

INTERNATIONAL TELECOMMUNICATION UNION

STUDY GROUP 17

**TELECOMMUNICATION
STANDARDIZATION SECTOR**

TD 1162

STUDY PERIOD 2001-2004

English only

Original: English

Question(s): 5/17

Geneva, 10-19 March 2004

TEMPORARY DOCUMENT

Source: Editor

Title: Draft new Recommendation X.84 – Support of frame relay services over MPLS core networks

Editor's note:

Attached is a clean copy of Draft Recommendation X.84. TD 1143 Rev3 is the source document used to produce this clean version.

Contact: Claude Kawa, Canada

Email: claude.kawa@hec.ca

© ITU 2004

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

ITU-T Recommendation X.84

SUPPORT OF FRAME RELAY SERVICES OVER MPLS CORE NETWORKS

(Geneva 2004)

Summary

MPLS has the potential to consolidate core networks and services, for example frame relay services, over a single common core infrastructure. This Recommendation defines frame relay over MPLS core network architecture, the transfer of frame relay control and user data information over the core MPLS network and the user plane protocol stack used at the edges of the MPLS core network. Two mapping modes for frame relay over MPLS are defined: the "one-to-one" mapping mode and the "many-to-one mapping mode". The signalling and management planes are outside of the scope of this Recommendation.

CONTENTS

1	SCOPE	4
2	REFERENCES.....	4
3	DEFINITIONS.....	4
4	ACRONYMS	5
5	CONVENTIONS	6
6	ARCHITECTURE	6
6.1	GENERAL	6
6.2	MPLS TUNNEL AND VC LSPs	7
6.3	RELATIONSHIP BETWEEN FRAME RELAY VC AND MPLS VC LSP.....	8
6.4	FRAME RELAY OVER MPLS MAPPING MODES.....	9
7	FRAME RELAY OVER MPLS REQUIREMENTS	9
8	PROTOCOL STACK AND FRAME FORMAT	10
8.1	DATA TRANSFER PROTOCOL STACK	10
8.2	FR-MPLS PACKET FORMAT FOR THE ONE-TO-ONE MAPPING MODE	10
9	FR-MPLS PACKET PROCESSING FOR ONE-TO-ONE MAPPING MODE.....	13
9.2	RECEPTION OF FR-MPLS PACKETS	15
9.3	HANDLING OF ERROR CONDITIONS	16
9.4	OPTIONAL FRAGMENTATION AND RE-ASSEMBLY PROCEDURES	17
10	FR PVC PROVISIONING	19
11	TRAFFIC MANAGEMENT ASPECTS	19
12	FRAME RELAY MANY-TO-ONE MAPPING MODE	20
12.1	GENERAL	20
12.2.	PACKET FORMAT FOR MANY-TO-ONE MAPPING MODE	21
12.3.	MANY-TO-ONE MODE PROCESSING.....	23
	APPENDIX I	24
	APPENDIX II	25

1 Scope

MPLS has the potential to consolidate core networks and services, for example frame relay services, over a single common core infrastructure. This Recommendation defines the architecture for frame relay over MPLS core network, the transfer of frame relay control and user-data information and the user plane protocol stack used at the edges of the MPLS core network.

This Recommendation defines two mapping modes for the provision of Frame Relay service over MPLS. One-to-one mode characterized by a one to one relation between a frame relay VC and a pair of unidirectional MPLS LSPs. The other mode is known as many-to-one mode. In this mode all FR VCs between a pair of frame relay devices (CEs) controlled by a single frame relay signalling channel may be transparently transported in one pair of MPLS LSPs.

Note: This Recommendation does not cover Switch Virtual Connection and Soft PVC signalling between PEs. The control protocol for PVC status monitoring is for further study. The signalling and management planes are outside of the scope of this Recommendation.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; all users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published.

- ITU-T Recommendation X.36 (2003), Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for public data networks providing frame relay data transmission service by dedicated circuit, Geneva.
- ITU-T Recommendation X.76 (2003), Network-to-network interface between public networks providing PVC and/or SVC frame relay data transmission service, Geneva.
- ITU-T Recommendation X.146, Performance Objectives and Quality of Service Classes Applicable to Frame Relay, Geneva, September 1998.
- IETF RFC 3031, Multiprotocol label switching architecture, January 2001.
- IETF RFC 3032, MPLS Label Stack encoding, January 2001.
- IETF RFC 3270, Multiprotocol label switching support of differentiated services, May 2002.

3 Definitions

Customer Edge: A Customer Edge is the customer device connected to a provider edge device. Also known as frame relay DTE.

Egress Provider edge: The Provider Edge device that is the receiver of FR-MPLS packets.

FR-MPLS packet: Packets exchanged between an ingress provider edge and an egress provider edge.

Ingress Provider edge: The Provider Edge that transmits FR-MPLS packets.

Provider Edge: A Provider Edge (PE) is a network edge device that provides a frame relay service over an MPLS network.

Label Switched Path: A Label Switched Path (LSP) is the path through one or more MPLS nodes at one level of the hierarchy over which packets in a particular forwarding equivalence class are transmitted.

MPLS node: An MPLS node is a device that is aware of MPLS control protocols, operates one or more layer three routing protocols and is capable of forwarding packets based on MPLS LSP labels.

Penultimate Hop Popping: In MPLS architecture, penultimate hop popping is the mechanism by which the penultimate node (the node immediately before the egress node) pops the top label of the label stack prior to forwarding the packet to the egress node. Performing label popping by the penultimate node provides the egress node with the possibility to process optimally the packets.

4 Acronyms

B _c	Committed Burst size
B _e	Excess Burst size
BECN	Backward Explicit Congestion Notification
CE	Customer Edge
CIR	Committed Information Rate
C/R	Command / Response indicator
DCE	Data Communication Equipment
DE	Discard Eligibility
DLCI	Data Link Connection Identifier
DTE	Data Terminal Equipment
FCS	Frame Check Sequence
FECN	Forward Explicit Congestion Notification
FR	Frame Relay
FRS	Frame Relay Service
HDLC	High-level Data Link Control
IETF	Internet Engineering Task Force
LSP	Label Switched Path
MPLS	Multi Protocol Label Switching
MTU	Maximum Transmission Unit
NNI	Network-to-Network Interface
PDU	Protocol Data Unit
PE	Provider Edge
PHB	Per Hop Behavior
PHP	Penultimate Hop Popping
PVC	Permanent Virtual Connection

QoS	Quality of Service
RFC	Request For Comments
SVC	Switched Virtual Connection
UNI	User-to-Network Interface
VC	Virtual Circuit / Virtual Connection

5 Conventions

PDU formats: In this recommendations PDU headers are structured as a group of four octets or words numbered from the left from 0 to 31.

Bit order within single octet and multiple octets fields: In a single octet field, the left most bit of an octet is the high order or most significant bit. Similarly, in a multi-octet field the left most bit of the whole field is the most significant bit.

6 Architecture

6.1 General

The reference model for frame relay services over MPLS core networks is shown in Figure 1/X.84. It consists of the following elements:

- An MPLS core network,
- Provider Edge (PE) devices providing interworking functions between frame relay and MPLS. PEs can support frame relay UNIs and/or NNIs,
- Frame relay devices (DTEs/CEs) and FR networks (DCEs) connected to PEs with frame relay UNIs and/or NNIs.

