

Internet Engineering Task Force (IETF)
Request for Comments: 7198
Category: Standards Track
ISSN: 2070-1721

A. Begen
Cisco
C. Perkins
University of Glasgow
April 2014

Duplicating RTP Streams

Abstract

Packet loss is undesirable for real-time multimedia sessions but can occur due to a variety of reasons including unplanned network outages. In unicast transmissions, recovering from such an outage can be difficult depending on the outage duration, due to the potentially large number of missing packets. In multicast transmissions, recovery is even more challenging as many receivers could be impacted by the outage. For this challenge, one solution that does not incur unbounded delay is to duplicate the packets and send them in separate redundant streams, provided that the underlying network satisfies certain requirements. 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 is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7198>.

Copyright Notice

Copyright (c) 2014 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	3
2. Terminology and Requirements Notation	4
3. Use Cases for Dual Streaming	4
3.1. Temporal Redundancy	4
3.2. Spatial Redundancy	5
3.3. Dual Streaming over a Single Path or Multiple Paths	5
3.4. Requirements	6
4. Use of RTP and RTCP with Temporal Redundancy	7
4.1. RTCP Considerations	7
4.2. Signaling Considerations	7
5. Use of RTP and RTCP with Spatial Redundancy	8
5.1. RTCP Considerations	9
5.2. Signaling Considerations	9
6. Use of RTP and RTCP with Temporal and Spatial Redundancy	10
7. Congestion Control Considerations	10
8. Security Considerations	11
9. Acknowledgments	11
10. References	12
10.1. Normative References	12
10.2. Informative References	12

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. This loss might be due to congestion; it might also be a result of an unplanned outage caused by a flapping link, a 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 is left up to the RTP layer and the applications to detect, and recover from, packet loss.

In a carefully managed network, congestion should not normally happen; however, network outages can still happen due to the reasons listed above. In such a managed network, one technique to recover from packet loss without incurring unbounded delay is to duplicate the packets and send them in separate redundant streams. As described later in this document, the probability that two copies of the same packet are lost in cases of non-congestive packet loss is quite small.

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 discusses the most common use cases and explains how duplication can be achieved for RTP streams in such use cases to address the immediate market needs. In the future, if there will be a different use case that is not covered by this document, a new specification that explains how RTP duplication should be done in such a scenario may be needed.

Stream duplication offers a simple way to protect media flows from packet loss. It has a comparatively high overhead in terms of bandwidth, since everything is sent twice, but with a low overhead in terms of processing. It is also very predictable in its overheads. Alternative approaches, for example, retransmission-based recovery [RFC4588] or Forward Error Correction [RFC6363], may be suitable in some other cases.

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. Use Cases for Dual Streaming

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 and the little additional protection it provides.

3.1. Temporal Redundancy

From a routing perspective, two streams are considered identical if the following two IP header fields are the same (in addition to the transport ports), 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 (SSRCs) 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 [RFC7197]. 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, the destination host could also have multiple IP addresses for an RTP source to send the redundant streams to.

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 flow labels will be the same. 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 are often small. This is what is also known as "time-shifted redundancy", "temporal redundancy" or simply "delayed duplication" [RFC7197] [IC2011]. This approach can be used with both types of dual streaming, described in Sections 3.1 and 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 ensured 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).

Temporal and spatial redundancy deal with different patterns of packet loss. The former helps with transient loss (within the duplication window), while the latter helps with longer-term packet loss that affects only one of the two redundant paths.

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].

3.4. Requirements

One of the following conditions is currently REQUIRED to hold in applications using this specification:

- o The original and duplicate RTP streams are carried (with their own SSRCs) in the same "m" line. (There could be other RTP streams listed in the same "m" line.)
- o The original and duplicate RTP streams are carried in separate "m" lines, and there is no other RTP stream listed in either "m" line.

When the original and duplicate RTP streams are carried in separate "m" lines in a Session Description Protocol (SDP) description and if the SDP description has one or more other RTP streams listed in either "m" line, duplication grouping is not trivial and further signaling will be needed; this is left for future standardization.

4. Use of RTP and RTCP with Temporal Redundancy

To achieve temporal redundancy, the main and duplicate RTP streams SHOULD be sent using the same 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 and flow label might also cause spatial redundancy (which might introduce an additional delay due to the delta between the path delays), and so it 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 Section 8 of [RFC3550].

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 mechanism based on RTCP Sender Reports, 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.)

Out-of-band signaling is needed for both features. The SDP attribute to signal duplication in the SDP offer/answer exchange ('duplication-delay') is defined in [RFC7197]. The required SDP grouping semantics are defined in [RFC7104].

