

Internet Engineering Task Force

B. St-Denis, D. Fedyk  
(Nortel Networks)

A. Bhatnagar, A. Siddiqui  
(Cable and Wireless)

R. Hartani, T. So  
(Caspian Networks)

Internet Draft

Document: <[draft-stdenis-ms-over-mpls-00.txt](#)>

Nov 2000

Category: Informational

## Multi-service over MPLS

### Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

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>.

### Abstract

This document describes a generalized approach to carrying Multi-service protocol data units (PDUs) over an MPLS Network. This proposal defines standard MPLS encapsulations that support permanent virtual circuit (PVC) and switch virtual circuit (SVC) networking. The goal of this draft is to provide a framework that allows an MPLS network to support a range of services from simple circuit emulation services, to complete network inter-working. There are two distinct aspects to these services: the data plane and the control plane. The data plane must be defined to support nailed-up and switched services. This architecture covers the data plane but not the complete procedures for the control plane. The control plane is only defined from a nailed up perspective. The control plane for dynamic services maybe supported by other signaling protocols such as PNNI. Non MPLS signaling is not covered by this draft.

## Table of Contents

1. Introduction
2. Requirements
3. Framework Overview
  - 3.1 Transport Label Details
  - 3.2 Service Label Details
  - 3.3 Multi-service Specific layer
4. Service Specifics
  - 4.1 ATM Framework Requirements
    - 4.1.1 ATM Service Specific Encoding
    - 4.1.2 ATM VC Service Specific Encoding
    - 4.1.3 ATM VP Service Specific Encoding
      - 4.1.3.2 ATM VP Service Specific Encoding
  - 4.2 Frame Relay Framework Requirements
    - 4.2.1 Frame Relay Service Specific Encoding
  - 4.3 Circuit Emulation Framework Requirements (CES)
    - 4.3.1 CES Requirements
5. Signaling
  - 5.1 Examples of Signaling Options
  - 5.2 ATM Switched Services
6. Security Considerations
7. References
8. Acknowledgments
9. Author's Addresses

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 [RFC-2119](#) [2].

### 1. Introduction

MPLS has the potential to reduce the complexity of service provider infrastructures by allowing virtually any existing digital service to be supported over a single networking infrastructure. However, many service providers have legacy network and Operational Support System (OSS) capabilities to support these existing service offerings. The benefit of MPLS to this type of network operator is as a common core transport for multiple existing legacy networks, not as a unifying control plane architecture for all services.

MPLS as a common transport infrastructure provides the ability to transparently carry existing networking systems and the evolution to a single networking core. The benefit of this model to the service provider is threefold:

- Leverage of the existing systems and services with increased capacity from an MPLS enabled core.
- Existing network operational processes and procedures used to maintain the legacy services are preserved.
- The common MPLS network infrastructure is used to support both the core capacity requirements of legacy networks and the requirements of new services which are supported natively over the MPLS network.

The requirements for legacy service network operation over MPLS are more comprehensive than simple PDU encapsulation. This is a different approach than in some earlier drafts such as [1]. In addition to the service attributes, the existing network operational capabilities must also be preserved across the MPLS infrastructure.

## 2. Requirements

This section defines the architectural framework for support of multiple services on MPLS. The framework drives the requirements to support generic network behaviors required by multi-service networks. However, while it does define requirements, the architectural framework also allows independence in the MPLS implementation used to support specific legacy services or networks.

