

FEC Framework Working Group	M. Watson	
Internet-Draft	Digital Fountain	
Intended status: Standards Track	July 12, 2008	
Expires: January 13, 2009		

[TOC](#)

Forward Error Correction (FEC) Framework draft-ietf-fecframe-framework-02

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 January 13, 2009.

Abstract

This document describes for a framework for using forward error correction (FEC) codes with applications in public and private IP networks to provide protection against packet loss. The framework supports applying Forward Error Correction to arbitrary packet flows over unreliable transport and is primarily intended for real-time, or streaming, media. This framework can be used to define Content Delivery Protocols that provide Forward Error Correction for streaming media delivery or other packet flows. Content Delivery Protocols defined using this framework can support any FEC Scheme (and associated FEC codes) which is compliant with various requirements defined in this document. Thus, Content Delivery Protocols can be defined which are not specific to a particular FEC Scheme and FEC Schemes can be defined which are not specific to a particular Content Delivery Protocol.

Table of Contents

1.	Introduction
2.	Definitions/Abbreviations
3.	Requirements notation
4.	Architecture Overview
5.	Procedural overview
5.1.	General
5.2.	Sender Operation
5.3.	Receiver Operation
6.	Protocol Specification
6.1.	General
6.2.	Structure of the source block
6.3.	Packet format for FEC Source packets
6.4.	Packet Format for FEC Repair packets
6.4.1.	Packet Format for FEC Repair packets over RTP
	(informative)
6.5.	FEC Framework Configuration Information
6.6.	FEC Scheme requirements
7.	Transport Protocols
8.	Congestion Control
8.1.	Normative requirements
9.	Security Considerations
10.	IANA Considerations
11.	Acknowledgments
12.	References
§	Author's Address
§	Intellectual Property and Copyright Statements

1. Introduction

[TOC](#)

Many applications have a requirement to transport a continuous stream of packetised data from a source (sender) to one or more destinations (receivers) over networks which do not provide guaranteed packet delivery. Primary examples are real-time, or streaming, media applications such as broadcast, multicast or on-demand audio, video or multimedia.

Forward Error Correction is a well-known technique for improving reliability of packet transmission over networks which do not provide guaranteed packet delivery, especially in multicast and broadcast applications. The FEC Building Block defined in [\[RFC5052\] \(Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction \(FEC\) Building Block," August 2007.\)](#) provides a framework for definition of Content Delivery Protocols (CDPs) for object delivery (including, primarily, file delivery) which make use of separately defined FEC Schemes. Any

CDP defined according to the requirements of the FEC Building Block can then easily be used with any FEC Scheme which is also defined according to the requirements of the FEC Building Block. (Note that the term "Forward Erasure Correction" is sometimes used, 'erasures' being a type of error in which data is lost and this loss can be detected, rather than being received in corrupted form - the focus of this document is strictly on erasures, however the term Forward Error Correction is more widely used).

This document defines a framework for the definition of CDPs which provide for FEC protection of arbitrary packet flows over unreliable transports such as UDP. As such, this document complements the FEC Building Block of [\[RFC5052\] \(Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction \(FEC\) Building Block," August 2007.\)](#), by providing for the case of arbitrary packet flows over unreliable transport, the same kind of framework as that document provides for object delivery. This document does not define a complete Content Delivery Protocol, but rather defines only those aspects that are expected to be common to all Content Delivery Protocols based on this framework.

This framework does not define how the flows to be protected are determined, nor how the details of the protected flows and the FEC streams which protect them are communicated from sender to receiver. It is expected that any complete Content Delivery Protocol specification which makes use of this framework will address these signalling requirements. However, this document does specify the information which is required by the FEC Framework at the sender and receiver - for example details of the flows to be FEC protected, the flow(s) that will carry the FEC protection data and an opaque container for FEC-Scheme-specific information.

FEC Schemes designed for use with this framework must fulfil a number of requirements defined in this document. Note that these requirements are different from those defined in [\[RFC5052\] \(Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction \(FEC\) Building Block," August 2007.\)](#) for FEC Schemes for object delivery. However there is a great deal of commonality and FEC Schemes defined for object delivery may be easily adapted for use with the framework defined here.

2. Definitions/Abbreviations

[TOC](#)

'FEC' Forward Error Correction.

'AL-FEC' Application Layer Forward Error Correction

'FEC Framework' A protocol framework for definition of Content Delivery Protocols using FEC, such as the framework defined in this document.

'Source data flow'

The packet flow or flows to which FEC protection is to be applied.

'Repair data flow' The packet flow or flows carrying forward error correction data

'Source protocol' A protocol used for the source data flow being protected - e.g. RTP.

'Transport protocol' The protocol used for transport of the source data flow being protected - e.g. UDP, DCCP.

'Transport payload' Data used as the payload for the transport layer (e.g. UDP or DCCP packet payload)

'Application protocol' Control protocols used to establish and control the source data flow being protected - e.g. RTSP.

'FEC Code' An algorithm for encoding data such that the encoded data flow is resilient to data loss or corruption.

'FEC Scheme' A specification which defines the additional protocol aspects required to use a particular FEC code with the FEC Framework, or (in the context of RMT), with the RMT FEC Building Block.

'Source Block' the group of source data packets which are to be FEC protected as a single block

'Protection amount' The relative increase in data sent due to the use of FEC.

FEC Framework Configuration Information: Information which controls the operation of the FEC Framework.

FEC Payload ID: Information which identifies the contents of a packet with respect to the FEC Scheme.

Source FEC Payload ID: An FEC Payload ID specifically for use with source packets.

Repair FEC Payload ID: An FEC Payload ID specifically for use with repair packets.

Content Delivery Protocol (CDP): A complete application protocol specification which, through the use of the framework defined in this document, is able to make use of FEC Schemes to provide Forward Error Correction capabilities.

3. Requirements notation

[TOC](#)

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 [\[RFC2119\] \(Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," March 1997.\)](#).