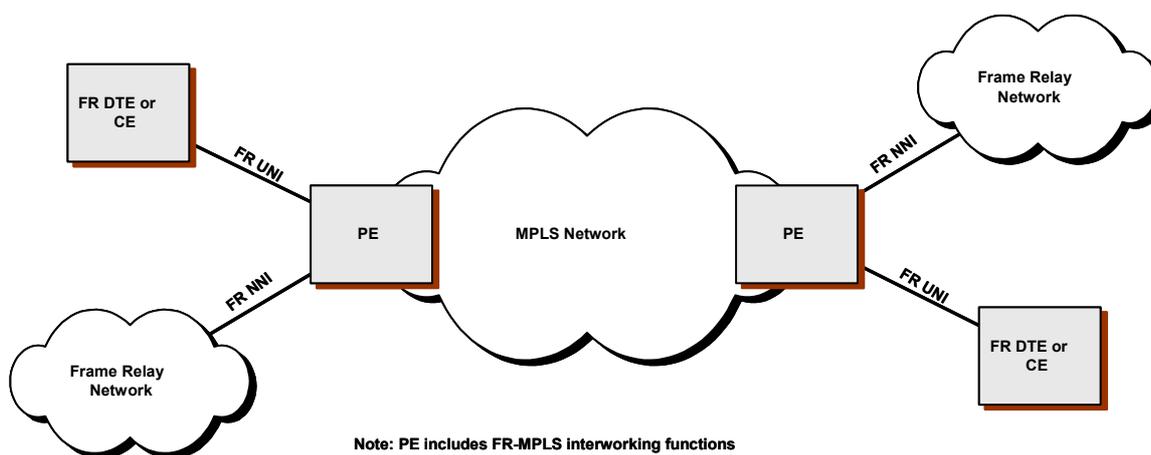


Figure 1/X.84 - Frame Relay over MPLS core network reference model

Frame relay over MPLS core network architecture allows the interconnection of frame relay networks (DCEs) and/or frame relay devices (DTEs) over an MPLS network. In this architecture, frame relay networks and DTEs act as CE devices attached to PEs of the MPLS network as shown

in Figure 1/X.84. The frame relay service is first provisioned between each frame relay DTE or DCE and the corresponding PE device. A Virtual Connection Label Switched Path (VC LSP) is then set up between the two PEs to complete the frame relay Virtual Connection (VC).

The use of the MPLS network by two frame relay networks and/or devices is not visible to the end users. The end user protocol stacks remain intact. The PE provides all mapping and encapsulation functions necessary to ensure that the service provided to the frame relay networks and/or devices is unchanged by the presence of an MPLS transport network.

6.2 MPLS Tunnel and VC LSPs

MPLS Label Switched Paths (LSPs) called “Tunnel LSPs” connect a pair of PEs. Several “Virtual connection LSPs” (VC LSPs) may be nested inside one Tunnel LSP (see Figure 2/X.84). Each VC LSP carries the traffic of a frame relay Permanent Virtual Connection (PVC) or Switched Virtual Connection (SVC) in one direction. Since LSPs are uni-directional, a pair of VC LSPs carrying traffic in opposite directions will usually be required for each frame relay PVC or SVC.

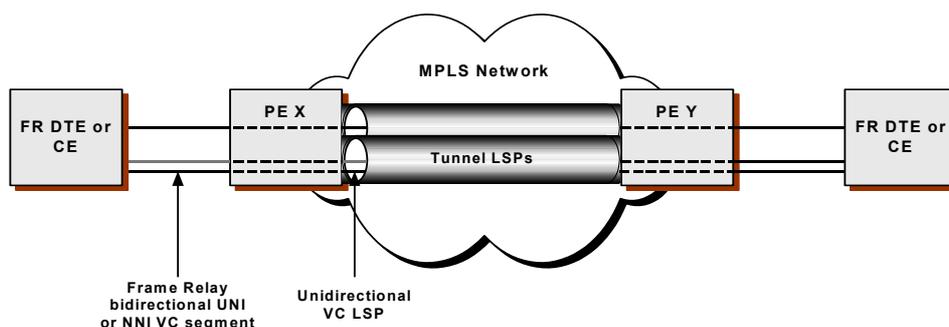


Figure 2/X.84 – Tunnel and VC LSPs

In the case of tunnel LSPs, one of the tunnel LSP transports the FR-MPLS packets, for example, from PE X to PE Y and the other tunnel LSP transports FR-MPLS packets in the opposite direction. The corresponding label does not tell PE Y about a frame relay virtual connection. If Penultimate Hop Popping (PHP) is used as per RFC 3031, PE Y will never see the tunnel label. If PE X itself is the penultimate hop, a tunnel label will not get pushed on. In this example, PE X is the ingress PE device and PE Y the egress PE device.

In the case when PE X has to send a frame relay frame to PE Y, it first pushes a VC label onto its label stack, and then pushes on a tunnel label. The VC label is not looked at until the FR-MPLS packet reaches PE Y. PE Y forwards the packet based on the VC label. The “VC label” identifies the frame relay virtual connection.

In general, the VC label must always be at the bottom of the label stack, and the tunnel label, if present, must be immediately above the VC label. As the packet is transported across the MPLS

network, additional labels may be pushed on (and then popped off) as needed. If PE X and PE Y are directly adjacent MPLS nodes, then it may not be necessary to use a tunnel label at all.

6.3 Relationship Between Frame Relay VC and MPLS VC LSP

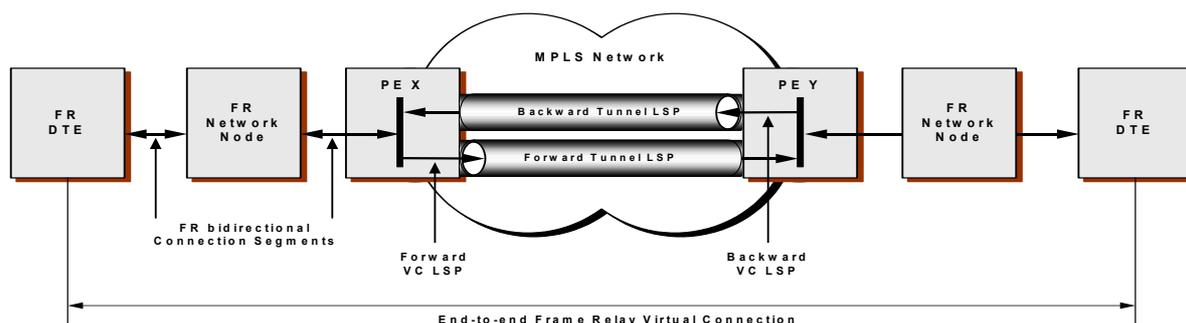
Frame relay VCs are considered to be bi-directional objects mainly because of the way they are created and identified. A single frame relay Data Link Connection Identifier (DLCI) refers to the two directions of a frame relay VC. Frame relay signalling establishes the two directions simultaneously. But in general, each direction of a frame relay VC may have different traffic and Quality of Service (QoS) characteristics. The resource management of a frame relay implementation treats each direction separately and independently. MPLS LSPs on the other hand are uni-directional and are established separately. The VC LSP in each direction must comply with the characteristics of the corresponding direction of the frame relay VC. Frame relay and MPLS interworking requires that frame relay VC segments interwork with a pair of MPLS LSPs.