The following requirements describe the set of information that is required to enable preservation of legacy service and networking attributes across an MPLS infrastructure:

- Payload types MUST be transparent so that none of the service information is lost within the MPLS Network.
- The capabilities of the service MUST be specified by signaling, provisioning, or a combination of the two techniques. By specifying modes and parameters that are expected to be constants during a flow, the requirements for including this information in a data header are relaxed.
- A consistent self-describing format for all services encapsulated in MPLS Multi-protocol format MUST be available. This will allow new formats to be added at any time in the future.
- Segmentation and Re-assembly SHOULD be supported to ensure that services are not limited by the MTU size of the MPLS network. Cell concatenation and frame fragmentation are well understood capabilities, which SHOULD be used to ensure that all service types, may be transported across an MPLS network.
- The MPLS network SHOULD support order preservation. The network SHOULD have the option to ensure that data exits the network in the order in which it arrived, on a per service channel basis. If order preservation is not required, then it can be omitted and the bandwidth and processing can be saved.
- Service endpoints of the MPLS network SHOULD be as flexible as possible, to allow for different behaviors on the edges of the network. For example, a receiver may drop 'out of order' frames when sequencing is enabled or ignore the sequencing information altogether.

## 3. Framework Overview

The architecture for Multi-service over MPLS must provide for a flexible set of label encapsulations to satisfy both aggregation and non-aggregation of services. Services in this context are traditional data services such as ATM and Frame Relay. Service connections are instances of a session over one of these services such as ATM VCs and Frame Relay connections. The intent is to extend MPLS for services other than IP. In order to facilitate this, a flexible label-stacking format has been adopted. The format is based on standard MPLS label(s) and a service specific shim. To describe this scheme two categories of labels have been defined, which correspond to two layers Service labels and Transport labels.

Figure 1 illustrates the generic encapsulation format. Note that this format uses a service specific "shim" between the MPLS label and the payload so it can work on any MPLS network.

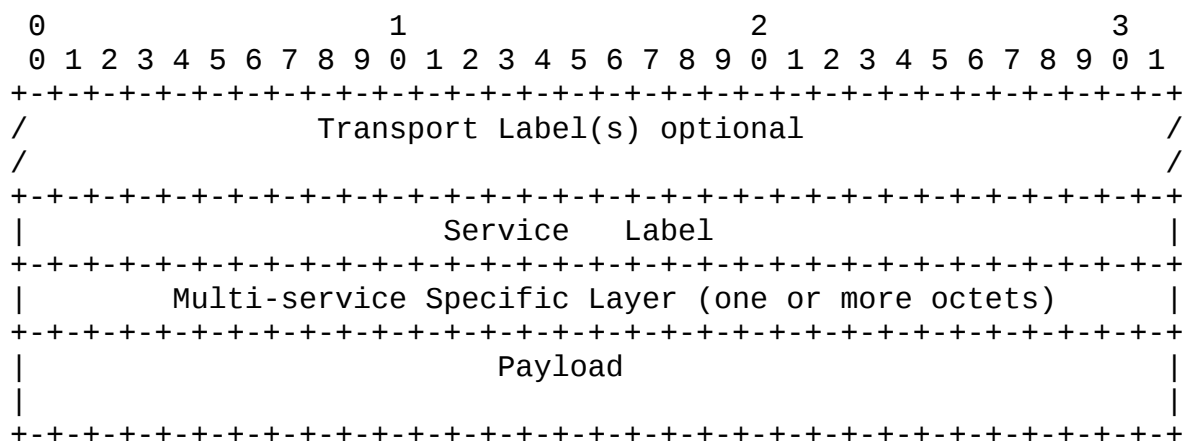


Figure 1 General Encapsulation format

The Service layer, represented by Service labels, encapsulates single service instances. Typically, they require provisioning or signaling or a combination of both to establish the LSP. A service may use a Service Label as the only label encapsulation in the non-aggregated mode.

The Transport Layer, represented by Transport labels, provides aggregation of several service instances between end points of a transport tunnel. A transport label is the outer stacked label(s) that are used to forward to the destination. Transport labels are optional; in other words, point to point non-aggregate Service labels can also be used.

Another term that is useful for categorizing the functions of Multi-service switches is the concept of a label edge router (LER). A Multi-service LER has to support the functions of the service and the conversion to MPLS encapsulation. Any label switch router (LSR) can transport the packets once they have been encapsulated. This is

analogous to IP where an IP switch can become an IP LER so it follows that a Multi-service switch can be a Multi-service LER.