4. Architecture Overview

[TOC](#)

The FEC Framework is described in terms of an additional protocol layer between the transport layer (e.g. UDP or DCCP) and Application and Transport Protocols running over this transport layer. Examples of such protocols are RTP, RTCP, etc. As such, the data path interface between the FEC Framework and both underlying and overlying layers can be thought of as being the same as the standard interface to the transport layer - i.e. the data exchanged consists of datagram payloads each associated with a single transport flow identified by the standard 5-tuple { Source IP Address, Source Transport Port, Destination IP Address, Destination Transport Port, Transport Protocol }.

The FEC Framework makes use of an FEC Scheme, in a similar sense to that defined in [\[RFC5052\] \(Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction \(FEC\) Building Block," August 2007.\)](#) and uses the terminology of that document. The FEC Scheme provides FEC encoding and decoding and describes the protocol fields and or procedures used to identify packet payload data in the context of the FEC Scheme. The interface between the FEC Framework and an FEC Scheme, which is described in this document, is a logical one, which exists for specification purposes only. At an encoder, the FEC Framework passes groups of transport packet payloads to the FEC Scheme for FEC Encoding. The FEC Scheme returns FEC repair packet payloads, encoded FEC Payload ID information for each of the repair packets and, in some cases, encoded FEC Payload ID information for each of the source packets. At a decoder, the FEC Framework passes transport packet payloads (source and repair) to the FEC Scheme and the FEC Scheme returns additional recovered source packet payloads.

This document defines certain FEC Framework Configuration Information which MUST be available to both sender and receiver(s). For example, this information includes the specification of the transport flows which are to be FEC protected, specification of the transport flow(s) which will carry the FEC protection (repair) data and the relationship(s) between these 'source' and 'repair' flows (i.e. which source flow(s) are protected by each repair flow. The FEC Framework

Configuration Information also includes information fields which are specific to the FEC Scheme. This information is analagous to the FEC Object Transmission Information defined in [\[RFC5052\] \(Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction \(FEC\) Building Block," August 2007.\)](#).

The FEC Framework does not define how the FEC Framework Configuration Information for the stream is communicated from sender to receiver. This must be defined by any Content Delivery Protocol specification as described in the following sections.

In this architecture we assume that the interface to the transport layer supports the concepts of payloads to be transported and identification transport flows on which those payloads are transported. Since this is an interface internal to the architecture, we do not specify this interface explicitly, except to say that transport flows which are distinct from the transport layer point of view (for example, distinct UDP flows as identified by the UDP source/destination ports/addresses) are also distinct on the interface between the transport layer and the FEC Framework.

The architecture outlined above is illustrated in the [Figure 1 \(FEC Framework Architecture\)](#).

