paths. This is referred to
    as Spatial Diversity or Spatial Redundancy [IC2011]. The
    techniques used to build diverse paths are beyond the scope of
    this document.

    Note that spatial redundancy often offers less delay in
    recovering from packet loss provided that the forwarding delay of
    the network paths are more or less the same. For both temporal
    and spatial redundancy approaches, packet misordering might still
    happen and needs to be handled using the sequence numbers of some
    sort (e.g., RTP sequence numbers).

To summarize, dual streaming allows an application and a network to
work together to provide a near zero-loss transport with a bounded or
minimum delay. The additional advantage includes a predictable
bandwidth overhead that is proportional to the minimum bandwidth

needed for the multimedia session, but independent of the number of
receivers experiencing a packet loss and requesting a retransmission.
For a survey and comparison of similar approaches, refer to [IC2011].

4. Use of RTP and RTCP with Temporal Redundancy

To achieve temporal redundancy, the main and redundant RTP streams
MUST be sent using the same 5-tuple of transport protocol, source and
destination IP addresses, and source and destination transport ports.
This is perhaps overly restrictive, but with the possible presence of
network address and port translation (NAPT) devices, using anything
other than an identical 5-tuple can also cause spatial redundancy.

Since main and redundant RTP streams follow an identical path, they
are part of the same RTP session. Accordingly, the sender MUST
choose a different SSRC for the redundant RTP stream than it chose
for the main RTP stream, following the rules in [RFC3550] Section 8.

4.1. RTCP Considerations

If RTCP is being sent for the main RTP stream, then the sender MUST
also generate RTCP for the redundant RTP stream. The RTCP for the
redundant RTP stream is generated exactly as-if the redundant RTP
stream were a regular media stream. The sender MUST NOT duplicate
the RTCP packets sent for the main RTP stream when sending the
duplicate stream, instead it MUST generate new RTCP reports for the
duplicate stream. The sender MUST use the same RTCP CNAME in the
RTCP reports it sends for the main and redundant streams, so that the
receiver can synchronize them.

Both the main and redundant RTP streams, and their corresponding RTCP
reports, will be received. If RTCP is used, receivers MUST generate
RTCP reports for both main and redundant streams in the usual way,
treating them as entirely separate media streams.

4.2. Signaling Considerations

Signaling is needed to allow the receiver to determine that an RTP
stream is a redundant copy of another, rather than a separate stream
that needs to be rendered in parallel. There are two parts to this:
an SDP extension is needed in the offer/answer exchange to negotiate
support for temporal redundancy; and signalling is needed to indicate
which stream is the duplicate (the latter can be done in-band using
an RTCP extension, or out-of-band by signalling the SSRCs used by the
duplicate streams in SDP).

We require out-of-band signalling for both features. The required

SDP attribute to signal duplication in the SDP offer/answer exchange
('duplication-delay') is defined in
[I-D.begen-mmusic-temporal-interleaving]. The required SDP grouping
semantics are defined in [I-D.begen-mmusic-redundancy-grouping].

In the following SDP example, a video stream is duplicated, and the
main and redundant streams are transmitted in two separate SSRCs
(1000 and 1010):

```
     v=0
     o=ali 1122334455 1122334466 IN IP4 dup.example.com
     s=Delayed Duplication
     t=0 0
     m=video 30000 RTP/AVP 100
     c=IN IP4 233.252.0.1/127
     a=source-filter:incl IN IP4 233.252.0.1 198.51.100.1
     a=rtpmap:100 MP2T/90000
     a=ssrc:1000 cname:ch1@example.com
     a=ssrc:1010 cname:ch1@example.com
     a=ssrc-group:DUP 1000 1010
     a=duplication-delay:100
     a=mid:Group1
```

It is RECOMMENDED that the SSRC listed first in the "a=ssrc-group:"
line is sent first, with the other RTP SSRC being the time-delayed
duplicate. This is not critical, however, and receivers should size
their playout buffers based on the "a=duplication-delay:" attribute,
and play the stream that arrives first in preference, with the other
stream acting as a repair stream, irrespective of the order in which
they are signalled.

5. **Use of RTP and RTCP with Spatial Redundancy**

When using spatial redundancy, the redundant RTP stream is sent on
using a different source and/or destination address/port pair. This
will be a separate RTP session to the session conveying the main RTP
stream.

The SSRCs used for the main and redundant streams MUST be chosen
randomly, following the rules in Section 8 of [RFC3550].
Accordingly, they will almost certainly not match each other. The
sender MUST, however, use the same RTCP CNAME for both the main and
redundant streams, and MUST include an "a=ssrc:... srcname:..."
attribute to correlate the flows. An "a=group:DUP" attribute is used
to indicate duplication.