Figure 2 Multi-Service Label Edge Routers

### [3.1 Transport Label Details](#)

Transport labels are simply MPLS labels or label stacks that are part of a transport Label switched path (LSP). These are optional as mentioned earlier. The most common use of a transport label is to aggregate a set of service LSPs, on an explicit path, between a pair of nodes in the network.

A transport label can also be represented as a Forwarding Adjacency (FA) with a set of attributes. The transport label could be advertised to the IGP or it could remain opaque to the IGP and local to the LSP head end node.

Transport labels may be for any flavor of LSP: control driven, point to point, merged or explicitly routed. How the transport labels are constructed is not explicitly specified by the framework. Transport tunnels would likely be engineered when trying to aggregate QoS applications with similar traffic requirements. A transport label may have bandwidth allocation for the whole tunnel and as a result connection admission control may be performed on the tunnel when it is established. Service labels can allocate bandwidth out of this transport tunnels' allocation but how this is achieved is a local issue not covered by the framework.

### [3.2 Service Label Details](#)

Service labels represent labels that have a one to one correspondence with a service connection. Service labels have two modes of operation aggregated and non-aggregated.

Non-aggregated Service labels can be formed from any type of LSP. Non-aggregated Service labels do not have a transport label. The service label has a service connection associated with it, which can be nailed up or configured hop by hop. The service connection may also be configured at the ends and signaled by MPLS (see the signaling section).

Aggregated Service Labels are always associated with co-located Transport labels and consequently may be simpler than non-aggregated

service labels. Aggregated service labels may be configured on the ends points when nailed up over a Transport LSP that extends multiple hops. Aggregated service labels can also have signaling associated with them.

A service label MAY have traffic parameters associated with the LER. If the service label is tunneled in a Transport label, the service label MAY reserve bandwidth from the transport labels' bandwidth allocation. As mentioned for Transport labels, this is outside the scope of the framework.

### 3.3 Multi-service Specific layer

This header contains service specific information. Each service has a defined encoding for this layer. All services share a common format.

This header is a variable length header depending on the service and the options on each LSP.

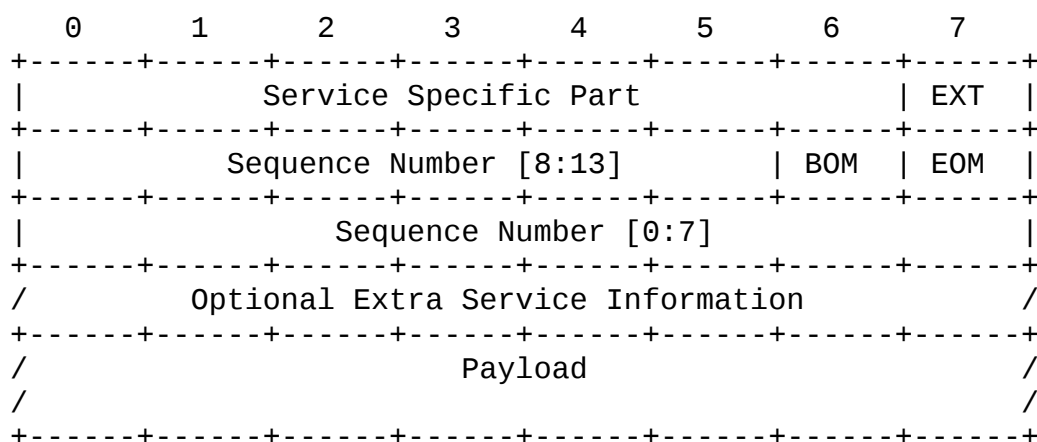


Figure 3 Multi-service Specific Layer

The Multi-service specific information is 8 bits that is specific to each service type. A common bit 7 is an extension bit (EXT) that specifies a 2 octet extension header. The service specific part is specified for each service.