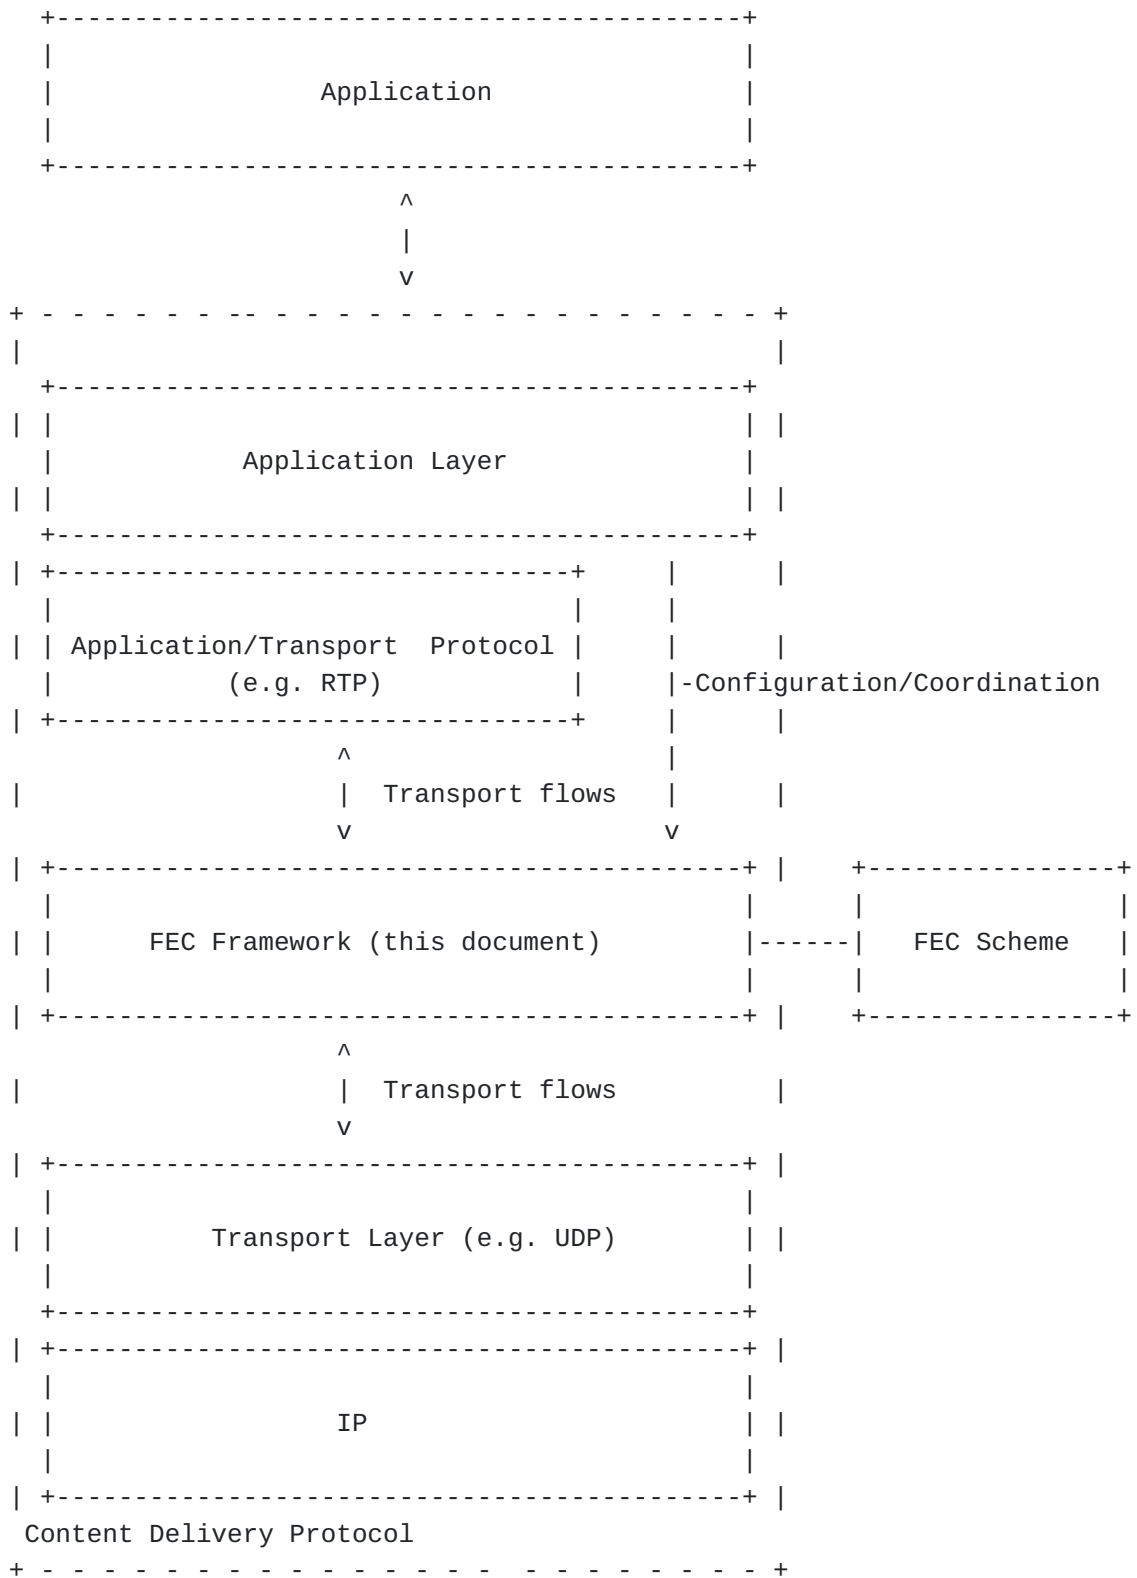


Figure 1: FEC Framework Architecture

The contents of the transport payload for repair packets is fully defined by the FEC Scheme. FEC Schemes MAY define a means for repair data to be carried over RTP, in which case the repair packet payload format starts with the RTP header. This specification provides guidelines for the use of RTP for the repair flows in this manner, however the explicit specification of this case is delegated to FEC Schemes.

The use of RTP for repair packets is independent of the protocols used for source packets: if RTP is used for source packets then repair packets may or may not use RTP and vice versa (although it is unlikely that there are useful scenarios where non-RTP source flows are protected by RTP repair flows). FEC Schemes are expected to recover entire transport payloads for recovered source packets in all cases. For example if RTP is used for source flows, the FEC Scheme is expected to recover the entire UDP payload, including the RTP header.

Editor's note: An alternative possibility would be to define a single RTP Payload Format (and associated MIME Type) for the carriage of FEC repair data for any FEC Scheme. In this case, FEC Scheme specifications need not mention RTP: the encapsulation of repair payloads in RTP packets will be defined for all schemes by a single RTP Payload Format specification.

5. Procedural overview

[TOC](#)

5.1. General

[TOC](#)

The mechanism defined in this document does not place any restrictions on the source transport payloads which can be protected together, except that the source transport payload is carried over a supported transport protocol (See [Section 7 \(Transport Protocols\)](#)). The data may be from multiple transport flows that are protected jointly. The FEC framework handles the packet flows as a sequence of 'source blocks' each consisting of a set of source transport payloads, possibly from multiple flows which are to be protected together. For example, each source block may be constructed from those source transport payloads related to a particular segment in time of the flow.

At the sender, the FEC Framework passes the payloads for a given block to the FEC Scheme for FEC encoding. The FEC Scheme performs the FEC encoding operation and returns the following information:

- *optionally, encoded FEC Payload IDs for each of the source payloads

- *one or more FEC repair packet payloads

- *encoded FEC Payload IDs for each of the repair packet payloads

The FEC framework then performs two operations: Firstly, it appends the FEC payload IDs, if provided, to each of the source transport payloads, and sends the resulting packets, known as 'FEC source packets', to the receiver and secondly it places the provided 'FEC repair packet payloads' and corresponding 'FEC Repair Payload IDs' appropriately to construct 'FEC repair packets' and send them to the receiver. Note that FEC repair packets MAY be sent to a different multicast group or groups from the source packets.

This document does not define how the sender determines which source transport payloads are included in which source blocks or the sending order and timing of FEC source and FEC repair packets. A specific Content Delivery Protocol MAY define this mapping or it MAY be left as implementation dependent at the sender. However, a CDP specification MUST define how a receiver determines the length of time it should wait to receive FEC repair packets for any given source block. The sequence of operations at the sender is described in more detail in [Section 5.2 \(Sender Operation\)](#).

At the receiver, original source transport payloads are recovered by the FEC Framework directly from any FEC Source packets received simply by removing the Source FEC Payload ID, if present. The receiver also passes the contents of the received FEC Source transport payloads, plus their FEC Payload IDs to the FEC Scheme for possible decoding.

If any FEC source transport payloads related to a given source block have been lost, then the FEC Scheme may perform FEC decoding to recover the missing source transport payloads (assuming sufficient FEC Source and FEC Repair packets related to that source block have been received).

Note that the receiver may need to buffer received source packets to allow time for the FEC Repair packets to arrive and FEC decoding to be performed before some or all of the received or recovered packets are passed to the application. If such a buffer is not provided, then the application must be able to deal with the severe re-ordering of packets that will be required. However, such buffering is Content Delivery Protocol and/or implementation-specific and is not specified here. The receiver operation is described in more detail in [Section 5.3 \(Receiver Operation\)](#)

The FEC Source packets MUST contain information which identifies the source block and the position within the source block (in terms

specific to the FEC Scheme) occupied by the packet. This information is known as the 'Source FEC Payload ID'. The FEC Scheme is responsible for defining and interpreting this information. This information MAY be encoded into a specific field within the FEC Source packet format defined in this specification, called the Explicit Source FEC Payload ID field. The exact contents and format of the Explicit Source FEC Payload ID field are defined by the FEC Scheme. Alternatively, the FEC Scheme MAY define how the Source FEC Payload ID is derived from other fields within the source packets. This document defines the way that the Explicit Source FEC Payload ID field is appended to source packets to form FEC Source packets.

