

AVTEXT
Internet-Draft
Intended status: Standards Track
Expires: March 24, 2014

A. Begen
Cisco
C. Perkins
University of Glasgow
September 20, 2013

Duplicating RTP Streams
draft-ietf-avtext-rtp-duplication-03

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 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 March 24, 2014.

Copyright Notice

Copyright (c) 2013 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 (<http://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	2
2. Terminology and Requirements Notation	3
3. Dual Streaming Use Cases	3
3.1. Temporal Redundancy	3
3.2. Spatial Redundancy	4
3.3. Dual Streaming over a Single Path or Multiple Paths	4
4. Use of RTP and RTCP with Temporal Redundancy	5
4.1. RTCP Considerations	6
4.2. Signaling Considerations	6
5. Use of RTP and RTCP with Spatial Redundancy	7
5.1. RTCP Considerations	7
5.2. Signaling Considerations	8
6. Use of RTP and RTCP with Temporal and Spatial Redundancy	8
7. Security Considerations	8
8. IANA Considerations	9
9. Acknowledgments	9
10. References	9
10.1. Normative References	9
10.2. Informative References	9
Authors' Addresses	10

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 processing 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 [RFC6363].

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 (the original plus its duplicate) 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 practice, 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 sender 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 stream and its duplicate stream can be sent to the same IP destination address and UDP destination port with a certain delay between them [I-D.ietf-mmusic-delayed-duplication]. 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 commonly called stream merging or de-duplication.

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 its duplicate stream, and with a delay, so that if there is a packet loss, such a delay would allow the same packet from the duplicate

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.ietf-mmusic-delayed-duplication] [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 (This is often made sure through careful network design). 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 duplicate RTP streams SHOULD be sent using the sample 5-tuple of transport protocol, source and destination IP addresses, and source and destination transport ports. Due to the possible presence of network address and port translation (NAPT) devices, load balancers, or other middleboxes, use of anything other than an identical 5-tuple might also cause spatial redundancy (which might introduce an additional delay due to the delta between the path delays), and so is NOT RECOMMENDED unless the path is known to be free of such middleboxes.

Since the main and duplicate RTP streams follow an identical path, they are part of the same RTP session. Accordingly, the sender MUST choose a different SSRC for the duplicate 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 duplicate RTP stream. The RTCP for the duplicate RTP stream is generated exactly as-if the duplicate 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 both streams, so that the receiver can synchronize them.

The main and duplicate streams are conceptually synchronized using the standard RTCP Sender Report-based mechanism, deriving a mapping between their timelines. However, the RTP timestamps and sequence numbers MUST be identical in the main and duplicate streams, making the mapping quite trivial.

Both the main and duplicate RTP streams, and their corresponding RTCP reports, will be received. If RTCP is used, receivers MUST generate RTCP reports for both the main and duplicate 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 duplicate 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 signaling is needed to indicate which stream is the duplicate (the latter can be done in-band using an RTCP extension, or out-of-band in the SDP description).

We require out-of-band signaling for both features. The required SDP attribute to signal duplication in the SDP offer/answer exchange ('duplication-delay') is defined in [I-D.ietf-mmusic-delayed-duplication]. The required SDP grouping semantics are defined in [I-D.ietf-mmusic-duplication-grouping].

In the following SDP example, a video stream is duplicated, and the main and duplicate 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:ch1a@example.com
a=ssrc:1010 cname:ch1a@example.com
a=ssrc-group:DUP 1000 1010
a=duplication-delay:50
a=mid:Ch1
```

As specified in Section 3.2 of [I-D.ietf-mmusic-duplication-grouping], it is advisable that the SSRC listed first in the "a=ssrc-group:" line (i.e., SSRC of 1000) is sent first, with the other SSRC (i.e., SSRC of 1010) being the time-delayed duplicate. This is not critical, however, and a receiving host should size its playout buffer based on the '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 signaled.

5. Use of RTP and RTCP with Spatial Redundancy

When using spatial redundancy, the duplicate RTP stream is sent 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. Thus, the SSRCs used for the main and duplicate 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 duplicate streams. An "a=group:DUP" line or "a=ssrc-group:DUP" line 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 duplicate RTP stream. The RTCP for the duplicate RTP stream is generated exactly as-if the duplicate 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 both streams, so that the receiver can synchronize them.

The main and duplicate streams are conceptually synchronized using the standard RTCP Sender Report-based mechanism, deriving a mapping between their timelines. However, the RTP timestamps and sequence numbers MUST be identical in the main and duplicate streams, making the mapping quite trivial.

Both the main and duplicate RTP streams, and their corresponding RTCP reports, will be received. If RTCP is used, receivers MUST generate RTCP reports for both the main and duplicate 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.ietf-mmusic-duplication-grouping]. In the following example, the redundant streams have different IP destination addresses. The example shows the same UDP port number and IP source address for each stream, 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=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=mid:S1b
```

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

This uses the same RTP/RTCP mechanisms from Sections Section 4 and Section 5, plus a combination of both sets of signaling.

7. Security Considerations

The security considerations of [RFC3550], [I-D.ietf-mmusic-delayed-duplication], and [I-D.ietf-mmusic-duplication-grouping] apply.

Stream de-duplication can be done by an in-network middlebox, rewriting the SSRC as appropriate. If the Secure RTP (SRTP) profile [RFC3711] is used to authenticate RTP packets, such rewriting is not possible without breaking the authentication, unless the de-duplication middlebox is trusted to re-authenticate the packets. This would require additional signaling that is not specified here. The use of the encryption features of SRTP does not affect stream de-duplication middleboxes, since the RTP headers are sent in the clear.

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.ietf-mmusic-delayed-duplication]
Begen, A., Cai, Y., and H. Ou, "Delayed Duplication Attribute in the Session Description Protocol", draft-ietf-mmusic-delayed-duplication-02 (work in progress), May 2013.
- [I-D.ietf-mmusic-duplication-grouping]
Begen, A., Cai, Y., and H. Ou, "Duplication Grouping Semantics in the Session Description Protocol", draft-ietf-mmusic-duplication-grouping-03 (work in progress), July 2013.
- [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.
- [RFC6363] Watson, M., Begen, A., and V. Roca, "Forward Error Correction (FEC) Framework", RFC 6363, October 2011.
- [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