The extension header is used for packet ordering, segmentation, and fragmentation. The extension header contains a beginning of frame marker (BOM) and an end of frame marker (EOM) which are used if this data has been fragmented. If the BOM and EOM are set to one then the frame has not been fragmented.

Sequence numbers are useful for ordering and also useful detecting loss in the network. This is important for circuit emulation types of service that may not have their own sequencing.

The Optional extra service field is available for further information if the 7 bit service specific part is not sufficient.

#### 4. Service Specifics

##### 4.1 ATM Framework Requirements

- 1) The ability to map an ATM connection to an MPLS LSP.
- 2) Support for both VC and VP ATM Connections.
- 3) The ability to transparently carry ATM User payload across an MPLS network.
- 4) The ability to transparently carry ATM Network Payload across an MPLS network. For Example:
  - a) OAM cells
  - b) PM cells
- 5) The ability to transparently carry both ATM User and Network payload across MPLS network on the same MPLS Service LSP.
- 6) The ability to carry ATM across the MPLS core efficiently for both bandwidth efficiency on the links and for performance (i.e: packet/sec) for the Core MPLS LSRs. This MAY imply the following:
  - a) Reassembling AAL5 frames before going in the MPLS Network
  - b) Segmenting AAL5 frame when exiting MPLS Network. Added bandwidth efficiency can be gained by:
    - i) Optionally/Dynamically not carrying AAL5 Trailer.
    - ii) Optionally/Dynamically not carrying AAL5 Padding.
  - c) Concatenation of ATM cells (for example: Grouping a collection of ATM cells on a single ATM connection).
- 7) When PM cells are transported then the following requirements are a MUST:
  - a) The Cells must be ordered. (for example: ATM User Cells and ATM Network Cells cannot be reordered)
  - b) The entire ATM User Payload must be carried.
- 8) When PM cells and Efficient ATM Transport (for example: AAL5 reassembly or concatenation cells) are transported, then following requirements are a MUST:
  - a) There must be no increased delay due to reassembly (for example: Don't wait for entire frame to be reassembled if a PM Cell arrives. PM Cells must cause the reassembly engine to send what it has currently reassembled as soon as a PM cell arrives, followed immediately by the PM Cell.

St-Denis

Informational-Expires May 2001  
Multi-service over MPLS

7  
November 2000

- b) The statement a) above implies that Frame Segmentation for AAL5 reassembly MUST be supported.
- c) When exiting the MPLS Network, AAL5 Padding and Trailer information must be reconstructed identically to how it looked before entering MPLS Network. (due to requirement 7b).

#### 4.1.1 ATM Service Specific Encoding

In order to accommodate the various modes of ATM cell formats the service types are broken into VC and VP modes of services. Each mode can be mixed on the same transport header but they will be described individually.

#### 4.1.2 ATM VC Service Specific Encoding

This service comprises of an individual ATM VC. The cell format of this encapsulation will be described first.




Figure 4 ATM VC Multi-service Specific Field

The CLP bit carries the CLP for the ATM cell. This information is copied from the cell when the cell is encapsulated. The network MAY use this bit directly to discard frames when buffers on queues become large. This encoding MAY be encoded in the EXP bits of the service label.

The PTI bits are the ATM Payload Type Indication bits. These bits carry the same meaning as the ATM encoding.

The Single/Concatenation indication indicates whether this is a single cell or a set of concatenated cells. Cell concatenation allows cells to be grouped together to allow more efficient transport. When cells are concatenated, each cell in the payload is 48 bytes in length. Cells can only be concatenated when all the other fields are identical.

The Cell/Frame indication indicates whether this is a single cell or that it is a frame which has been reassembled. When set to 0 it indicates a cell. The frame encoding is described below.

The RES bit is reserved for future expansion.



The EXT bit describes if the optional sequence number is carried. A zero indicates no sequence number.

Since the cell format has a slightly different meaning when the cells are concatenated together this format is outlined below.