The FEC Repair packets MUST contain information which identifies the source block and the relationship between the contained repair payloads and the original source block. This is known as the 'Repair FEC Payload ID'. This information MUST be encoded into a specific field, the Repair FEC Payload ID field, the contents and format of which are defined by the FEC Scheme.

The FEC Scheme MAY use different FEC Payload ID field formats for FEC Source packets and FEC Repair packets.

5.2. Sender Operation

[TOC](#)

It is assumed that the sender has constructed or received original data packets for the session. These may be RTP, RTCP, MIKEY or indeed any other type of packet. The following operations, illustrated in [Figure 2 \(Sender operation\)](#) describe a possible way to generate compliant FEC Source packet and FEC repair packet streams:

1. Source transport payloads are provided by the application.
2. A source block is constructed as specified in [Section 6.2 \(Structure of the source block\)](#).
3. The source block is passed to the FEC Scheme for FEC encoding. The Source FEC Payload ID information of each Source packet is determined by the FEC Scheme. If required by the FEC Scheme the Source FEC Payload ID is encoded into the Explicit Source FEC Payload ID field.
4. The FEC Scheme performs FEC Encoding, generating repair packet payloads from a source block and a Repair FEC Payload ID field for each repair payload.
5. The Explicit Source FEC Payload IDs (if used), Repair FEC Payload IDs and repair packet payloads are provided back from the FEC Scheme to the FEC Framework.

6. The FEC Framework constructs FEC Source packets according to [Section 6.3 \(Packet format for FEC Source packets\)](#) and FEC Repair packets according to [Section 6.4 \(Packet Format for FEC Repair packets\)](#)

using the FEC Payload IDs and repair packet payloads provided by the FEC Scheme.

7. The FEC Source and FEC Repair packets are sent using normal transport layer procedures. The port(s) and multicast group(s) to be used for FEC Repair packets are defined in the FEC Framework Configuration Information. The FEC Source packets are sent using the same transport flow identification information as would have been used for the original source packets if the FEC Framework were not present (for example, in the UDP case, the UDP source and destination addresses and ports on the eventual IP FEC Source Packet will be the same whether or not the FEC Framework is applied).

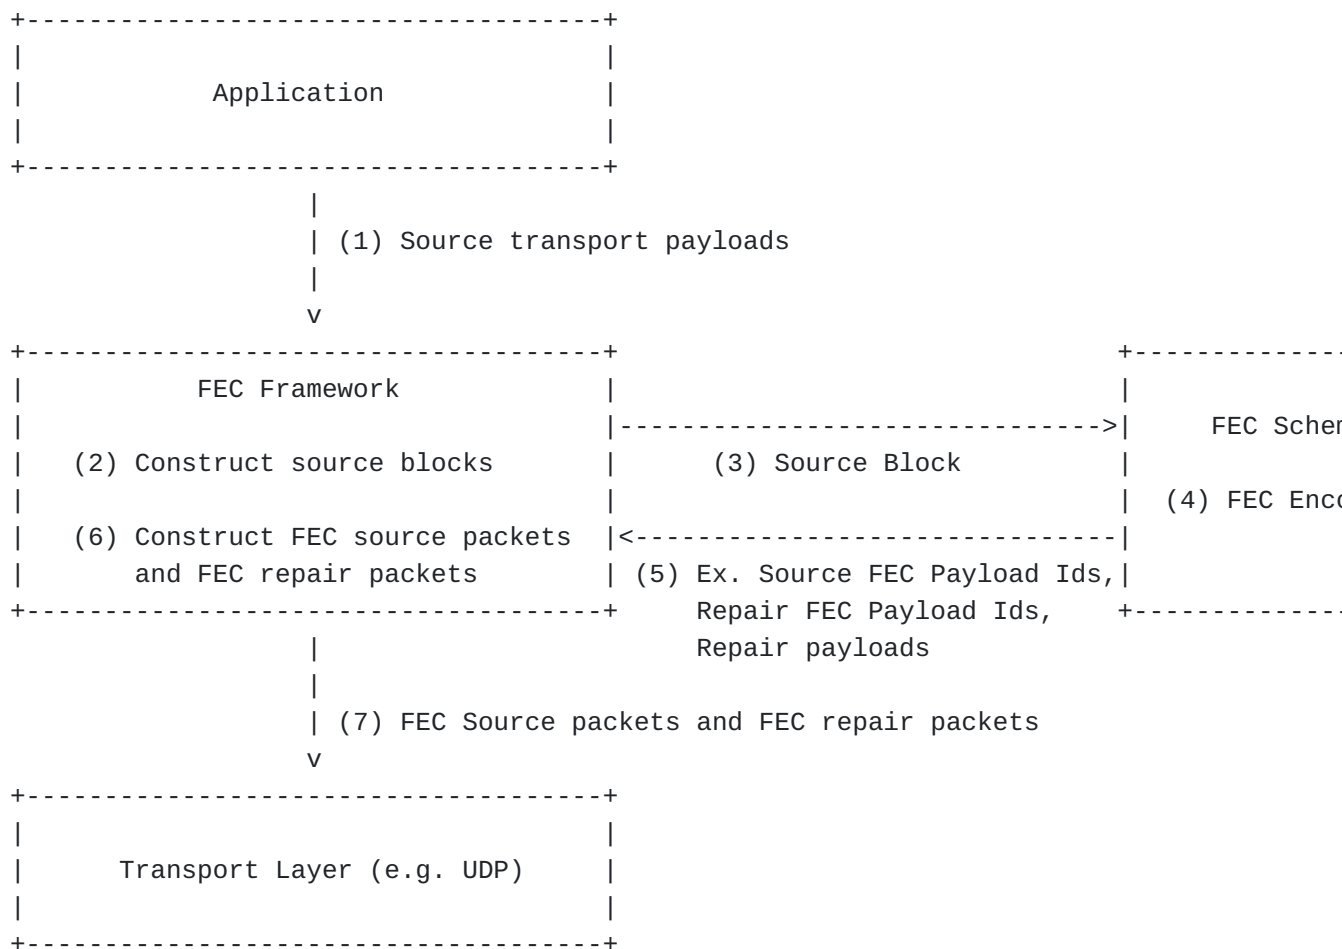


Figure 2: Sender operation

5.3. Receiver Operation

[TOC](#)