In general, a frame relay VC consists of several segments: one at each UNI, one or more segments inside a frame relay network or between frame relay networks and in the case of FR-MPLS interworking there is a pair of VC LSP between PEs.

During the creation of a frame relay VC, a pair of VC LSPs will have to be established between two PEs. Figure 3/X.84 illustrates the relationship between a Tunnel LSP, a VC LSP and frame relay VCs. For one end-to-end frame relay VC, two VC LSPs exist; the X-to-Y LSP transports the traffic from PE X to PE Y and the Y-to-X LSP transports the traffic in the opposite direction. Multiple VC LSPs may be nested inside one Tunnel LSP.

The X-to-Y VC LSP is the transmit VC LSP for PE X and the receive LSP for PE Y. Likewise, the Y-to-X LSP is the transmit LSP for Y and the receive LSP for X.

In the frame relay domain, a DLCI identifies a frame relay VC and in the MPLS domain, VC LSP labels with possibly different values identify the pair of VC LSPs, one label value for each LSP.



Note: The PEs contain frame relay MPLS interworking functions.

Figure 3/X.84- Frame relay end-to-end VC and MPLS LSP

In the PE X-to-PE Y direction a tunnel LSP transports FR-MPLS packets from PE X to PE Y, and the corresponding label is not related to any frame relay VC.

6.4 Frame Relay over MPLS mapping modes

Two mapping modes are defined for frame relay over MPLS networks:

1. The first one is called "one-to-one" mapping mode. With one-to-one mapping mode, for each frame relay VC, a pair of MPLS LSPs (one for each direction of the traffic) is established between a pair of PEs as described in the previous sub-clause.
2. The second mapping mode is called "many-to-one mode". With this mapping mode multiple frame relay VCs between a pair of frame relay devices (CEs) controlled by the same frame relay signalling VC are assigned to one pair of MPLS LSPs. In this mode, all the frame relay VCs (including the signalling channel assigned to DLCI = 0) are transparently transported between PEs.

As specified in clauses 8, 9 and 12, the encapsulation of frame relay information is different for the two mapping modes.

7 Frame Relay over MPLS Requirements

This clause lists the frame relay requirements to be met in a frame relay over MPLS network configuration.

1. Frame Transport: must have the ability to carry both user-data frames and network management (e.g., OAM) frames on the same VC LSP
2. Frame Length: must transport variable length frame relay frames without being limited by MPLS network Maximum Transmission Unit (MTU) size.
3. VC Mapping: must support a 1:2 correspondence between a frame relay VC segment and a pair of VC LSPs because frame relay VCs are bi-directional entities and MPLS LSPs are unidirectional entities.
4. Frame Ordering: must deliver frame relay frames in the same order as they were transmitted. Correct ordering will be achieved by using either sequence numbers or other means,
5. Control bits: must support the transport of the DE (Discard Eligibility), FECN (Forward Explicit Congestion Notification), BECN (Backward Explicit Congestion Notification) and C/R (Command / Response indicator) bits.
6. PVC Status signalling: must support the mapping and transport of PVC active and inactive indications. The support of continuity check should be provided. Note PVC status signalling is for further study.
7. Traffic Management: should be able to map the following frame relay traffic management parameters to VC LSP parameters:
 - a) Committed Information Rate (CIR) or throughput,
 - b) Committed Burst size (B_c),

- c) Excess Burst size (B_e),
- d) Maximum frame size.

Must support the frame relay service characteristics as defined by frame relay traffic parameters with the establishment of suitably engineered Tunnel LSPs.

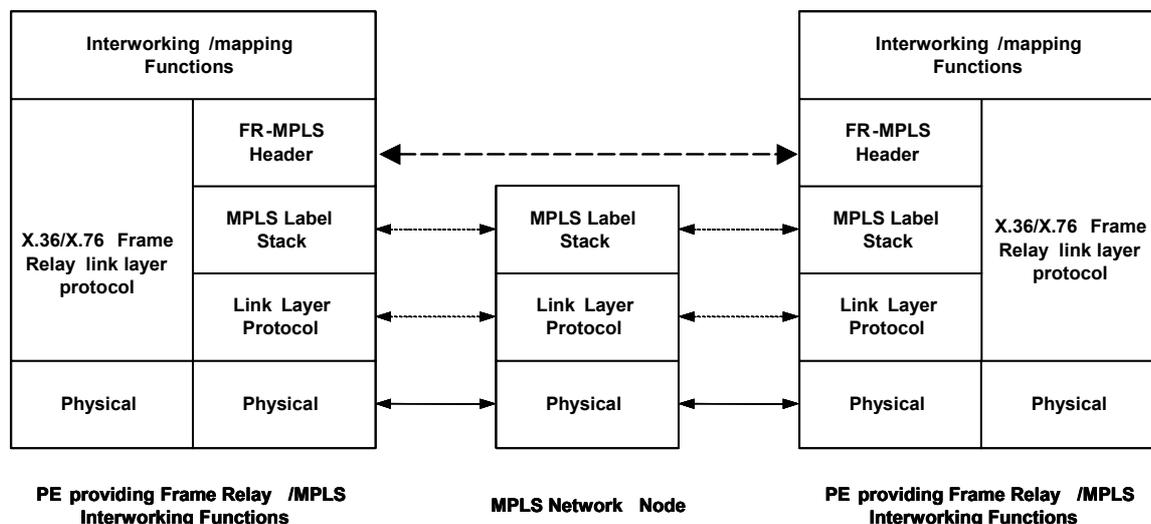
- 8. Frame priority and QoS: should support the ability to map different frame relay priorities or QoS classes to appropriately engineered Tunnels and VC LSPs.

8 Protocol Stack and Frame Format

8.1 Data Transfer Protocol Stack

The MPLS side of a PE consists of multiple protocol layers as shown in Figure 4/X.84. Above the physical layer is the link layer protocol. Above the link layer is the MPLS header encapsulation and processing layer performing MPLS functions, for example, label stack processing as defined in RFC 3032 and traffic management and queuing. The MPLS processing layer interacts with the MPLS network. The FR-MPLS header is exchanged between two PEs.

The frame relay side of a PE consists of the physical layer and the frame relay link layer protocol. Both frame relay UNIs and NNIs are supported as specified in Recommendations X.36 and X.76. The interworking or mapping functions performs the actions necessary to be able to transfer frames from the frame relay side to the MPLS side and vice versa.