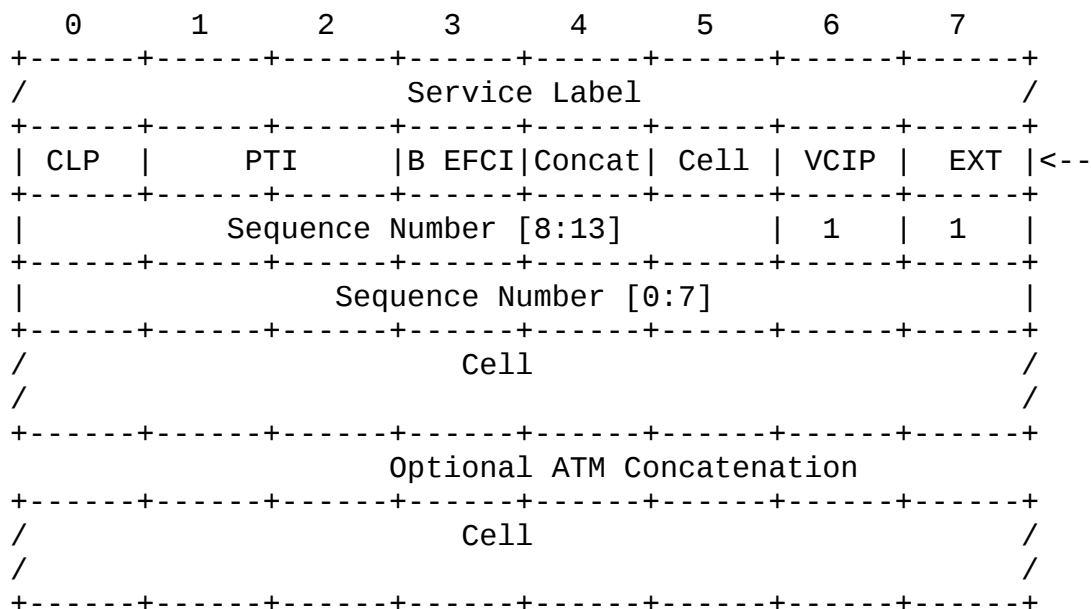
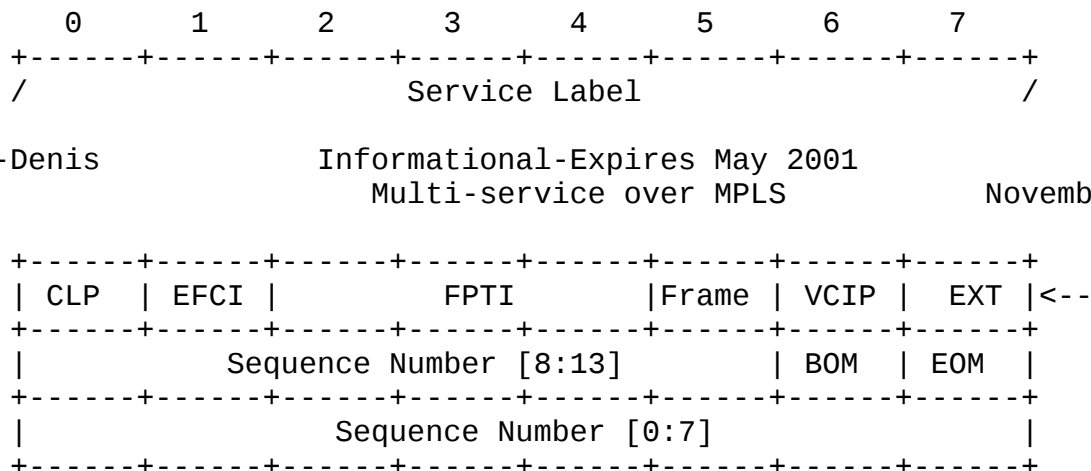


Figure 5 ATM VC Multi-service Specific Information Concatenated cell format

Cell concatenation can only be performed on user cells. Therefore, the bit 3 which indicates OAM cells is reused to carry the concatenated EFCI. This bit is set by the network if there is congestion when entering the MPLS Network. The MPLS network cannot set this field since it does not know it is carrying ATM service. This bit should be logically OR'ed into each cell when de-concatenating.