The following describes a possible receiver algorithm, illustrated in [Figure 3 \(Receiver Operation\)](#), when receiving an FEC source or repair packet:

1. FEC Source Packets and FEC Repair packets are received and passed to the FEC Framework. The type of packet (Source or Repair) and the transport flow to which it belongs (in the case of source packets) is indicated by the transport flow information which identifies the flow at the transport layer (for example source and destination ports and addresses in the case of UDP).
 2. The FEC Framework extracts the Explicit Source FEC Payload ID field (if present) from FEC Source Packets and the Repair FEC Payload ID from FEC Repair Packets.
 3. The Explicit Source FEC Payload IDs (if present), Repair FEC Payload IDs, FEC Source payloads and FEC Repair payloads are passed to the FEC Scheme.
 4. The FEC Scheme uses the received FEC Payload IDs (and derived FEC Source Payload IDs in the case that the Explicit Source FEC Payload ID field is not used) to group source and repair packets into source blocks. If at least one source packet is missing from a source block, and at least one repair packet has been received for the same source block then FEC decoding may be performed in order to recover missing source payloads. The FEC Scheme determines whether source packets have been lost and whether enough data for decoding of any or all of the missing source payloads in the source block has been received.
 5. The FEC Scheme returns the source transport payload to the FEC Framework in the form of source blocks containing received and decoded source packets and indications of any source packets which were missing and could not be decoded.
 6. The FEC Framework passes the received and recovered source packet payloads to the application.
-

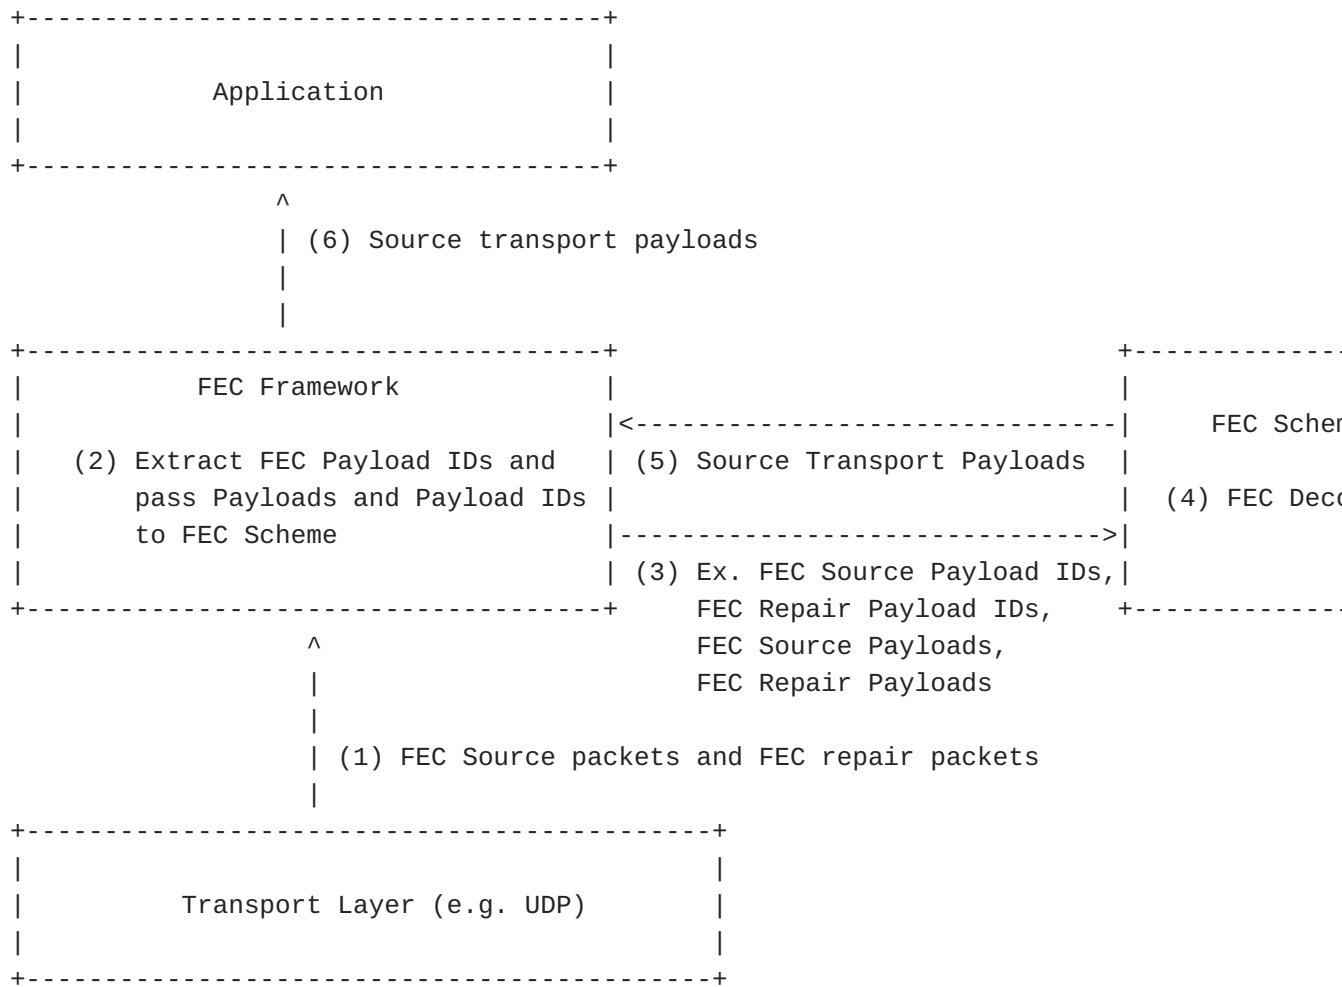


Figure 3: Receiver Operation

Note that the above procedure may result in a situation in which not all original source packets are recovered. Source packets which are correctly received and those which are reconstructed MAY be delivered to the application out of order and in a different order from the order of arrival at the receiver. Alternatively, buffering and packet re-ordering MAY be required to re-order received and reconstructed source packets into the order they were placed into the source block, if that is necessary according to the application.

