

MMUSIC
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2009

A. Begen
Cisco Systems
October 27, 2008

**Forward Error Correction Grouping Semantics in Session Description
Protocol
draft-begen-mmusic-rfc4756bis-00**

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 April 30, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

The Session Description Protocol (SDP) supports grouping media lines. SDP also has semantics defined for grouping the associated source and Forward Error Correction (FEC)-based repair flows. However, the semantics that were defined in RFC 4756 generally fail to provide the specific grouping relationships between the source and repair flows when there are more than one source and/or repair flows in the same group. Furthermore, the existing semantics also do not support

additive repair flows and prioritization among the repair flows protecting one or more source flows. This document addresses these issues by introducing new FEC grouping semantics.

Table of Contents

- 1. Introduction 3
- 2. Requirements Notation 5
- 3. Requirements and Issues with RFC 4756 5
 - 3.1. Source and Repair Flow Association 5
 - 3.2. Support for Additivity 6
 - 3.3. Support for Prioritization 6
- 4. FEC Grouping 7
 - 4.1. New Grouping Semantics 7
 - 4.2. Offer / Answer Consideration 7
 - 4.3. Example of FEC Grouping 8
- 5. Security Considerations 9
- 6. IANA Considerations 9
- 7. Acknowledgments 9
- 8. References 9
 - 8.1. Normative References 9
 - 8.2. Informative References 10
- Author's Address 10
- Intellectual Property and Copyright Statements 11

1. Introduction

Any application that needs a reliable transmission over an unreliable packet network has to cope with packet losses. Forward Error Correction (FEC) is an effective approach that provides reliable transmission particularly in multicast and broadcast applications where the feedback from the receiver(s) is potentially limited.

In a nutshell, FEC groups source packets into blocks and applies protection to generate a desired number of repair packets. These repair packets may be sent on demand or independently of any receiver feedback. The choice depends on the FEC scheme or the Content Delivery Protocol used by the application, the packet loss characteristics of the underlying network, the transport scheme (e.g., unicast, multicast and broadcast), and the application. At the receiver side, lost packets can be recovered by erasure decoding provided that a sufficient number of source and repair packets have been received.

For example, one of the most basic FEC schemes is the parity codes, where an XOR operation is applied to a group of packets (i.e., source block) to generate a single repair packet. At the receiver side, this scheme achieves a full recovery if only one packet is lost within the source block and the repair packet is received. There are various other ways of generating repair packets, possibly with different loss-recovery capabilities.

The FEC Framework [[I-D.ietf-fecframe-framework](#)] outlines a general framework for using FEC codes in multimedia applications that stream audio, video or other types of multimedia content. The FEC Framework specification states that source and repair packets MUST be carried in different streams, which are referred to as the source and repair flows, respectively. At the receiver side, the receivers should know which flows are the source flows and which flows are the repair flows. The receivers should also know the exact association of the source and repair flows so that they can use the correct data to repair the original content in case there is a packet loss. Currently, SDP [[RFC4566](#)] uses [[RFC3388](#)] and [[RFC4756](#)] for this purpose.

In order to provide applications more flexibility, the FEC Framework [[I-D.ietf-fecframe-framework](#)] allows a source flow to be protected by multiple FEC schemes, each of which requires an instance of the FEC Framework. Thus, multiple instances of the FEC Framework MAY exist at the sender and the receiver(s). Furthermore, within a single FEC Framework instance, multiple source flows MAY be grouped and protected by one or more repair flows.

Begen

Expires April 30, 2009

[Page 3]

A basic example scenario is shown in Figure 1. Here, source flow S1 is protected by repair flow R1. Also, source flows S1 and S2 are grouped and protected together by repair flow R2.




Figure 1: Example scenario with two FEC Framework instances where R1 protects S1, and R2 protects the group of S1 and S2

Grouping source flows before applying FEC protection may allow us to achieve a better coding performance. As a typical scenario, suppose that source flows S1 and S2 in Figure 1 correspond to the base and enhancement layers in a layered video content, respectively. Repair flow R2 protects the combination of the base and enhancement layer for the receivers who receive both layers, and repair flow R1 protects the base layer only, for the receivers who want the base layer only, or who receive both layers but prefer FEC protection for the base layer only due to a bandwidth and/or processing-power limitation.

