

Internet Engineering Task Force (IETF)
Request for Comments: 7197
Category: Standards Track
ISSN: 2070-1721

A. Begen
Cisco
Y. Cai
Microsoft
H. Ou
Cisco
April 2014

Duplication Delay Attribute in the Session Description Protocol

Abstract

A straightforward approach to provide protection against packet losses due to network outages with a longest duration of T time units is to duplicate the original packets and send each copy separated in time by at least T time units. This approach is commonly referred to as "time-shifted redundancy", "temporal redundancy", or simply "delayed duplication". This document defines an attribute to indicate the presence of temporally redundant media streams and the duplication delay in the Session Description Protocol.

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/rfc7197>.

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	2
2. Requirements Notation	4
3. The 'duplication-delay' Attribute	5
4. SDP Examples	6
5. Security Considerations	7
6. IANA Considerations	8
6.1. Registration of SDP Attributes	9
7. Acknowledgements	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10

1. Introduction

Inside an IP network, packet delivery may be interrupted due to failure of a physical link, interface, or device. To reduce the impact of such interruptions, some networks are built in a resilient manner, allowing for multiple alternative paths between two endpoints. However, if there is no resiliency in the network or the failure happens in a non-resilient part of the network, a temporary outage will occur (i.e., packets will get dropped). The outage will last until network reconvergence takes place (i.e., until connectivity is restored) around the failure. Typically, network reconvergence takes between tens and hundreds of milliseconds, depending on the size and features of the network.

There are a number of network-reconvergence technologies available today, such as IP Fast Convergence, MPLS Traffic Engineering Fast Reroute, and Multicast Only Fast Reroute. These technologies can be augmented by different types of application-layer loss-repair methods such as Forward Error Correction (FEC), retransmission, temporal

redundancy, and spatial redundancy to minimize (and sometimes totally eliminate) the impact of outages. Each combination has its distinct requirements in terms of bandwidth consumption and results in a different network complexity. Thus, a network operator has to carefully consider what combination to deploy for different parts of a network (e.g., core vs. edge). A detailed overview of network-convergence technologies and loss-repair methods is provided in [IC2011].

One of the loss-repair methods is temporal redundancy, also known as delayed duplication. A media sender using this method transmits an original source packet and transmits its duplicate after a certain delay following the original transmission. If a network outage hits the original transmission, the expectation is that the second transmission arrives at the receiver (with a high probability). Alternatively, the second transmission may be hit by an outage and so gets dropped, and the original transmission completes successfully. Also, both transmissions can arrive on the receiver side; in that case, the receiver (or the node that does the duplicate suppression) needs to identify the duplicate packets and discard them appropriately, thereby producing a duplicate-free stream.

Delayed duplication can be used in a variety of multimedia applications where there is sufficient bandwidth for the duplicated traffic and the application can tolerate the introduced delay. However, it must be used with care, since it might easily result in a new series of denial-of-service attacks. Delayed duplication is harmful in cases where the primary cause of packet loss is congestion, rather than a network outage due to a temporary link or network element failure. Duplication should only be used by endpoints that want to protect against network failures; protection against congestion must be achieved through other means, as duplication will only make congestion worse.

One particular use case for delayed duplication is to improve the reliability of real-time video feeds inside a core IP network where bandwidth is plentiful and maximum reliability (preferably zero loss) is desired [IC2011]. Compared to other redundancy approaches such as FEC [RFC6363] and redundant data encoding (e.g., [RFC2198]), delayed duplication is easy to implement, since it does not require any special type of encoding or decoding.

For duplicate suppression, the receiver has to be able to identify the identical packets. This is straightforward for media packets that carry one or more unique identifiers such as the sequence number field in the RTP header [RFC3550]. In non-RTP applications, the receiver can use unique sequence numbers if available or other alternative approaches to compare the incoming packets and discard the duplicate ones.

This specification introduces a new Session Description Protocol (SDP) [RFC4566] attribute for applications/services using the delayed duplication method to indicate the relative delay for each additional duplication. The attribute is used with the duplication grouping semantics defined in [RFC7104].

This specification does not explain how to select the duplication delay that a sender should use; the selection technique depends on the underlying network and the reconvergence technologies used inside such a network. This specification does not explain how the receiver should suppress the duplicate packets and merge the incoming streams to produce a loss-free and duplication-free output stream (a process commonly called "stream merging"), either. An application or a transport service that will use the delayed duplication method must determine its own rules about stream merging.

