**Internet Engineering Task Force** **RMT WG**
**INTERNET-DRAFT**
**draft-ietf-rmt-bb-fec-03.ps**

**M.Luby/Digital Fountain**
**L.Vicisano/Cisco**
**J.Gemmell/Microsoft**
**L.Rizzo/ACIRI and Univ. Pisa**
**M.Handley/ACIRI**
**J. Crowcroft/UCL**

**19 July 2001**
**Expires: January 2002**

# RMT BB Forward Error Correction Codes

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 of RFC2026.

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 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 a "work in progress".

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

To view the list Internet-Draft Shadow Directories, see http://www.ietf.org/shadow.html.

This document is a product of the IETF RMT WG. Comments should be addressed to the authors, or the WG's mailing list at rmt@isi.edu.

**Abstract**

This memo describes the abstract packet formats and IANA registration procedures for use of Forward Error Correction (FEC) codes within the context of reliable IP multicast transport. This memo should be read in conjunction with and uses the terminology of the companion memo [1], which describes the use of Forward Error Correction (FEC) codes within the context of reliable IP multicast transport and provides an introduction to some commonly used FEC codes.

## 1. FEC Abstract Packet Fields and Out-of-Band Information

This section specifies the information that protocol packets must carry to implement the various forms of FEC-based reliability. A session is defined to be all the information associated with a transmission of data about a particular object by a single sender. There are three classes of packets that may contain FEC information within a session: data packets, session-control packets and feedback packets. They generally contain different kinds of FEC information. Note that some protocols do not use feedback packets.

Data packets may sometime serve as session-control packets as well; both data and session-control packets generally travel downstream (from the sender towards receivers) and are

addressed to a multicast IP address (sometime they might be addressed to the unicast address of a specific receiver). In the following, for simplicity we will refer to both data and session control packets as downstream-traveling packets, or simply downstream packets.

As a general rule, feedback packets travel upstream (from receivers to the sender) and are addressed to the unicast address of the sender. Sometimes, however, they might be addressed to a multicast IP address or to the unicast address of a receiver or also the the unicast address of some different node (an intermediate node that provides recovery services or a neighboring router).

The FEC-related information that can be present in downstream packets can be classified as follows:

1) FEC Encoding Identifier

   Identifies the FEC encoding being used and has the purpose of allowing receivers to select the appropriate FEC decoder. As a general rule, the "FEC Encoding Identifier" MUST be the same for a given session, i.e., for all transmission of data related to a particular object, but MAY vary across different transmissions of data about different objects in different sessions, even if transmitted using the same set of multicast groups and/or using a single upper-layer session.

2) FEC payload ID

   Identifies the symbol(s) in the payload of the packet. The content of this piece of information depends on the encoder being used (e.g. in Block FEC codes this may be the combination of block index and encoding symbol identifier; in ideal expandable FEC codes this may be just a flat encoding symbol identifier).

3) FEC Object Transmission Information

   This is information regarding the encoding of a specific object needed by the FEC decoder (e.g. for Block FEC codes this may be the combination of block length and object length). This might also include general parameters of the FEC encoder. Note that, in strict terms, the "FEC Encoding Identifier" belongs to this class of information, however, for practical purposes, we will consider it separately.

All the classes of information above, except 2), can either be transmitted within the transport session (using protocol packet-header fields) or out of band. The information described in 2) MUST be transmitted in data-packet header fields, as it provides a description of the data contained in the packet. In the following we will specify the content of 1), 2) and 3) regardless of whether these pieces of information are transmitted in protocol packets or out of band. This document neither specifies out-of-band methods to transport the information nor does it specify the way out-of-band information is associated with packet payloads. This last task is devolved to an upper-layer protocol.

Within the context of FEC repair schemes, feedback packets are (optionally) used to request FEC retransmission. The FEC-related information present in feedback packets usually contains an FEC Block Identifier, that defines the block that is being repaired, and the number of Repair Symbols requested. Although this is the most common case, variants are possible in which the receivers provide more specific information about the Repair Symbols requested (e.g. an index range or a list of symbols accepted). It is also possible to include multiple of these request in a single feedback packet.