Using multiple FEC Framework instances for a single source flow provides flexibility to the receivers. An example scenario is sketched in Figure 2. Different instances may offer repair flows that are generated by different FEC schemes, and receivers choose receiving the appropriate repair flow(s) that they can support and decode. Alternatively, different instances (whether they use the same FEC scheme or not) may use larger and smaller source block sizes, which accommodate the receivers that have looser and tighter latency requirements, respectively. In addition, different instances may also provide FEC protection at different redundancy levels. This is particularly useful in multicast scenarios where different receivers might experience different packet loss rates and each receiver can choose the repair flow that is tailored to its needs.




Figure 2: Example scenario with two FEC Framework instances, each with a single repair flow protecting the same source flow S3

To summarize, the FEC Framework supports the following:

1. A source flow MAY be protected by multiple different FEC schemes.
2. An FEC scheme MAY generate multiple repair flows.
3. Source flows MAY be grouped prior to FEC protection. That is, one or more repair flows MAY protect a group of source flows.

To fully benefit from the flexibility provided by the FEC Framework, the grouping semantics for FEC MUST support these features.

2. Requirements Notation

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

3. Requirements and Issues with RFC 4756

3.1. Source and Repair Flow Association

Currently, the 'group' attribute and the FEC grouping semantics defined in [RFC3388] and [RFC4756], respectively, are used to associate source and repair flows together.

The 'group' attribute is used to group multiple repair flows with one or more source flows. However, [RFC3388] prohibits an "m" line identified by its 'mid' attribute from appearing in more than one "a=group" line using the same semantics. This limitation prevents us from indicating specific associations between the source and repair flows by using an "a=group:FEC" line per FEC Framework instance. For example, for the scenario sketched in Figure 1, [RFC3388] mandates us to write

a=group:FEC S1 S2 R1 R2

Clearly, this "a=group:FEC" line does not say anything specific about which repair flows are protecting which source flows.

A new work ([I-D.ietf-mmusic-rfc3388bis]) is currently in progress in the MMUSIC WG to remove this limitation in [RFC3388]. However, [RFC4756] also needs to be updated according to the FEC Framework requirements.

3.2. Support for Additivity

The FEC Framework also supports additive repair flows. Additivity among the repair flows means that multiple repair flows may be decoded jointly to improve the recovery chances of the missing packets in a single or the same set of source flows. Additive repair flows can be generated by the same FEC scheme or different FEC schemes.

For example, in Figure 4, repair flows R5 and R6 may be additive within the FEC Framework instance #1. Alternatively, all three repair flows R5, R6 and R7 could be additive, too.




Figure 4: Example scenario with two FEC Framework instances, where two repair flows in the first instance and a single repair flow in the second instance protect the same source flow S4

3.3. Support for Prioritization

When there are multiple repair flows protecting a single or the same set of source flows, the sender MAY assign different levels of priority to each repair flow. The sender uses prioritization to let the receivers know in which order the receivers MUST receive and decode the repair flows. Note that the repair flows that are assigned a priority level may be additive or non-additive.

A possible scenario where prioritization can be used is that the sender generates two repair flows and sends them in two separate multicast groups. It might be highly desirable to keep all the receivers in the same multicast group to save from the network traffic. The sender can achieve this by assigning a higher priority

to one of the repair flows. Another scenario is where one of the repair flows is more expensive in terms of complexity, cost, etc. and the sender wants every receiver to use the less expensive repair flow first. This is again possible by using prioritization. For example, in Figure 4, the sender MAY assign a higher priority to repair flow R5 and a lower priority to repair flow R6.

Currently, there is no SDP semantics for indicating the priority levels of the individual flows that are grouped in an "a=group" line. Instead, within the FEC Framework, the priority of a repair flow is described with the optional parameter 'priority' in the "a=fec-repair-flow" line [[I-D.ietf-fecframe-sdp-elements](#)], which is used to describe the individual repair flows. However, there might be other benefits of defining more general semantics for flow prioritization in SDP that can be used by other applications and frameworks besides the FEC Framework.

4. FEC Grouping

4.1. New Grouping Semantics

Each "a=group" line is used to indicate an association relationship between the source and repair flows. The flows included in one "a=group" line are called an "FEC Group."

By extending [[I-D.ietf-mmusic-rfc3388bis](#)] we define "FEC-XR" as the new grouping semantics that can support the features of the FEC Framework.

4.2. Offer / Answer Consideration

The backward compatibility in offer / answer is generally handled as specified in [[RFC3388](#)].