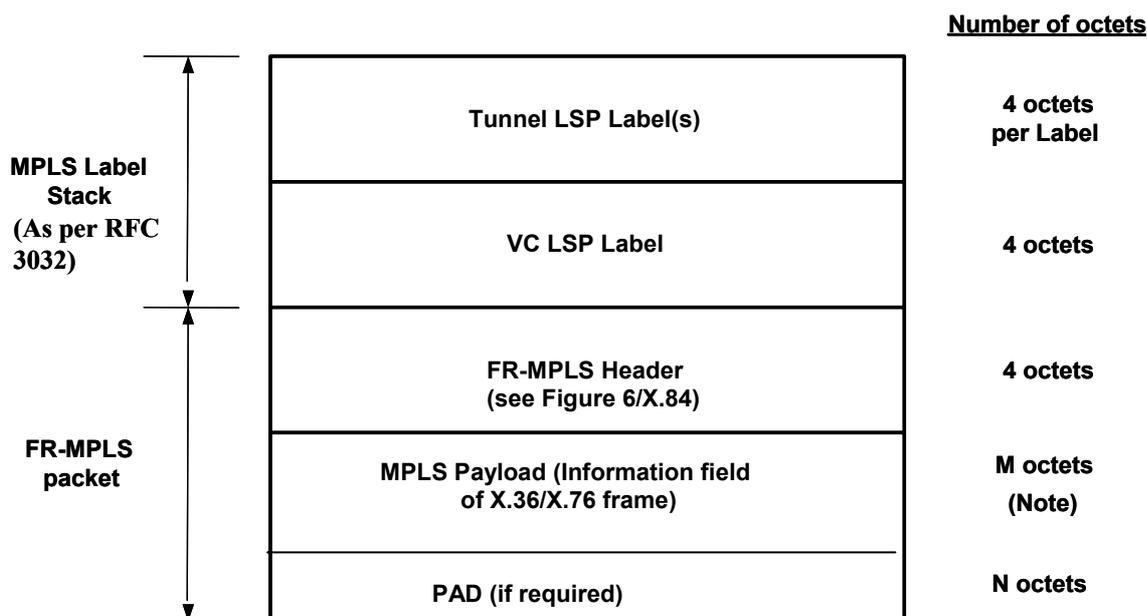
Note: Recommendations X.36 and X.76 supersedes Recommendation Q.922 and contain the same frame relay link layer protocol as defined in Recommendation Q.922 link core service.

Figure 4/X.84 – Data transfer protocol stacks

8.2 FR-MPLS Packet Format for the one-to-one mapping mode

Figure 5/X.84 shows FR-MPLS packet format for the one-to-one mapping mode. The FR-MPLS packet consists of a FR-MPLS header followed by the payload field and a padding field if required.

The payload may contain user or network data. In the case of user's data, the payload field contains the frame relay information field. An MPLS LSP label stack as defined in RFC 3032 precedes the FR-MPLS packet. The MPLS label stack and FR-MPLS packets are encapsulated in a link layer frame. The specification of the link layer protocol is outside of the scope of this Recommendation. It can be any link layer used between a PE and the MPLS network.



Note: The default number of octets in the payload field (X.36/X.76 frame relay information field only) is defined in Recommendations X.36 and X.76. If it is greater than the MTU, fragmentation may be used.

Figure 5/X.84 – Format of FR-MPLS packet for the one to one mode

The meaning of the fields of the FR-MPLS packet (Figure 5/X.84) for the one-to-one mapping mode is as follows:

8.2.1 Tunnel LSP label(s)

The Tunnel LSP label(s) is (are) used by MPLS network nodes to forward a FR-MPLS packet from one PE to the other. Since MPLS LSPs are uni-directional, a pair of Tunnel LSPs carrying traffic in opposite directions is required to create a bi-directional transport. A Tunnel LSP label is a standard MPLS label as defined in RFC 3032. There can be more than one tunnel LSP since the MPLS network can use more than one tunnel as per RFC 3031.

The S bit shall be set to 0 indicating that this is not the bottom of the label stack.

The setting of the EXP and TTL fields of the tunnel label(s) is outside the scope of this Recommendation.

8.2.2 VC LSP Label

The VC LSP label identifies one LSP assigned to a FR VC in one direction. Together the Tunnel LSP label(s) and VC LSP label form an MPLS label stack. More than one VC LSP may be supported by one MPLS tunnel LSP. A VC LSP label is a standard MPLS label as defined in RFC 3032.

Since MPLS LSP is unidirectional, for the case of bi-directional FR VCs, there will be two different VC LSPs, one for each direction of the connection. These may have different label values.

The S bit shall be set to 1, to indicate that it is the bottom of the stack.

The TTL value in the VC label and the setting of the EXP bits are for further study.

8.2.3 FR-MPLS header

FR-MPLS header contains protocol control information. Its structure is shown in Figure 6/X.84.

The FR-MPLS header is shown in Figure 6/X.84.

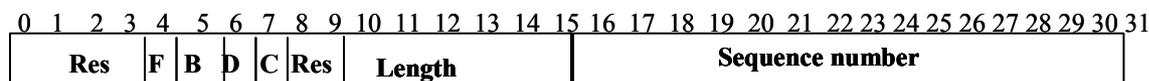


Figure 6/X.84 – FR-MPLS header structure for the one to one mode

The meaning of the fields of the FR-MPLS packet header (Figure 6/X.84) is as follows:

Res (bits 0 to 3):

Reserved bits. They are set to zero on transmission and ignored on reception.

F (bit 4):

FR FECN (Forward Explicit Congestion Notification) bit.

B (bit 5):

FR BECN (Backward Explicit Congestion Notification) bit.

D (bit 6):

FR DE bit indicates the discard eligibility.

C (bit 7):

FR frame C/R (command/response) bit.

Res (bits 8 and 9):

Reserved for optional fragmentation and re-assembly procedures. When fragmentation and re-assembly procedures are not supported, they are set to zero on transmission and ignored on reception.

Length (bits 10 to 15):

The length field is used in conjunction with the padding of short FR-MPLS packets when the link layer protocol (a notable example is Ethernet) requires a minimum frame length.

If the total length of a FR-MPLS packet is less than 64 bytes, padding has to be performed.

When padding of a short FR-MPLS packet is performed the length field is the sum of the lengths of the following fields of the FR-MPLS packet (Figure 5/X.84) specified in bytes: FR-MPLS Header, Payload. Otherwise the length field must be set to zero. The value of the length field, if non-zero, is used to remove the padding characters by the ingress PE.

Sequence number (Bit 16 to 31):

Sequence numbers provide one possible mechanism to ensure the ordered delivery of FR-MPLS packets. The sequence number field specifies the sequence number of a FR-MPLS packet. A circular list of sequence numbers is used. A sequence number takes a value from 1 to 65535 ($2^{16}-1$). **Successive sequence numbers are generated by adding 1 to the previous sequence number, up to a maximum value of 65535. The next sequence number after 65535 is 1.** The value of zero indicates that the sequence number field is not used.

Note: The use of sequence numbers may not be required if the network itself ensures ordered delivery of FR-MPLS packets.

8.2.4 Payload

The payload field corresponds to frame relay frame information field as defined in Recommendations X.36 and X.76 with bit/octet stuffing removed. The default for the number of bytes in an information field is 262 octets. Recommendations X.36 and X.76 recommend to support a size of at least 1600 octets. The maximum length of the payload field should be agreed by the two PEs when the VC LSP is established.

Note: Frame relay frame opening and closing flags, the address and the FCS fields are not included in the payload.