This document does not provide any detail about the format of FEC information in feedback packets.

## 1.1. FEC Encoding Identifier

This is a numeric index that identifies a specific FEC scheme OR a class of encoding schemes that share the same format of "FEC Payload ID" and "FEC Object Transmission Information".

The FEC Encoding Identifier identifies a specific FEC scheme when the encoding scheme is formally and fully specified, in a way that independent implementors can implement both encoder and decoder from the specification. Companion documents of this specification may specify such FEC schemes and associate them with "FEC Encoding Identifier" values. These documents MUST also specify a correspondent format for the "FEC Payload ID" and "FEC Object Transmission Information". Currently FEC Encoding Identifiers in the range 0-127 are reserved for this class of encoding schemes (fully-specified schemes).

It is possible that a FEC scheme cannot be completely specified or that such a specification is simply not available or also that a party exists that owns the encoding scheme and it is not willing to disclose its algorithm. We refer to these encoding schemes as "Under-Specified" schemes. Under-specified schemes can still be identified as follows:

- A format of the fields "FEC Payload ID" and "FEC Object Transmission Information" MUST be defined for the encoding scheme.
- A value of "FEC Encoding Identifier" MUST be reserved and associated to the format definitions above. An already reserved "FEC Encoding Identifier" MUST be reused if it is associated to the same format of "FEC Payload ID" and "FEC Object Transmission Information" as the ones needed for the new under-specified FEC scheme.
- A value of "FEC Encoding Name" must be reserved (see below).

An Under-specified FEC scheme is completely identified by the tuple (FEC Encoding Identifier, FEC Encoding Name). The tuple MUST identify a single scheme that has at least one implementation. The party that owns this tuple MUST be able to provide information on how to obtain the under-specified FEC scheme identified by the tuple (e.g. a pointer to a publicly available reference-implementation or the name and contacts of a company that sells it, either separately or embedded in another product).

"FEC Encoding Names" are numeric identifiers scoped by a FEC Encoding Identifier.

The FEC Encoding Name MUST be part of the "FEC Object Transmission Information" and must be communicated to receivers together with the FEC Encoding Identifier.

Different FEC schemes that share the same FEC Encoding Identifier -- but have different FEC Encoding Names -- also share the same format of FEC Payload ID and FEC Object Transmission Information.

This specification reserves the range 0-127 of FEC Encoding Identifiers for fully-specified encoding schemes and the range 128-255 for under-specified encoding schemes.

### 1.2. FEC Payload ID and FEC Object Transmission Information

A document that specifies an encoding scheme and reserves a value of FEC Encoding Identifier MUST define a packet-field format for FEC Payload ID and FEC Object Transmission Information according to the need of the encoding scheme. This also applies to documents that reserves values of FEC Encoding Identifiers for under-specified encoding schemes. In this case the FEC Object Transmission Information must also include a field to contain the "FEC Encoding Name".

A packet field definition of FEC Object Transmission Information MUST be provided despite the fact that a protocol instantiation may decide to communicate this information out of band.

All packet field definitions (FEC Payload ID and FEC Object Transmission Information) MUST be fully specified at the level of bit-fields and they must have a length that is a multiple of a 4-byte word (this is to facilitate the alignment of packet fields in protocol instantiations).

Note that the specification of FEC Payload ID MUST allow an implementation-independent identification of the packet payload even in the case of Under-Specified encoding schemes.

## 2. Preassigned FEC Encoding Identifiers

This section specifies the FEC Encoding Identifier and the relative packets field for a number of known schemes that follow under the class of under-specified FEC schemes. Others may be specified in companion documents.

### 2.1. Small Block, Large Block and Expandable FEC Codes

This section reserves a FEC Encoding Identifier for the families of codes described in [1], and specifies the relative packet fields. Under-specified FEC schemes that belong to this class MUST use this identifier and packet field definitions.