6.1. General

[TOC](#)

This section specifies the protocol elements for the FEC Framework. Three components of the protocol are defined in this document and are described in the following sections:

1. Construction of a source block from source payloads. The FEC code will be applied to this source block to produce the repair payloads.
2. A format for packets containing source data.
3. A format for packets containing repair data.

The operation of the FEC Framework is governed by certain FEC Framework Configuration Information. This configuration information is also defined in this section. A complete protocol specification that uses this framework MUST specify the means to determine and communicate this information between sender and receiver.

6.2. Structure of the source block

[TOC](#)

The FEC Framework and FEC Scheme exchange source transport payload in the form of source blocks. A source block is generated by the FEC Framework from an ordered sequence of source transport payloads. The allocation of transport payloads to blocks is dependent on the application. Note that some source transport payloads may not be included in any block. For each source transport payload included in a source block, the following information is provided to the FEC Scheme:

- *A description of the source transport flow with which the transport payload is associated (See 6.5)
 - *The source transport payload itself
 - *The length of the source transport payload
-

6.3. Packet format for FEC Source packets

[TOC](#)

The packet format for FEC Source packets MUST be used to transport the payload of an original source packet. As depicted in [Figure 4 \(Structure of the FEC packet format for FEC Source packets\)](#), it consists of the original packet, optionally followed by the Explicit Source FEC Payload ID field. The FEC Scheme determines whether the

Explicit Source FEC Payload ID field is required. This determination is specific to each transport flow.




Figure 4: Structure of the FEC packet format for FEC Source packets

The FEC Source packets MUST be sent using the same transport flow as would have been used for the original source packets if the FEC Framework were not present. The Original transport Payload field MUST be identical to the source transport payload. The transport payload of the FEC Source packet MUST consist of the Original Transport Payload followed by the Explicit Source FEC Payload ID field, if required. The Explicit Source FEC Payload ID field contains information required to associate the source packet with a source block and for the operation of the FEC algorithm and is defined by the FEC Scheme. The format of the Source FEC Payload ID field is defined by the FEC Scheme. Note that in the case that the FEC Scheme or CDP defines a means to derive the Source FEC Payload ID from other information in the packet (for example the a sequence number of some kind used by the application protocol), then the Source FEC Payload ID field is not included in the packet. In this case the original source packet and FEC Source Packet are identical.

Since the addition of the Explicit Source FEC Payload ID increases the packet length, then in applications where avoidance of IP packet fragmentation is a goal, Content Delivery Protocols SHOULD consider the Explicit Source FEC Payload ID size when determining the size of source transport payloads that will be delivered using the FEC Framework.

Note: The Explicit Source FEC Payload ID is placed at the end of the packet so that in the case that Robust Header Compression [\[RFC3095\]](#) (Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed," July 2001.) or other header compression mechanisms are used and in the case that a ROHC profile is defined for the protocol carried within the transport payload (for example RTP),

then ROHC will still be applied for the FEC Source packets. Applications that may be used with this Framework should consider that FEC Schemes may add this Explicit Source FEC Payload ID and thereby increase the packet size.

6.4. Packet Format for FEC Repair packets

[TOC](#)

The packet format for FEC Repair packets is shown in [Figure 5 \(Packet format for repair packets\)](#). The transport payload consists of a Repair FEC Payload ID field followed by repair data generated in the FEC encoding process.

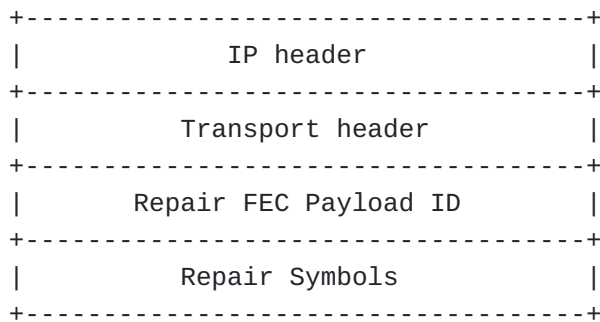


Figure 5: Packet format for repair packets

The Repair FEC Payload ID field contains information required for the operation of the FEC algorithm at the receiver. This information is defined by the FEC Scheme. The format of the Repair FEC Payload ID field is defined by the FEC Scheme.

6.4.1. Packet Format for FEC Repair packets over RTP (informative)

[TOC](#)

For FEC Schemes which specify the use of RTP for repair packets, the packet format for repair packets includes an RTP header as shown in [Figure 6 \(Packet format for repair packets\)](#).




Figure 6: Packet format for repair packets

6.5. FEC Framework Configuration Information

[TOC](#)

The FEC Framework Configuration Information is information that the FEC Framework needs in order to apply FEC protection to the transport flows. A complete Content Delivery Protocol specification that uses the framework specified here **MUST** include details of how this information is derived and communicated between sender and receiver.

The FEC Framework Configuration Information includes identification of a set of source packet flows. For example, in the case of UDP, each packet flow is uniquely identified by a tuple { Source IP Address, Destination IP Address, Source UDP port, Destination UDP port }. Note that in some applications some of these fields may be wildcarded, so that the flow is identified by a subset of the fields and in particular in many applications the limited tuple { Destination IP Address, Destination UDP port } is sufficient.

A single instance of the FEC Framework provides FEC protection for all packets of a specified set of source packet flows, by means of one or more packet flows consisting of repair packets. The FEC Framework Configuration Information includes, for each instance of the FEC Framework:

1. Identification of the packet flow(s) carrying FEC Repair packets, known as the FEC repair flow(s).
2. For each source packet flow protected by the FEC repair flow(s):
 - a. Definition of the packet flow carrying source packets (for example, by means of a tuple as describe above for UDP).