Another option at the transmitting of the service tunnel is to reassemble the cells into frames. The following is the format when carrying frames. This format can improve the efficiency of carrying frames over MPLS. This may be valuable when the speed of the ATM interface is close to the speed of the MPLS interface for example.





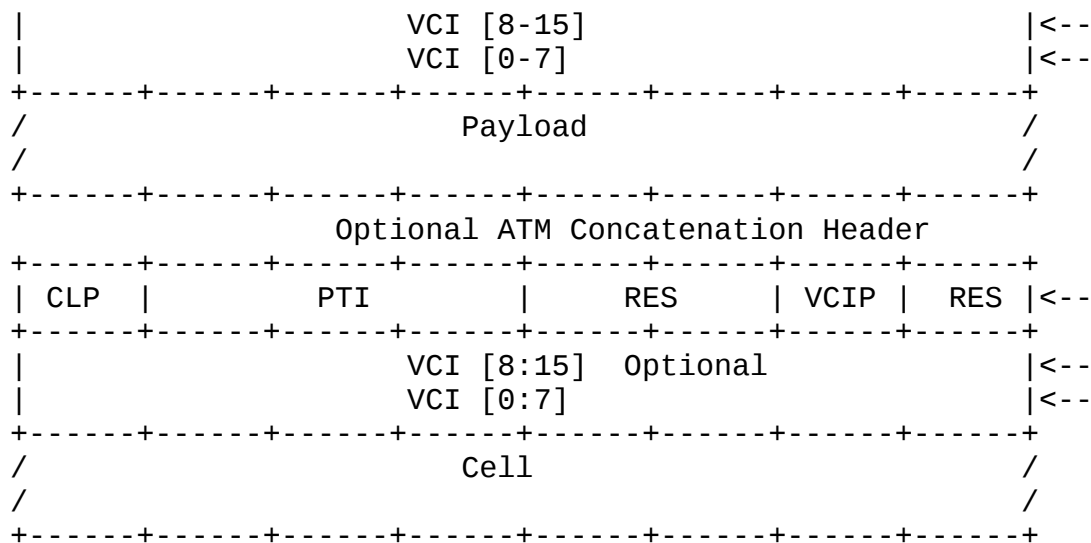


Figure 7 ATM VP Optional Extra Service Information

Figure 7 illustrates the the placement of the VCI field when included for a VP connection.

The VP encoding for concatenation is similar to the VC encoding with the exception that the VCI has to be indicated for the Cells within the payload. Each cell that is concatenated can be from the same VCI or a different VCI in the case of a VP connection.

When concatenating cells if the same VCI is concatenated in adjacent cells the VCI field may be omitted. The VCI Present (VCIP) field indicates whether the VCI is present.

## 4.2 Frame Relay Framework Requirements

- 1) Frame Transport: MUST have the ability to carry both User and Network traffic on the same Service LSP. MUST accommodate variable length Frame Relay PDUs.
- 2) Connection Mapping: MUST support a one to one mapping of a Frame Relay connection to a Service LSP.

- 3) Frame Ordering: MUST support the reliable ordering of frames so that frame order is always preserved across an end to end service connection. SHOULD support a non-ordered mode of operation for frames whose payload does not demand ordered service across the end to end network. In this case, the requirement for sequence number generation and checking is removed with appropriate improvements in bandwidth efficiency and processing requirement. MAY support the interleaving of ordered and non-ordered mode within the same LSP for the efficient carriage of multiple

traffic types.