In practice, more than two redundant streams are unlikely to be used, since the additional delay and increased overhead are not easily justified. However, we define the new attribute in a general way so that it could be used with more than two redundant streams (i.e., multiple duplications), if needed. While the primary focus in this specification is the RTP-based transport, the new attribute is applicable to both RTP and non-RTP streams. Protocol issues and details on duplicating RTP streams are presented in [RFC7198].

2. 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. The 'duplication-delay' Attribute

The following ABNF [RFC5234] syntax formally describes the 'duplication-delay' attribute:

```
delaying-attribute    = "a=duplication-delay:" periods CRLF
periods               = period *( SP period)
period                = 1*DIGIT ; in milliseconds
```

ABNF Syntax for the 'duplication-delay' Attribute

The 'duplication-delay' attribute is defined as both a media-level and session-level attribute. It specifies the relative delay with respect to the previous transmission of each duplication in milliseconds (ms) at the time of transmission. The following rules apply:

- o If used as a media-level attribute, it MUST be used with the 'ssrc-group' attribute and "DUP" grouping semantics as defined in [RFC7104]. When used as a media-level attribute, the relative delay value(s) it specifies SHALL apply to every Synchronization Source (SSRC)-based duplication grouping in the same media description. In other words, one cannot specify different duplication delay values for different duplication groups in the same media description.
- o If used as a session-level attribute, it MUST be used with 'group' attribute and "DUP" grouping semantics as defined in [RFC7104]. When used as a session-level attribute, the relative delay value(s) it specifies SHALL apply to every duplication grouping in the same SDP description. In other words, one cannot specify different duplication delay values for different duplication groups in the same SDP description. If one needs to specify different duplication delay values for different duplication groups, then one MUST use different SDP descriptions for each or MUST use the 'duplication-delay' attribute at the media level. In that case, the 'duplication-delay' attribute MUST NOT be used at the session level.
- o For offer/answer model considerations, refer to [RFC7104].

4. SDP Examples

In the first example below, the multicast stream consists of two RTP streams, each duplicated once, resulting in two sets of two-stream groups. The same duplication delay of 100 ms is applied to each grouping. The first set's streams have SSRCs of 1000 and 1010, and the second set's streams have SSRCs of 1020 and 1030.

```
v=0
o=ali 1122334455 1122334466 IN IP4 dup.example.com
s=Delayed Duplication
t=0 0
m=video 30000 RTP/AVP 100 101
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=rtpmap:101 MP2T/90000
a=ssrc:1020 cname:ch1b@example.com
a=ssrc:1030 cname:ch1b@example.com
a=ssrc-group:DUP 1020 1030
a=duplication-delay:100
a=mid:Ch1
```

Note that in actual use, SSRC values, which are random 32-bit numbers, could be much larger than the ones shown in this example.

In the second example below, the multicast stream is duplicated twice. 50 ms after the original transmission, the first duplicate is transmitted, and 100 ms after that, the second duplicate is transmitted. In other words, the same packet is transmitted three times over a period of 150 ms.

```
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:ch1c@example.com
a=ssrc:1010 cname:ch1c@example.com
a=ssrc:1020 cname:ch1c@example.com
```

```
a=ssrc-group:DUP 1000 1010 1020
a=duplication-delay:50 100
a=mid:Ch1
```

In the third example below, the multicast UDP stream is duplicated with a duplication delay of 50 ms. Redundant streams are sent in separate source-specific multicast (SSM) sessions, so the receiving host has to join both SSM sessions if it wants to receive both streams.

```
v=0
o=ali 1122334455 1122334466 IN IP4 dup.example.com
s=Delayed Duplication
t=0 0
a=group:DUP S1a S1b
a=duplication-delay:50
m=audio 30000 udp mp4
c=IN IP4 233.252.0.1/127
a=source-filter:incl IN IP4 233.252.0.1 198.51.100.1
a=mid:S1a
m=audio 40000 udp mp4
c=IN IP4 233.252.0.2/127
a=source-filter:incl IN IP4 233.252.0.2 198.51.100.1
a=mid:S1b
```

5. Security Considerations

The 'duplication-delay' attribute is not believed to introduce any significant security risk to multimedia applications. A malevolent third party could use this attribute to misguide the receiver(s) about the duplication delays and/or the number of redundant streams. For example, if the malevolent third party increases the value of the duplication delay, the receiver(s) will unnecessarily incur a longer delay, since they will have to wait for the entire period. Or, if the duplication delay is reduced by the malevolent third party, the receiver(s) might not wait long enough for the duplicated transmission and incur unnecessary packet losses. However, these require intercepting and rewriting the packets carrying the SDP description; if an interceptor can do that, many more attacks are also possible.