5.1. RTCP Considerations

If RTCP is being sent for the main RTP stream, then the sender MUST
also generate RTCP for the redundant RTP stream. The RTCP for the
redundant RTP stream is generated exactly as-if the redundant RTP
stream were a regular media stream; the sender MUST NOT duplicate the
RTCP packets sent for the main RTP stream. The sender MUST use the
same RTCP CNAME in the RTCP reports it sends for the main and
redundant streams, so that the receiver can synchronize them.

The main and redundant streams are conceptually synchronised using
the standard RTCP SR-based mechanism, deriving a mapping between
their timelines. The RTP timestamps and sequence numbers SHOULD be
identical in the main and redundant streams, however, making the
mapping trivial in most cases.

Both main and redundant streams, and their corresponding RTCP, will
be received. If RTCP is used, receivers MUST generate RTCP reports
for both main and redundant streams in the usual way, treating them
as entirely separate media streams.

5.2. Signaling Considerations

The required SDP grouping semantics have been defined in
[I-D.begen-mmusic-redundancy-grouping]. In the following example,
the redundant streams have different IP destination addresses. The
example shows the same UDP port number and IP source addresses, but
either or both could have been different for the two streams.

```
v=0
o=ali 1122334455 1122334466 IN IP4 dup.example.com
s=DUP Grouping Semantics
t=0 0
a=group:DUP S1a S1b
m=video 30000 RTP/AVP 100
c=IN IP4 233.252.0.1/127
a=source-filter:incl IN IP4 233.252.0.1 198.51.100.1
a=rtpmap:100 MP2T/90000
a=ssrc:1000 cname:ch1@example.com
a=ssrc:1000 srcname:45:a8:f4:19:b4:c3
a=mid:S1a
m=video 30000 RTP/AVP 101
c=IN IP4 233.252.0.2/127
a=source-filter:incl IN IP4 233.252.0.2 198.51.100.1
a=rtpmap:101 MP2T/90000
a=ssrc:1010 cname:ch1@example.com
a=ssrc:1010 srcname:45:a8:f4:19:b4:c3
a=mid:S1b
```

6. Use of RTP and RTCP with Temporal and Spatial Redundancy

   This uses the same RTP/RTCP mechanisms, plus a combination of both
   sets of signaling.

7. Security Considerations

   The security considerations of [RFC3550],
   [I-D.begen-mmusic-temporal-interleaving], and
   [I-D.begen-mmusic-redundancy-grouping] apply.

   If stream de-duplication is done by an in-network middlebox, rather
   than by an end system, that middlebox can work if Secure RTP (SRTP)
   encryption is used [RFC3711], since the RTP headers are in the clear.
   Doing so would break the authentication when the SSRC is rewritten,
   unless the de-duplication middlebox were trusted to re-authenticate
   the packets. This would require additional signalling which is not
   specified here, since de-duplication in the receiver end system is
   expected to be the more common use case.

8. IANA Considerations

   No IANA actions are required.

9. Acknowledgments

Thanks to Magnus Westerlund for his suggestions.

10. References

10.1. Normative References

[RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V.
           Jacobson, "RTP: A Transport Protocol for Real-Time
           Applications", STD 64, RFC 3550, July 2003.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
           Requirement Levels", BCP 14, RFC 2119, March 1997.

[I-D.begen-mmusic-temporal-interleaving]
           Begen, A., Cai, Y., and H. Ou, "Delayed Duplication
           Attribute in the Session Description Protocol",
           draft-begen-mmusic-temporal-interleaving-03 (work in
           progress), October 2011.

[I-D.begen-mmusic-redundancy-grouping]
           Begen, A., Cai, Y., and H. Ou, "Duplication Grouping
           Semantics in the Session Description Protocol",
           draft-begen-mmusic-redundancy-grouping-02 (work in
           progress), October 2011.

[RFC3711] Baugher, M., McGrew, D., Naslund, M., Carrara, E., and K.
           Norrman, "The Secure Real-time Transport Protocol (SRTP)",
           RFC 3711, March 2004.

10.2. Informative References

[RFC2354] Perkins, C. and O. Hodson, "Options for Repair of
           Streaming Media", RFC 2354, June 1998.

[RFC4588] Rey, J., Leon, D., Miyazaki, A., Varsa, V., and R.
           Hakenberg, "RTP Retransmission Payload Format", RFC 4588,
           July 2006.

[RFC5109] Li, A., "RTP Payload Format for Generic Forward Error
           Correction", RFC 5109, December 2007.

[IC2011]   Evans, J., Begen, A., Greengrass, J., and C. Filsfils,
           "Toward Lossless Video Transport (to appear in IEEE
           Internet Computing)", November 2011.

Authors' Addresses

    Ali Begen
    Cisco
    181 Bay Street
    Toronto, ON M5J 2T3
    CANADA

    Email: abegen@cisco.com

    Colin Perkins
    University of Glasgow
    School of Computing Science
    Glasgow,   G12 8QQ
    UK

    Email: csp@csperkins.org