8.2.5 Pad

The Pad consists of a number of characters (including zero character) to bring the FR over MPLS packet size to the minimum size as required by the underlying link layer protocol, in particular IEEE 802.3/Ethernet. Any 8-bit character with a decimal value from 0 to 255 may be used as a padding character.

Note: The limit of 64 octets for the packet length below which padding is required, aligns with the minimum Ethernet frame size.

9 FR-MPLS packet processing for one-to-one mapping mode

9.1 Generating FR-MPLS packets

The generation process of a FR-MPLS packet is initiated when a PE receives a frame relay frame from one of its frame relay UNI or NNI. The PE takes the following actions (not necessarily in the order shown):

- It generates the following fields of the FR-MPLS header from the corresponding fields of the frame relay frame as follows:
 - Command/Response (C/R or C) bit: The C bit is copied unchanged in the FR-MPLS header.

- Discard eligibility indicator (DE or D): The D bit is set as follows in the FR-MPLS header: This bit, if used, is set to 1 to indicate a request that a frame should be discarded in preference to other frames in a congestion situation.
 - Setting of the D bit by a PE is optional. However, no PE shall clear this bit (set it to 0 if it was received with the value of 1). A PE that does not provide discard eligibility notification shall pass this bit unchanged. Networks are not constrained to discard only frames with D = 1 in the presence of congestion.
 - Forward explicit congestion notification (FECN or F bit): FECN may be set by a congested PE to notify the user that congestion avoidance procedures should be initiated where applicable for traffic in the direction of the FR-MPLS packet carrying the FECN.
 - The F bit is set to 1 to indicate to the destination that the frames it received have encountered congested resources. This bit may be used by a destination to adjust its transmission rate.
 - While setting the F bit by a PE is optional, no PE shall clear this bit (set it to 0 if it was received with the value of 1). PEs that do not provide FECN shall pass this bit unchanged.
 - Backward explicit congestion notification (BECN or B bit): BECN follows the same processing rules as FECN, except that it applies to the opposite direction.
 - Length: If the packet's length (defined as the length of the frame relay frame information field plus the length of the FR-MPLS header) is less than 64 bytes, the length field MUST be set to the packet's length. Otherwise the length field MUST be set to zero. The value of the length field, if non-zero, can be used to remove any padding, see subclause 8.2.5.
 - Sequence number: See sub-clause 9.1.1.
- It processes the payload and padding fields as follows: The payload of the FR-MPLS packet is the contents of frame relay frame information field stripped from any bit or octet stuffing. FCS is removed prior to MPLS encapsulation. Padding characters may follow the payload field if required by the link layer protocol to bring the FR-MPLS packet to a minimum length.

Additional processing is performed by the lower protocol layers in order to transmit the FR-MPLS packet to its next destination.

9.1.1 Setting the sequence number

The sequence number field is set according to whether the sequence number is used or not.

If the ingress PE supports the sequence number capability then the following procedure to number FR-MPLS packets is used:

- The initial FR-MPLS packet transmitted MUST use sequence number 1,
- For a subsequent frame, the sequence number corresponds to the sequence number of the preceding frame incremented by 1 up to the maximum value of 65535,
-
- When the sequence number reaches the maximum 16 bit value (65535) the next sequence number wraps to 1 (the value of 0 is skipped).

If the PEs do not support sequence number processing, then the sequence number field must be set to 0.

9.2 Reception of FR-MPLS packets

When an egress PE receives a FR-MPLS packet, it processes the different fields of the FR-MPLS header in order to synthesize a new frame relay frame for transmission to a CE on a FR UNI or NNI. The PE performs the following actions (not necessarily in the order shown):

- It generates the following FR frame header fields from the corresponding fields of the FR-MPLS packet as follows:

- C/R bit is copied unchanged in the frame relay header.
- D bit is copied as follows into the frame relay header DE bit: If it was set to one in the incoming FR-MPLS packet, it must be copied unchanged in the FR frame header or, depending on the traffic policing performed by the PE and its state of congestion, the FR-MPLS packet may be dropped.

Otherwise if the D bit was set to zero, it may be set to zero or one, depending on the traffic policing performed by the PE device. Setting of this bit by a PE is optional.

- The F bit is copied as follows in the frame relay header FECN bit: If it was set to one in the incoming FR-MPLS packet, it must be copied unchanged in the frame relay header. Otherwise it was set to zero.

The F bit may be set to zero or one, depending on the congestion state of the PE device in the forward direction. Setting this bit by a PE is optional, if the PE does not support FECN, it shall pass this bit unchanged.

- BECN follows the same processing rules as FECN, except that it applies to the opposite direction.
- It processes the length and sequence field, the details are in the subsequent sub-sections.

- It regenerates the frame relay information field from the contents of the FR-MPLS packet payload after removing any padding character and retrieves the appropriate DLCI.

Once the above fields of a FR frame have been generated, the FCS has to be computed, HDLC flags have to be added and any bit or byte stuffing has been performed. The FR frame is queued for transmission on the selected frame relay UNI or NNI.

9.2.1 Checking the sequence number by the egress PE

When a FR-MPLS packet is received the sequence number is processed as follows:

- If the sequence number of the packet is 0, then the packet passes the sequence number check.

Note: A sequence number equal to 0 means that sequence numbers are not used. Through management or signalling, the two PEs determine whether sequence numbers are used or not.

- Otherwise if the packet sequence number \geq the expected sequence number and the packet sequence number - the expected sequence number < 32768 , then the packet is in order.
- Otherwise if the packet sequence number $<$ the expected sequence number and the expected sequence number - the packet sequence number ≥ 32768 , then the packet is in order.

- Otherwise the packet is out of order.

If the packet is in order, then it passes the sequence number check and the expected sequence number is set as per the following assignment:

expected_sequence_number := packet_sequence_number + 1, up to a maximum of 65535 after which it skips zero and wraps from 65535 to 1.

if (expected_sequence_number = 0) then expected_sequence_number := 1;

FR-MPLS packets, that are received out of order should be dropped unless they can be re-ordered without introducing significant delays.

If an egress PE receives an excessive number of out-of-sequence FR-MPLS packets, it should inform the management plane responsible for FR-MPLS interworking function in the PE and take the appropriate actions. The threshold for declaring that out-of-sequence FR-MPLS packets are excessive is not defined in this Recommendation.

Note: Excessive out-of-sequence FR-MPLS packets has an effect similar to excessive packet loss on a frame relay VC.

9.2.2 Processing of the length field by the receiver

Any padding character, if present, in the payload field of a FR-MPLS packet received must be removed before forwarding the data to the next destination.

The procedure described here is used to remove padding characters.

If Length_field is set to zero then there are no padding characters following the payload field.

Else padding characters are included and their length is computed as follows:

The length of FR-MPLS packet consists of FR-MPLS header and payload. It does not include MPLS label stack. Hence, Padding-length = Length of FR-MPLS packet - Length_field;

After computing the length of the padding field, Padding-length characters are removed from the end of the FR-MPLS packet.

9.3 Handling of error conditions

If a PE receives a FR-MPLS packet with a header containing invalid contents, it shall be discarded. For example:

- An invalid or unassigned tunnel or VC label,
- A value different than zero for the first four bits of the FR-MPLS header,
- A length field with a contents greater or equal to 64,
- If fragmentation is not used a value different than zero for the two bits reserved for fragmentation (bits 8 and 9 of the FR-MPLS header).