The FEC Encoding Identifier for under specified codes assigned to Small Block, Large Block, and Expandable FEC Codes is 128.

The FEC Payload ID is composed of an encoding symbol identifier and an encoding block number structured as follows:

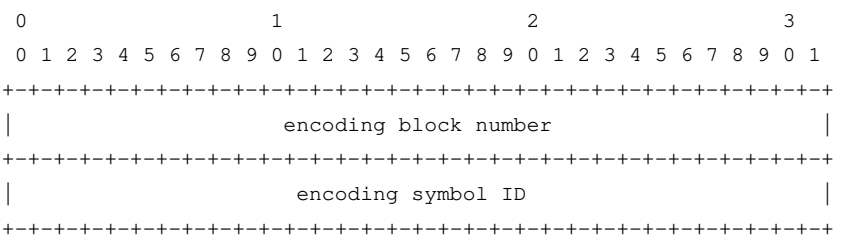


In addition, a one bit FEC Encoding Flag MAY be included, and this flag indicates whether the encoding symbol(s) in the payload of the packet are source symbol(s) or redundant symbol(s).

The FEC Object Transmission Information has the following structure:

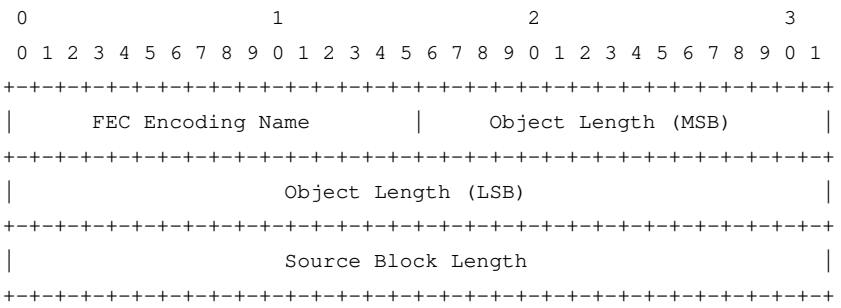


Note that this structure limits the range of possible FEC Encoding Names to 0-:-65536, despite the fact that the FEC Object Transmission Information can also be transmitted out of band.

The Object Length, composed of a Most Significant Bytes portion (MSB) and a Least Significant Bytes portion (LSB), is expressed in bytes.

## 2.2. Small Block Systematic FEC Codes

The following definitions apply to systematic codes of small length (total block length < 2^16).

The FEC Encoding Identifier for under specified codes assigned to Small Block Systematic FEC Codes is 129.

Although these codes can generally be accommodated by the FEC Encoding Identifier described in Section 2.1, a specific Encoding-ID is defined for systematic codes to allow better flexibility and to retain header compactness. The small source block length and small exapansion factor that often characterize systematic codes may require that the data source changes frequently the source block length. To allow the dynamic variation of the source block length and to communicate it to the receivers with low overhead, the block length is included in the FEC Payload ID.

The FEC Payload ID is composed of the encoding block number, the encoding symbol identifier and the block length:

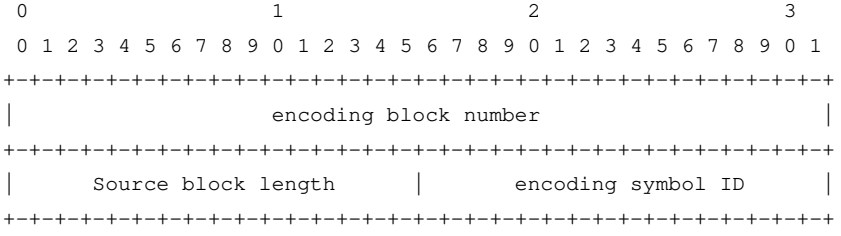


The FEC Object Transmission Information, when used, has the following structure:

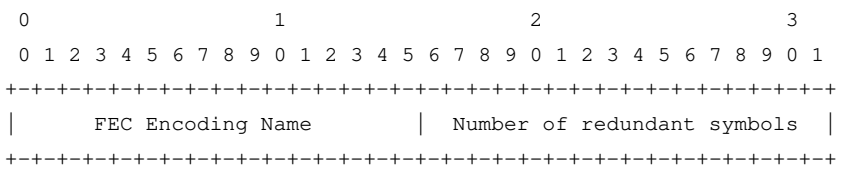


When variable block-length is used, the number of redundant symbols is to be interpreted as a maximum value (upper bound). This field is provided as an indication to allow receives to preallocate a decoder suitable for all the object blocks.

Note that this structure limits the range of possible FEC Encoding Names to 0-:-65536, despite the FEC Object Transmission Information can also be transmitted out of band.

## 3. IANA Considerations

Values of FEC Encoding Identifiers and FEC Encoding Names are subject to IANA registration. FEC Encoding Identifiers and FEC Encoding Names are hierarchical: FEC Encoding Identifiers (at the top level) scope ranges of FEC Encoding Names. Not all FEC Encoding Identifiers have a corresponding FEC Encoding Name scope (see below).

A FEC Encoding Identifier is a numeric non-negative index. Value from 0 to 127 are reserved for FEC encoders that are fully specified, as described in section 3.1. The assignment of a FEC Encoding Identifier in this range can only be granted if the requestor can provide such a specification published as an IETF RFC. Values greater than 127 can be assigned to under-specified encoding schemes. Note: this specification already assigns the values 128 and 129.

In any case values of FEC Encoding Identifiers can only be assigned if the required FEC packet fields corresponding to it (see section 1.2) are specified in a RFC.

Each FEC Encoding Identifier assigned to an under-specified encoding scheme scopes an independent range of FEC Encoding Names (i.e. the same value of FEC Encoding Name can be reused for different FEC Encoding Identifiers). An FEC Encoding Name is a numeric non-negative index. The document that reserves a FEC Encoding Identifier MAY also specify a range for the subordinate FEC Encoding Names.

Under the scope of a FEC Encoding Identifier, FEC Encoding Names are assigned on a First Come First Served base to requestors that are able to provide point of contact information and a pointer to publicly accessible documentation describing the FEC scheme and ways to obtain it (e.g. a pointer to a publicly available reference-implementation or the name and contacts of a company that sells it, either separately or embedded in another product). The requestor is responsible for keeping this information up to date.

## 4. Acknowledgments

Brian Adamson contributed to this document by shaping Section 2.2 and providing general feedback. We also wish to thank Vincent Roca for his comments.

## 5. References

[1] Luby, M., Vicisano, Gemmell, J., L., Rizzo, L., Handley, M., Crowcroft, J., "The use of Forward Error Correction in Reliable Multicast", Internet draft draft-ietf-rmt-info-fec-00.txt, November 2000.

## 6. Authors' Addresses

Michael Luby
luby@digitalfountain.com
Digital Fountain, Inc.
39141 Civic Center Drive
Suite 300
Fremont, CA 94538

Lorenzo Vicisano
lorenzo@cisco.com
cisco Systems, Inc.
170 West Tasman Dr.,
San Jose, CA, USA, 95134

Jim Gemmell
jgemmell@microsoft.com
Microsoft Research
301 Howard St., #830
San Francisco, CA, USA, 94105

Luigi Rizzo
luigi@iet.unipi.it
ACIRI, 1947 Center St., Berkeley CA 94704
and
Dip. di Ing. dell'Informazione
Universita' di Pisa
via Diotisalvi 2, 56126 Pisa, Italy

Mark Handley
mjh@aciri.org
ACIRI
1947 Center St.
Berkeley CA, USA, 94704

Jon Crowcroft
J.Crowcroft@cs.ucl.ac.uk
Department of Computer Science
University College London
Gower Street,
London WC1E 6BT, UK

## 7. Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS 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."