- b. An integer identifier for this flow definition (i.e. tuple). This identifier MUST be unique amongst all source packet flows which are protected by the same FEC repair flow.
3. The FEC Encoding ID, identifying the FEC Scheme
4. The length of the Explicit Source FEC Payload Id, in bytes
5. An opaque container for FEC-Scheme-specific information

Multiple instances of the FEC Framework, with separate and independent FEC Framework Configuration Information, may be present at a sender or receiver. A single instance of the FEC Framework protects all packets of all the source packet flows identified in (2) above i.e. all packets sent on those flows MUST be FEC Source packets as defined in [Section 6.3 \(Packet format for FEC Source packets\)](#). A single source packet flow may be protected by multiple instances of the FEC Framework.

The integer flow identifier identified in 2(b) is a "shorthand" to identify source flows between the FEC Framework and the FEC Scheme. The reason for defining this as an integer, and including it in the FEC Framework Configuration Information is so that the FEC Scheme at the sender and receiver may use it to identify the source flow with which a recovered packet is associated. The integer flow identifier may therefore take the place of the complete flow description (e.g. UDP 4-tuple).

Whether and how this flow identifier is used is defined by the FEC Scheme. Since source packets are directly associated with a flow by virtue of their packet headers, this identifier need not be carried in source packets. Since repair packets may provide protection for multiple source flows, this flow identifier would likely not be carried directly in repair packets. However, the flow identifier associated with a particular source packet may be recovered from the repair packets as part of an FEC decoding operation. Integer flow identifiers SHOULD be allocated starting from zero and increasing by one for each flow.

A single FEC repair flow provides repair packets for a single instance of the FEC Framework. Other packets MUST NOT be sent within this flow i.e. all packets in the FEC repair flow MUST be FEC repair packets as defined in [Section 6.4 \(Packet Format for FEC Repair packets\)](#) and MUST relate to the same FEC Framework instance.

In the case that RTP is used for repair packets, the identification of the repair packet flow MAY also include the RTP Payload Type to be used for repair packets.

6.6. FEC Scheme requirements

In order to be used with this framework, an FEC Scheme MUST be capable of processing data arranged into blocks of source transport packet payloads (source blocks).

A specification for a new FEC scheme MUST include the following things:

1. The FEC Encoding ID value that uniquely identifies the FEC scheme. This value MUST be registered with IANA as described in [Section 10 \(IANA Considerations\)](#).
2. The type, semantics and encoding format of the Repair FEC Payload ID.
3. The type, semantics and encoding format of the FEC Scheme-specific FEC Framework Configuration Information.
4. A full specification of the FEC code.

This specification MUST precisely define the valid FEC-Scheme-Specific FEC Framework Configuration Information values, the valid FEC Payload ID values and the valid packet payload sizes (where packet payload refers to the space - not necessarily contiguous - within a packet dedicated to carrying encoding symbol bytes).

Furthermore, given a source block as defined in [Section 6.2 \(Structure of the source block\)](#), valid values of the FEC-Scheme-Specific FEC Framework Configuration Information, a valid Repair FEC Payload ID value and a valid packet payload size, the specification MUST uniquely define the values of the encoding symbol bytes to be included in the repair packet payload of a packet with the given Repair FEC Payload ID value.

A common and simple way to specify the FEC code to the required level of detail is to provide a precise specification of an encoding algorithm which, given a source block, valid values of the FEC-Scheme-Specific FEC Framework Configuration Information, a valid Repair FEC Payload ID value and a valid packet payload size as input produces the exact value of the encoding symbol bytes as output.

5. A description of practical encoding and decoding algorithms.

This description need not be to the same level of detail as for the encoding above, however it must be sufficient to demonstrate that encoding and decoding of the code is both possible and practical.

FEC scheme specifications MAY additionally define the following:

1. Type, semantics and encoding format of an Explicit Source FEC Payload ID.

Whenever an FEC scheme specification defines an 'encoding format' for an element, this must be defined in terms of a sequence of bytes which can be embedded within a protocol. The length of the encoding format MUST either be fixed or it must be possible to derive the length from examining the encoded bytes themselves. For example, the initial bytes may include some kind of length indication.

FEC scheme specifications SHOULD use the terminology defined in this document and SHOULD follow the following format:

- 1. Introduction** <describe the use-cases addressed by this FEC scheme>
- 2. Formats and Codes**
 - 2.1 Source FEC Payload ID(s)** <Either, define the type and format of the Explicit Source FEC Payload ID, or define how Source FEC Payload ID information is derived from source packets>
 - 2.2 Repair FEC Payload Id** <Define the type and format of the Repair FEC Payload ID>
 - 2.3 FEC Framework Configuration Information** <Define the type and format of the FEC Scheme-specific FEC Framework configuration information>
- 3. Procedures** <describe any procedures which are specific to this FEC scheme, in particular derivation and interpretation of the fields in the FEC Payload ID and FEC Scheme-specific FEC Framework configuration information.>
- 4. FEC code specification** <provide a complete specification of the FEC Code>

Specifications MAY include additional sections, for example, examples. Each FEC scheme MUST be specified independently of all other FEC schemes; for example, in a separate specification or a completely independent section of larger specification (except, of course, a specification of one FEC Scheme may include portions of another by reference).

7. Transport Protocols

The following transport protocols are supported:

- *User Datagram Protocol (UDP)

- *Datagram Congestion Control Protocol (DCCP)

Editor's note: This section will contain transport-specific considerations, if any.

8. Congestion Control

[TOC](#)

This section starts with a informative section on the motivation of the normative requirements for congestion control, which are spelled out in [Section 8.1 \(Normative requirements\)](#).

Informative Note: The enforcement of Congestion Control (CC) principles has gained a lot of momentum in the IETF over the recent years. While the need of CC over the open Internet is unquestioned, and the goal of TCP friendliness is generally agreed for most (but not all) applications, the subject of congestion detection and measurement in heterogenous networks can hardly be considered as solved. Most congestion control algorithms detect and measure congestion by taking (primarily or exclusively) the packet loss rate into account. This appears to be inappropriate in environments where a large percentage of the packet losses are the result link-layer errors and independent of the network load. Note that such environments exist in the "open Internet", as well as in "closed" IP based networks. An example for the former would be the use of IP/UDP/RTP based streaming from an Internet-connected streaming server to a device attached to the Internet using cellular technology.