Note: Link layer errors are not addressed in this Recommendation.

9.4 Optional fragmentation and re-assembly procedures

A FR-MPLS packet payload is normally relayed across a VC LSP as a single PDU. However, there are cases where the combined size of the payload and its associated headers may exceed the network path Maximum Transmission Unit (MTU). When a packet exceeds the MTU of a given network, fragmentation and reassembly will allow the packet to traverse the network and reach its intended destination.

The fragmentation procedures for the one-to-one mode are optional and may be negotiated between the local and remote PEs through signalling or may be provisioned at both the PEs. These procedures do not apply to many-to-one mode.

Fragmentation and reassembly in network equipment generally requires significantly more resources than sending a packet as a single unit. As such, fragmentation and reassembly should be avoided whenever possible. Solutions for avoiding fragmentation include:

- Proper configuration and management of MTU sizes between the CE, PE and across the MPLS network,
- Adaptive measures which operate with the originating CE to reduce the packet sizes at the source as defined in RFC 1191 and 1981.

9.4.1 One-to-one mapping mode with fragmentation

The fragmentation procedure for the one-to-one mapping mode use the existing sequence number and two reserved bits (bits 8 and 9) in the length field as control bits.

Figure 7/X.84 shows FR-MPLS header structure for one-to-one mapping mode with fragmentation. Description of the control bits for segmentation (bit 8 and 9 of length field) are given below:

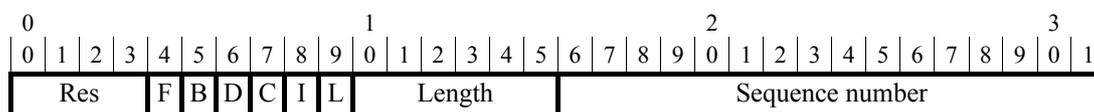


Figure 7/X.84 – FR-MPLS header structure with fragmentation

Initial fragment bit (I):

The initial (I) fragment bit is set to zero on the first fragment derived from an frame relay frame and set to one for all other fragments derived from the same frame.

Last fragment bit (L):

The last (L) fragment bit is set to zero on the last fragment and set to one in all the other fragments.

A packet containing a complete non-fragmented frame relay frame has both the I and L bits set to zero. PEs that do not support the fragmentation functionality should set both the I and L bits to zero.

Sequence Number:

PEs that support one-to-one mode fragmentation support sequence numbers as described in Clauses 9.1.1 and 9.2.1. The sequence number feature is used by the egress PE to detect out of order or lost fragments.

Since the value zero indicates that the sequence number is not in use, its use for fragmentation must follow the same rule, as the sequence number is incremented. It skips zero and wraps from 65535 to 1.

Fragmentation takes place in the egress PE and reassembly takes place in the ingress PE after the reception of a FR-MPLS packet fragment.

A series of data fragments is created by fragmenting a frame relay payload (Information field) into multiple fragments. The resulting fragments must be transmitted in the same sequence as they occurred in the frame prior to being fragmented.

Every fragment in the series contains the Frame Relay congestion bits (F, B, D) and C in the control bits of the FR-MPLS header.

The first fragment sent on a VC (following a VC becoming active) may have the sequence number set to any value (other than zero), and the sequence number must subsequently be incremented by one for each fragment sent.

A PE supporting fragmentation of FR frame encapsulated packets should support the following procedures.

- The FR-MPLS packet carrying the first fragment of a frame relay frame should have the Initial Fragment bit (I) set to zero. For the subsequent fragments the (I) bit will be set to one.
- The FR-MPLS packet carrying the last fragment of a FR frame should have the Last fragment bit (L) set to zero. For the packets carrying the other fragments of the same FR frame, the L bit will be set to one.
- The FR-MPLS packet carrying a complete non-fragmented FR frame will have both I and L bits set to zero.
- The FR-MPLS packet carrying neither the Initial nor the last fragment will have both I and L bits set one.

9.4.3 Reassembly procedures

For each VC LSP, the receiver must keep track of the incoming sequence numbers and maintain the most recently received sequence number. The receiver detects the end of a reassembled frame when it receives a fragment bearing the (L)ast bit set to "0". Reassembly of the frame is complete if all sequence numbers up to that fragment have been received.

Note that the Frame Relay congestion bits (F, B, D) must be logically ORed for all fragments, and the results included in the reassembled frame.

The receiver detects lost fragments when one or more sequence numbers are skipped. When a lost fragment or fragments are detected on a VC, the receiver must discard currently unassembled and subsequently received fragments for that VC until it receives the first fragment that bears the (I) initial bit set to "0". The fragment bearing the (I) initial bit set to "0" is used to begin accumulating a new frame.

In the event of an error (e.g., one or more fragments lost due to transmission error or reassembly buffer overflow), fragments which cannot be reconstructed back into the original frame must be discarded by the receiver.

10 FR PVC provisioning

Provisioning of FR PVCs requires the following actions: The PEs and CEs are configured independently for each FR UNI or NNI PVC segments. Some of the FR PVC configuration parameters may include:

- Outgoing and incoming throughputs (CIR),
- Outgoing and incoming committed burst sizes (Bc),
- Outgoing and incoming excess burst sizes (Be),
- Outgoing and incoming maximum frame lengths,
- The DLCI assigned to the FR PVC locally,
- If used, FR transfer and discard priority class or FR service class assigned to the FR VC.

The establishment of a frame relay VC across an MPLS network requires, a pair of VC LSPs to be established between the two PEs as described in clause 6.3. The traffic capacity of the VC LSP must accommodate the traffic and QoS requirements of the FR PVC as defined by the FR PVC configuration parameters.

11 Traffic management aspects

In Recommendations X.36 and X.146, a number of different traffic parameters and Quality of Service (QoS) classes are defined.

When a tunnel LSP is used to carry multiple frame relay VCs with different combinations of traffic parameters and QoS classes, the tunnel LSP shall be capable of providing the required QoS for all frame relay VCs. In a MPLS network that does not support QoS differentiation on a per packet

basis, the tunnel LSP shall meet the most stringent QoS requirements of the frame relay VCs carried by the tunnel LSP.

Use of Differentiated Services for Frame relay over MPLS:

If the MPLS network supports Differentiated Services (DiffServ) Behaviour Aggregates defined in informational RFCs 2475 and 3260, MPLS packets can be treated with different priorities on a Per Hop Behaviour (PHB). In this case, two different types of LSPs are defined in RFC 3270 which can both be used for the tunnel LSP:

- Label-Only-Inferred-PSC LSPs (L-LSP);
- EXP-Inferred-PSC LSPs (E-LSP).

If a L-LSP is used as a tunnel LSP, the PHB scheduling class of each packet is inferred from the label without any other information (e.g., regardless of the EXP field value). In that case, the LSP shall meet the most stringent QoS requirements of the frame relay VCs tunneled by the LSP.

If an E-LSP is used as a tunnel LSP, the EXP field of the tunnel label is used to determine the PHB to be applied to each packet, i.e., different packets in one LSP may receive a different QoS. The 3-bit EXP field of the tunnel label can represent eight different combinations of Per Hop Behaviour (PHB) and drop precedence levels. The mapping of the PHB to EXP fields is either explicitly signalled at label set-up or relies on a pre-configured mapping.