Depending on the implementation, a node that does not understand FEC grouping (either does not understand line grouping at all, or just does not understand the FEC semantics) SHOULD respond to an offer containing FEC grouping either (1) with an answer that ignores the grouping attribute or (2) with a refusal to the request (e.g., 488 Not acceptable here or 606 Not acceptable in SIP).

In the first case, the original sender of the offer MUST establish the connection without FEC. In the second case, if the sender of the offer still wishes to establish the session, it SHOULD re-try the request with an offer without FEC.

Editor's note: This section may need to be updated according to

[I-D.ietf-mmusic-rfc3388bis].

4.3. Example of FEC Grouping

For the scenario sketched in Figure 1, we can write the following SDP:

```
v=0
o=ali 1122334455 1122334466 IN IP4 fec.example.com
s=New FEC Grouping Semantics
t=0 0
a=group:FEC-XR S1 R1
a=group:FEX-XR S1 S2 R2
m=video 30000 RTP/AVP 100
c=IN IP4 224.1.1.1/127
a=rtpmap:100 MP2T/90000
a=mid:S1
m=video 30000 RTP/AVP 101
c=IN IP4 224.1.1.2/127
a=rtpmap:101 MP2T/90000
a=mid:S2
m=application 30000 RTP/AVP 110
c=IN IP4 224.1.2.1/127
a=rtpmap:110 1d-interleaved-parityfec/90000
a=fmtp:110 L:5; D:10; repair-window: 200000
a=mid:R1
m=application 30000 RTP/AVP 111
c=IN IP4 224.1.2.2/127
a=rtpmap:111 1d-interleaved-parityfec/90000
a=fmtp:111 L:10; D:10; repair-window: 200000
a=mid:R2
```

For the additivity and prioritization issues, let us consider the scenario sketched in Figure 4. Suppose that repair flows R5 and R6 are additive but repair flow R7 is not additive with any of the other repair flows. In this case, we can write

```
a=group:FEC-XR S4 R5 R6
a=group:FEC-XR S4 R7
```

If none of the repair flows were additive, we could have written

```
a=group:FEC-XR S4 R5
a=group:FEC-XR S4 R6
a=group:FEC-XR S4 R7
```

If the sender wants to assign priority levels to the repair flows, the priority order may be indicated with the order of the 'mid'

values of the repair flows. However, this approach is potentially prone to errors and SHOULD be evaluated carefully whether it covers the corner cases or not. One example is the case when priorities are assigned to a mix of cross-instance additive and non-additive repair flows.

5. Security Considerations

There is a weak threat for the receiver that the FEC grouping can be modified to indicate FEC relationships that do not exist. Such attacks may result in failure of FEC to protect, and/or mishandling of other media payload streams. It is recommended that the receiver SHOULD do integrity check on SDP and follow the security considerations of SDP [RFC4566] to only trust SDP from trusted sources.

6. IANA Considerations

This document registers the following semantics with IANA in Semantics for the "group" SDP Attribute under SDP Parameters.

Semantics	Token	Reference
-----	-----	-----
Forward Error Correction XR	FEC-XR	This document

7. Acknowledgments

Some parts of this document is based on [RFC4756]. Thus, the author would like to thank those who contributed to [RFC4756].

8. References

8.1. Normative References

[I-D.ietf-fecframe-framework]
 Watson, M., "Forward Error Correction (FEC) Framework",
draft-ietf-fecframe-framework-03 (work in progress),
 October 2008.

[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.

[I-D.ietf-mmusic-rfc3388bis]

Camarillo, G., "The SDP (Session Description Protocol) Grouping Framework", [draft-ietf-mmusic-rfc3388bis-01](#) (work in progress), July 2008.

8.2. Informative References

[RFC4756] Li, A., "Forward Error Correction Grouping Semantics in Session Description Protocol", [RFC 4756](#), November 2006.

[RFC3388] Camarillo, G., Eriksson, G., Holler, J., and H. Schulzrinne, "Grouping of Media Lines in the Session Description Protocol (SDP)", [RFC 3388](#), December 2002.

[I-D.ietf-fecframe-sdp-elements]

Begen, A., "SDP Elements for FEC Framework", [draft-ietf-fecframe-sdp-elements-01](#) (work in progress), July 2008.

Author's Address

Ali Begen
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
USA

Email: abegen@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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.

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, THE IETF TRUST 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.

Intellectual Property

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.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