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
```

Section 3.2 of [RFC7104] states that 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

Assuming the network is structured appropriately, 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 from 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 mechanism based on RTCP Sender Reports, 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 [RFC7104]. 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 4 and 5, plus a combination of signaling provided in each of these sections.

7. Congestion Control Considerations

Duplicating RTP streams has several considerations in the context of congestion control. First of all, RTP duplication **MUST NOT** be used in cases where the primary cause of packet loss is congestion since duplication can make congestion only worse. Furthermore, RTP duplication **SHOULD NOT** be used where there is a risk of congestion upon duplicating an RTP stream. Duplication is **RECOMMENDED** only to be used for protection against network outages due to a temporary link or network element failure and where it is known (e.g., through explicit operator configuration) that there is sufficient network capacity to carry the duplicated traffic. The capacity requirement constrains the use of duplication to managed networks and makes it unsuitable for use on unmanaged public networks.

It is essential that the nodes responsible for the duplication and de-duplication are aware of the original stream's requirements and the available capacity inside the network. If there is an adaptation capability for the original stream, these nodes have to assume the same adaptation capability for the duplicated stream, too. For example, if the source doubles the bitrate for the original stream, the bitrate of the duplicate stream will also be doubled.

Depending on where de-duplication takes place, there could be different scenarios. When the duplication and de-duplication take place inside the network before the ultimate endpoints that will consume the RTP media, the whole process is transparent to these endpoints. Thus, these endpoints will apply any congestion control, if applicable, on the de-duplicated RTP stream. This output stream will have fewer losses than either the original or duplicated stream will have, and the endpoint will make congestion control decisions accordingly. However, if de-duplication takes place at the ultimate endpoint, this endpoint **MUST** consider the aggregate of the original and duplicated RTP stream in any congestion control it wants to apply. The endpoint will observe the losses in each stream separately, and this information can be used to fine-tune the duplication process. For example, the duplication interval can be adjusted based on the duration of a common packet loss in both streams. In these scenarios, the RTP Monitoring Framework [RFC6792] can be used to monitor the duplicated streams in the same way an ordinary RTP would be monitored.

8. Security Considerations

The security considerations of [RFC3550], [RFC7104], [RFC7197], and any RTP profiles and payload formats in use apply.

Duplication can be performed end-to-end, with the media sender generating a duplicate RTP stream, and the receiver(s) performing de-duplication. In such cases, if the original media stream is to be authenticated (e.g., using Secure RTP (SRTP) [RFC3711]), then the duplicate stream also needs to be authenticated, and duplicate packets that fail the authentication check need to be discarded.

Stream duplication and de-duplication can also be performed by in-network middleboxes. Such middleboxes will need to rewrite the RTP SSRC such that the RTP packets in the duplicate stream have a different SSRC to the original stream, and such middleboxes will need to generate and respond to RTCP packets corresponding to the duplicate stream. This sort of in-network duplication service has the potential to act as an amplifier for denial-of-service attacks if the attacker can cause attack traffic to be duplicated. To prevent this, middleboxes providing the duplication service need to authenticate the traffic to be duplicated as being from a legitimate source, for example, using the SRTP profile [RFC3711]. This requires the middlebox to be part of the security context of the media session being duplicated, so it has access to the necessary keying material for authentication. To do this, the middlebox will need to be privy to the session setup signaling. Details of how that is done will depend on the type of signaling used (SIP, Real Time Streaming Protocol (RTSP), WebRTC, etc.), and is not specified here.

Similarly, to prevent packet injection attacks, a de-duplication middlebox needs to authenticate original and duplicate streams, and ought not use non-authenticated packets that are received. Again, this requires the middlebox to be part of the security context and to have access to the appropriate signaling and keying material.

The use of the encryption features of SRTP does not affect stream de-duplication middleboxes, since the RTP headers are sent in the clear.

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.
- [RFC7197] Begen, A., Cai, Y., and H. Ou, "Duplication Delay Attribute in the Session Description Protocol", RFC 7197, April 2014.
- [RFC7104] Begen, A., Cai, Y., and H. Ou, "Duplication Grouping Semantics in the Session Description Protocol", RFC 7104, January 2014.
- [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.
- [RFC6792] Wu, Q., Hunt, G., and P. Arden, "Guidelines for Use of the RTP Monitoring Framework", RFC 6792, November 2012.
- [IC2011] Evans, J., Begen, A., Greengrass, J., and C. Filsfils, "Toward Lossless Video Transport", IEEE Internet Computing, Vol. 15, No. 6, pp. 48-57, November 2011.

Authors' Addresses

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

EEmail: abegen@cisco.com

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

EEmail: csp@csperskins.org
URI: <http://csperskins.org/>