12 Frame relay many-to-one mapping mode

12.1 General

Frame relay many-to-one mapping mode is a network option.

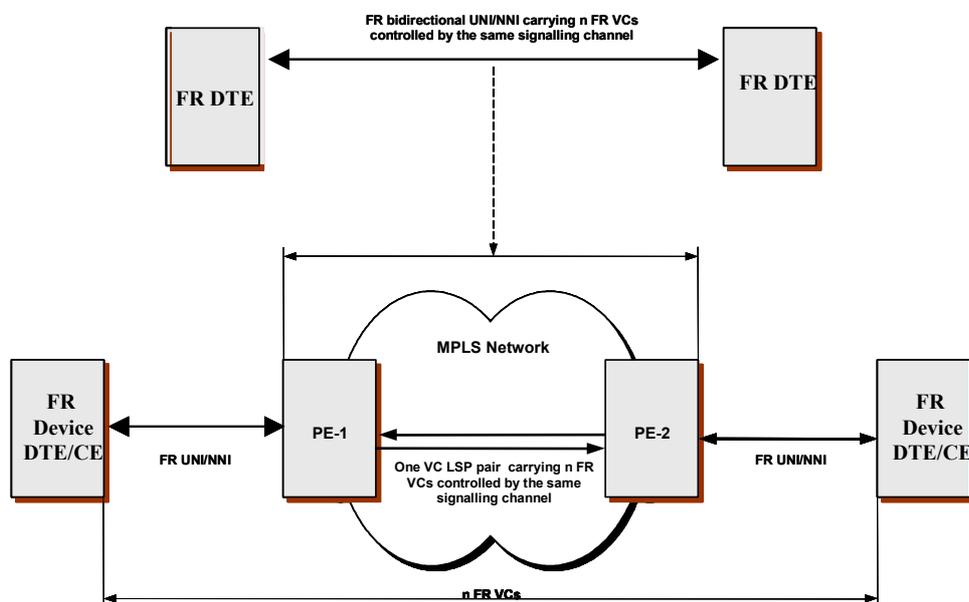


Figure - 9/X.84 illustrates the concept of frame relay many-to-one mapping mode

Note 1: The n FR VCs between the pair of frame relay devices (CEs), aggregated on the same MPLS VC LSP pair are controlled by the same signalling channel.

Note 2: The sequence number processing and the interworking functions are located in the PEs.

Figure 9/X.84 - Concept of frame relay many-to-one mode

Figure 9/X.84 upper part shows two frame relay devices connected via a frame relay UNI or NNI. Between them n frame relay VCs are configured. The n frame relay VCs are controlled by the same signalling channel using DLCI = 0.

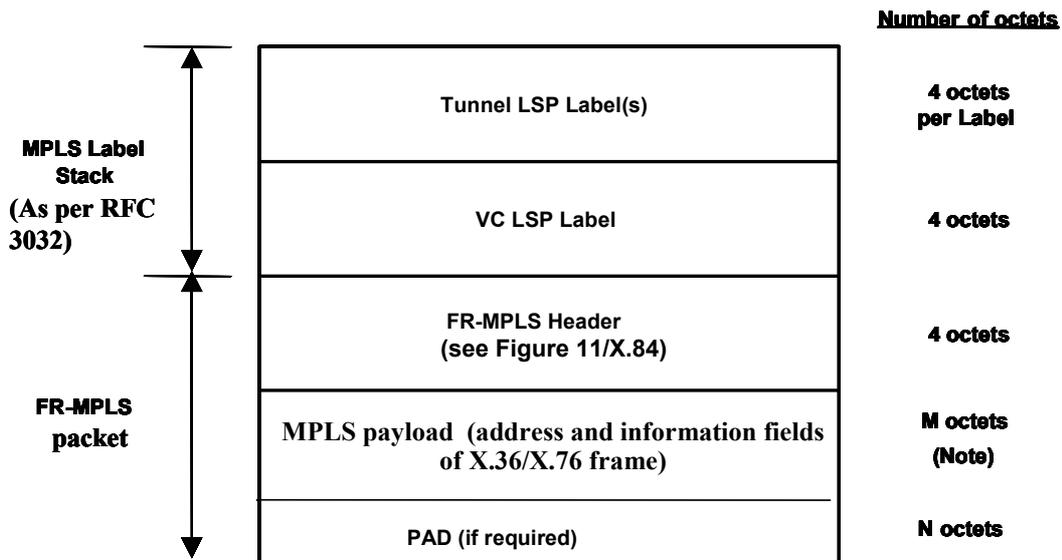
Figure 9/X.84 lower part shows the replacement of the physical frame relay interface with a pair of PEs and a pair of VC LSPs (one VC LSP for each traffic direction between PE1 and PE2). The interface between a FR device and a PE is either a FR UNI or NNI. The set of n FR VCs which are controlled by the same signalling channel using DLCI = 0, between the two PEs are mapped to one pair of VC LSPs. In the many-to-one mapping mode, FR VCs are not visible individually to a PE, there is no configuration of individual FR VC in a PE. A PE processes the set of FR VCs as an aggregate. FR traffic and QoS parameters listed in Clauses 10 and 11 are assigned to the aggregate traffic flowing on an interface between a CE and a PE and not to individual FR VC and policing may be performed on the aggregate.

FR many-to-one mode provides transport between two PEs of a complete FR frame excluding the opening and closing flags and the Frame check sequence (FCS) and with bit/byte stuffing undone.

Note: As an analogy, FR-MPLS many-to-one mapping mode is somewhat similar to a virtual path connection in an ATM network which includes the signalling channel. The intermediate MPLS nodes do not handle any frame relay status monitoring functions.

12.2. Packet format for many-to-one mapping mode

The packet format for many-to-one mode and mapping of frame relay frame is shown in Figure 10/X.84.



Note: The the payload field includes X.36/X.76 frame relay information field and the address field (including the DLCI and control fields).

Figure 10/X.84 – FR-MPLS packet format for the many-to-one mapping mode

The meaning of the fields of the FR-MPLS packet (Figure 10/X.84) for the many-to-one mode is as follows:

Tunnel LSP label(s):

See section 8.2.1.

VC LSP label:

The VC LSP label identifies one LSP assigned to a set of FR VCs controlled by the same frame relay signalling channel. There is a pair of VC LSPs for the two traffic directions. See section 8.2.2.

FR-MPLS header:

FR-MPLS header contains protocol control information. Its structure is shown in Figure 11/X.84. Frame relay control bits (F, B, D and C) for the many to one mode are not encoded in this header and bits 4-7 are set to zero.

The use of the length and sequence number fields is the same as for the one-to-one mode, with the following exceptions:

There is one sequence number counter for the set of FR VCs and not one for each individual FR VC. To compute the FR-MPLS packet size to determine whether padding is needed or not, the length of FR frame is used.

MPLS Payload:

The payload consists of the address field (including the DLCI and the control bits) and information field of a frame relay frame with the opening and closing flags, bit/octet stuffing and FCS removed.