The authors of this draft are primarily interested in applications where the application reliability requirements and end-to-end reliability of the network differ, such that it warrants higher layer protection of the packet stream - for example due to the presence of unreliable links in the end-to-end path - and where real-time, scalability or other constraints prohibit the use of higher layer (transport or application) feedback. A typical example for such applications is multicast and broadcast streaming or multimedia transmission over heterogenous networks. In other cases, application reliability requirements may be so high that the required end-to-end reliability is difficult to achieve even over wired networks. Furthermore the end-to-end network reliability may not be known in advance.

This FEC framework is not proposed, nor intended, as a QoS enhancement tool to combat losses resulting from highly congested networks. It should not be used for such purposes.

In order to prevent such mis-use, one approach would be to leave standardisation to bodies most concerned with the problem described above. However, the IETF defines base standards used by several bodies, including DVB, 3GPP, 3GPP2, all of which appear to share the environment and the problem described.

Another approach would be to write a clear applicability statement - for example restricting use of the framework to networks with wireless links. However, there may be applications where the use of FEC may be justified to combat congestion-induced packet losses - particularly in lightly loaded networks, where congestion is the result of relatively rare random peaks in instantaneous traffic load - thereby intentionally violating congestion control principles. One possible example for such an application could be a no-matter-what, brute-force FEC protection of traffic generated as an emergency signal.

We propose a third approach, which is to require at a minimum that the use of this framework with any given application, in any given environment, does not cause congestion issues which the application alone would not itself cause i.e. the use of this framework must not make things worse.

Taking above considerations into account, the normative text of this section implements a small set of constraints for the FEC, which are mandatory for all senders compliant with this FEC framework. Further restrictions may be imposed for certain Content Delivery Protocols. In this it follows the spirit of the congestion control section of RTP and its Audio-Visual Profile (RFC3550/STD64 and RFC3551/STD65).

One of the constraints effectively limits the bandwidth for the FEC protected packet stream to be no more than roughly twice as high as the original, non-FEC protected packet stream. This disallows the (static or dynamic) use of excessively strong FEC to combat high packet loss rates, which may otherwise be chosen by naively implemented dynamic FEC-strength selection mechanisms. We acknowledge that there may be a few exotic applications, e.g. IP traffic from space-based senders, or senders in certain hardened military devices, which would warrant a higher FEC strength. However, in this specification we give preference to the overall stability and network friendliness of the average application, and for those a factor of 2 appears to be appropriate.

A second constraint requires that the FEC protected packet stream be in compliance with the congestion control in use for the application and network in question.

8.1. Normative requirements

[TOC](#)

The bandwidth of FEC Repair packet flows MUST NOT exceed the bandwidth of the source packet flows being protected. In addition, whenever the source packet flow bandwidth is adapted due to the operation of congestion control mechanisms, the FEC repair packet flow bandwidth MUST be similarly adapted.

9. Security Considerations

[TOC](#)

The application of FEC protection to a stream does not provide any kind of security protection.

If security services are required for the stream, then they MUST either be applied to the original source transport payload before FEC protection is applied, or to both the source and repair data, after FEC protection has been applied.

If integrity protection is applied to source packets before FEC protection is applied, and no further integrity protection is applied to repair packets, then a denial of service attack is possible if an attacker is in a position to inject fake repair transport payloads. If received by a receiver, such fake repair transport payloads could cause incorrect FEC decoding resulting in incorrect source transport payloads being passed up to the application protocol. Such incorrect packets would then be detected by the source integrity protection and discarded, resulting in partial or complete denial of service.

Therefore, in such environments, integrity protection MUST also be applied to the FEC repair transport payloads, for example using IPsec. Receivers MUST also verify the integrity of source transport payloads before including the source transport payload into the source block for FEC purposes.

It is possible that multiple streams with different confidentiality requirements (for example, the streams may be visible to different sets of users) can be FEC protected by a single repair stream. This scenario is not recommended, since resources will be used to distribute and decode data which cannot then be decrypted by at least some receivers. However, in this scenario, confidentiality protection MUST be applied before FEC encoding of the streams, otherwise repair transport payload may be used by a receiver to decode unencrypted versions of source streams which they do not have permission to view.

10. IANA Considerations

[TOC](#)

tbd

11. Acknowledgments

[TOC](#)

This document is based in large part on [\[I-D.watson-tsvwg-fec-sf\] \(Watson, M., "Forward Error Correction \(FEC\) Streaming Framework," July 2005.\)](#) and so thanks are due to the additional authors of that document, Mike Luby, Magnus Westerlund and Stephan Wenger. That document was in turn based on the FEC streaming protocol defined by 3GPP in [\[MBMSTS\] \(3GPP, "Multimedia Broadcast/Multicast Service \(MBMS\); Protocols and codecs," April 2005.\)](#) and thus thanks are also due to the participants in 3GPP TSG SA working group 4.

12. References

[TOC](#)

[RFC2119]	Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC3095]	Bormann, C., Burmeister, C., Degermark, M., Fukushima, H., Hannu, H., Jonsson, L-E., Hakenberg, R., Koren, T., Le, K., Liu, Z., Martensson, A., Miyazaki, A., Svanbro, K., Wiebke, T., Yoshimura, T., and H. Zheng, "RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed," RFC 3095, July 2001 (TXT).
[RFC5052]	Watson, M., Luby, M., and L. Vicisano, "Forward Error Correction (FEC) Building Block," RFC 5052, August 2007 (TXT).
[I-D.watson-tsvwg-fec-sf]	Watson, M., "Forward Error Correction (FEC) Streaming Framework," draft-watson-tsvwg-fec-sf-00 (work in progress), July 2005 (TXT).
[MBMSTS]	3GPP, "Multimedia Broadcast/Multicast Service (MBMS); Protocols and codecs," 3GPP TS 26.346, April 2005.

Author's Address

[TOC](#)

	Mark Watson
	Digital Fountain
	39141 Civic Center Drive
	Suite 300
	Fremont, CA 94538
	U.S.A.
Email:	mark@digitalfountain.com

Full Copyright Statement

[TOC](#)

Copyright © 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.