In order to avoid attacks of this sort, the SDP description needs to be integrity protected and provided with source authentication. This can, for example, be achieved on an end-to-end basis using S/MIME [RFC5652] [RFC5751] when SDP is used in a signaling packet using MIME types (application/sdp). Alternatively, HTTPS [RFC2818] or the authentication method in the Session Announcement Protocol (SAP) [RFC2974] could be used as well.

Another security risk is due to possible software misconfiguration or a software bug where a large number of duplicates could be unwillingly signaled in the 'duplication-delay' attribute. Similarly, an attacker can use this attribute to start a denial-of-service attack by signaling and sending too many duplicated streams. In applications where this attribute is to be used, it is a good practice to put a hard limit on both the number of duplicate streams and the total delay introduced due to duplication, regardless of what the SDP description specifies.

Since this mechanism causes duplication of media packets, if those packets are also cryptographically protected (e.g., encrypted) then such duplication could act as an accelerator if any Million Message [RFC3218] or similar attack such as Lucky 13 [Lucky13] exists against the security mechanism that is in use. Such acceleration could turn an otherwise infeasible attack into one that is practical; however, assuming that the amount of duplication is small and that such weak or broken security mechanisms should really not be used, the overall security impact of the duplication should be minimal. If, however, a bad actor were in control of the SDP but did not have access to the keying material used for media, then such a bad actor could potentially use the SDP to cause the media handling to use a weak or broken mechanism with a lot of duplication, in which case the duplication could be significant. Deployments where the SDP is controlled by an actor who should not have access to the media keying material should therefore be cautious in their use of this duplication mechanism.

If this mechanism were used in conjunction with a source description (SDS) and if the key being used for media protection is derived from a human-memorable or otherwise dictionary-attackable secret, then the duplication done here could allow for a more efficient dictionary attack against the media. The right countermeasure is to use proper keying or, if using an SDS, to ensure that the keys used are not dictionary-attackable.

6. IANA Considerations

The following contact information shall be used for the registration in this document:

Ali Begen
abegen@cisco.com

6.1. Registration of SDP Attributes

This document registers a new attribute name in SDP.

SDP Attribute ("att-field"):

Attribute name:	duplication-delay
Long form:	Duplication delay for temporally redundant streams
Type of name:	att-field
Type of attribute:	Media or session level
Subject to charset:	No
Purpose:	Specifies the relative duplication delay(s) for redundant stream(s)
Reference:	[RFC7197]
Values:	See [RFC7197]

7. Acknowledgements

The authors would like to thank Colin Perkins, Paul Kyzivat, and Stephen Farrell for their suggestions and reviews.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [RFC7104] Begen, A., Cai, Y., and H. Ou, "Duplication Grouping Semantics in the Session Description Protocol", RFC 7104, January 2014.

8.2. Informative References

- [RFC6363] Watson, M., Begen, A., and V. Roca, "Forward Error Correction (FEC) Framework", RFC 6363, October 2011.
- [RFC2198] 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.
- [RFC7198] Begen, A. and C. Perkins, "Duplicating RTP Streams", RFC 7198, April 2014.
- [IC2011] Evans, J., Begen, A., Greengrass, J., and C. Filmsfils, "Toward Lossless Video Transport", IEEE Internet Computing, Vol. 15, No. 6, pp. 48-57, November 2011.
- [RFC5652] Housley, R., "Cryptographic Message Syntax (CMS)", STD 70, RFC 5652, September 2009.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [RFC2818] Rescorla, E., "HTTP Over TLS", RFC 2818, May 2000.
- [RFC2974] Handley, M., Perkins, C., and E. Whelan, "Session Announcement Protocol", RFC 2974, October 2000.
- [RFC3218] Rescorla, E., "Preventing the Million Message Attack on Cryptographic Message Syntax", RFC 3218, January 2002.
- [Lucky13] AlFardan, N. and K. Paterson, "Lucky Thirteen: Breaking the TLS and DTLS Record Protocols", IEEE Symposium on Security and Privacy, May 2013, <<http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6547131&queryText%3DLucky+Thirteen>>.

Authors' Addresses

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

E-Mail: abegen@cisco.com

Yiqun Cai
Microsoft
1065 La Avenida
Mountain View, CA 94043
USA

E-Mail: yiqunc@microsoft.com

Heidi Ou
Cisco
170 W. Tasman Dr.
San Jose, CA 95134
USA

E-Mail: hou@cisco.com