Pad:

The Pad consists of a number of characters (including zero character) to bring the FR-MPLS packet size to the minimum size required by the link layer protocol, in particular IEEE 802.3/Ethernet. Any 8-bit character with a decimal value from 0 to 255 may be used as a padding character.

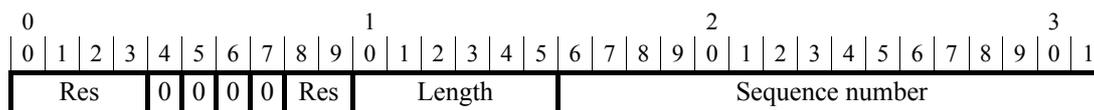


Figure 11/X.84 – FR- MPLS header structure for the many-to-one mode

The two peer PEs must configure the maximum FR frame size such that it can be supported on the LSP MTU size.

12.3. Many-to-one mode processing

When a PE receives a FR frame from a FR device (FR DTE or CE), it shall remove the flags, undo bit/byte stuffing and check the FCS field to determine whether transmission errors occurred or not. If transmission errors occurred, the frame is discarded. Otherwise, the FR frame is encapsulated as a MPLS payload to be forwarded to the remote PE. The FCS and the flags are removed prior to MPLS encapsulation. A PE shall not modify any of the fields in the frame relay frame, they shall be forwarded to the remote PE as received from the FR device.

The processing of the length and sequence number fields is similar to that of Sub-clauses 9.1 and 9.2 with the following difference. There is one sequence number counter for the set of FR VCs and not one counter for each individual FR VC. The sequence numbers checks only the ordered reception of FR-MPLS packets belonging to the aggregated FR VCs.

If an egress PE receives an excessive number of out-of-sequence FR-MPLS packets, it should inform the management plane responsible for FR-MPLS interworking function in the PE and take the appropriate actions. The threshold for declaring that out-of-sequence FR-MPLS packets are excessive is not defined in this Recommendation.

Upon receiving a FR-MPLS packet, the remote PE shall extract the payload field, encapsulate the result in a FR frame (adding flags and FCS) for transmission to the local FR device.

- 1.

APPENDIX I

Example of Fragmentation for the one-to-one mapping mode

An example of the one-to-one mode fragmentation procedure is shown in Figure III.1/X.84. The octets in white indicate the data portion of the original frame that is split into fragments (three fragments in this example). For this example, the starting sequence number of 42 was chosen at random.

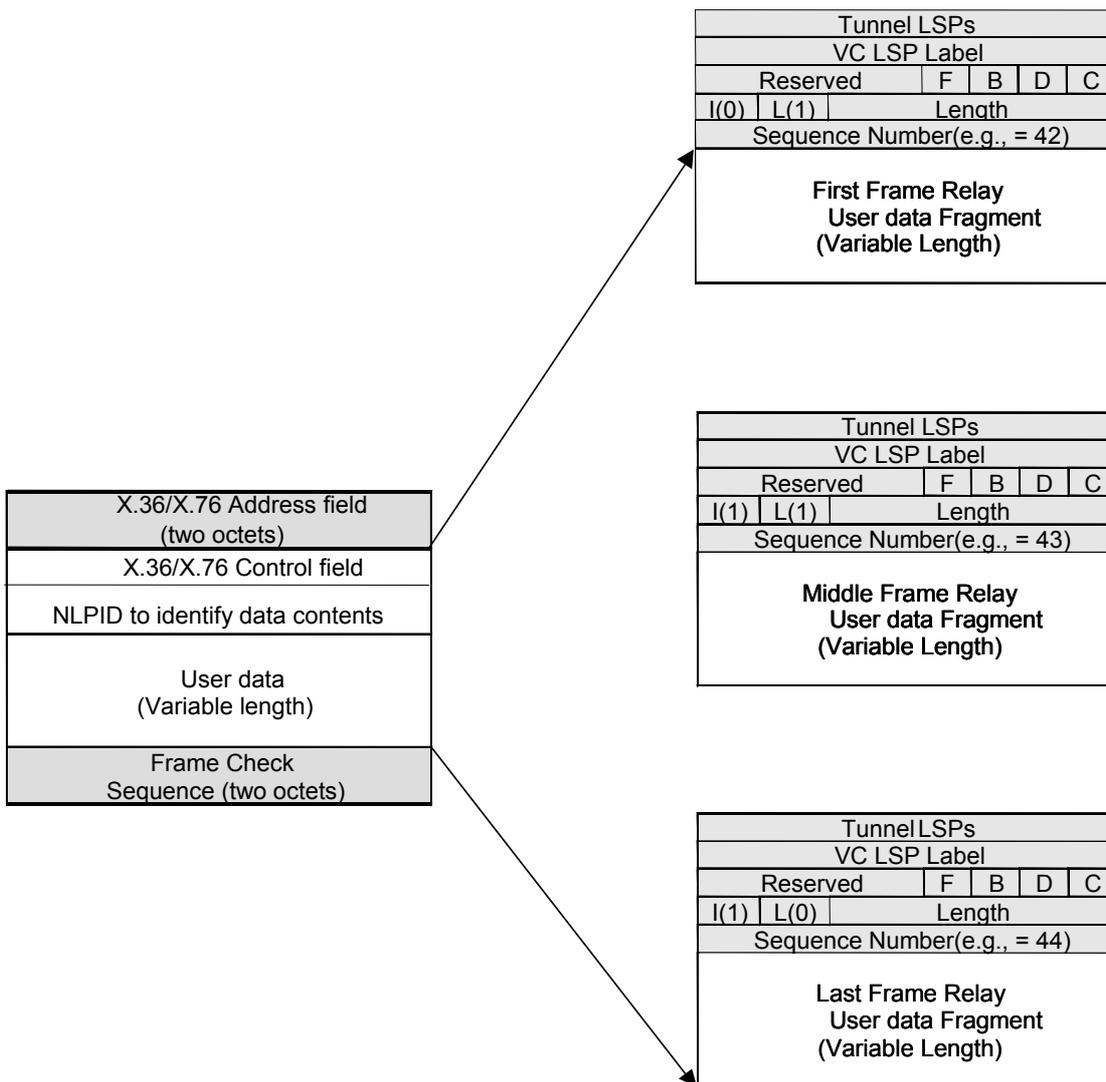


Figure I.1/X.84 – Example of fragmentation for the one to one mapping mode

APPENDIX II

Bibliography

- IETF RFC 1191, Path MTU Discovery, November 1990.
 - IETF RFC 1981, Path MTU Discovery for IP version 6, August 1996.
 - RFCs 2475, An architecture for differentiated services, December 1998.
 - RFC 3260, New terminology and classification for DiffServ. April 2002.
 - ITU-T Recommendation G.805 (2000), Generic functional architecture of transport networks.
 - ITU-T Recommendation M.20 (1992), Maintenance philosophy for telecommunication networks.
 - ITU-T Recommendation Y.1710 (2002), Requirements for Operation & Maintenance functionality for MPLS networks.
 - ITU-T Recommendation Y.1711 (2002), Operation & Maintenance mechanism for MPLS networks.
 - ITU-T Recommendation Y.1720 (2003), Protection switching for MPLS networks.
-