- 4) Frame Priority: MUST support the ability to map different priority PVCs to appropriately engineered LSPs. SHOULD support the ability to map the per-frame indication of discard eligibility (DE) to the EXP bits of the service label and transport label as appropriate.
- 5) Congestion Indication: Must support the transparent transport of FECN (Forward Explicit Congestion Notification) and BECN (Backward Explicit Congestion Notification) to allow end-to-end congestion management.
- 6) PVC Status Management: MUST Support the mapping and transport of PVC active and inactive indications. The support of end-to-end connection integrity (continuity and performance) mechanisms SHOULD be provided.
- 7) Traffic Management: Mapping of the following Frame Relay traffic management parameters to the attributes of the Service LSP MAY be provided:
  - a) CIR (Committed Information Rate) or throughput.
  - b) Bc (Committed burst size).
  - c) Be (Excess Burst Size).
  - d) Access Rate.
  - e) Maximum frame size.

#### 4.2.1 Frame Relay Service Specific Encoding

St-Denis

Informational-Expires May 2001  
Multi-service over MPLS

12  
November 2000

This service comprises an individual Frame Relay Connection. The format of this encapsulation will be described first.

0	1	2	3	4	5	6	7
+-----+-----+-----+-----+-----+-----+-----+-----+							
/		Service Label				/	
+-----+-----+-----+-----+-----+-----+-----+-----+							
DE	C/R	BECN	FECN	OAM	Res(0)	Res(0)	EXT
+-----+-----+-----+-----+-----+-----+-----+-----+							




Figure 8 FR Connection Multi-service Specific Field

The DE bit indicates the discard eligibility of the encapsulated frame. This information is copied from the frame when the frame is encapsulated. The network MAY use this bit directly to discard frames when buffers on queues become large. This encoding MAY be encoded in the EXP bits of the service label.

The C/R (command/response) bit is carried end-to-end without modification.

FECN and BECN (Forward and Backward Explicit Congestion Notification) bits are set when congestion is encountered. They are never reset (to zero) by the network.


The EXT bit describes if the optional sequence number is carried. A zero indicates no sequence number.

### [4.3](#) Circuit Emulation Framework Requirements (CES)

#### [4.3.1](#) CES Requirements

- 1) CES Transport: MUST have the ability to carry both Payload and Out-of-band information. MUST accommodate variable length PDUs. Examples of out-of-band information are loop back signaling or A-B bit signaling.

0	1	2	3	4	5	6	7		
St-Denis								Informational-Expires May 2001	13
								Multi-service over MPLS	November 2000



+-----+-----+-----+-----+-----+-----+-----+-----+

Figure 9 CES Multi-service Specific Field

CES for low speed connections (DS0 to DS3) is different than for high-speed connections (OC-1 to OC-192). For low speed CES, it follows the same format as ATM's AAL1 standard.

## 5. Signaling

This section presents a functional description of the signaling information required to support multi-service encapsulation over MPLS. MPLS signaling specific formats and procedures are to be included in a subsequent version of this draft.

Multi-service encapsulation over MPLS requires information to be exchanged or coordinated by the service endpoints to ensure that compatible capabilities are at each endpoint. This information may be signaled using any of the traditional label distribution protocols (LDP, CR-LDP, RSVP-TE) by simply enhancing the protocol(s) with new TLVs/objects. The signaled information is to be processed by the service endpoints (MS LERs) and be passed transparently by each tandem LSR.

The following list indicates the information to be signaled by the MS LERs (transmitter and receiver) during LSP establishment phase. In its simplest form, the transmitter indicates the capability required for a service, and the receiver accepts or rejects the request. Optionally, a negotiation mechanism may be used in order to have the transmitter and receiver agree on a capability set or subset for a service.

### 1) Service Type

The service type indicates the service requested.

Value	Service Type
-----	-----
0	ATM
1	Frame Relay
2	CES
3	S-CES
4	Ethernet
5	VLAN

6	PPP
64-127	Vendor Specific

## 2)Service Sub Type

The service sub-type indicates the corresponding connection type of the service.

For ATM service:

Value	Service Sub-type
0	VCC
1	VPC

For Frame Relay service:

Value	Service Sub-type
0	DLCI

For CES:

Value	Service Sub-type
0	DS-0
1	DS-1
2	DS-2

For S-CES:

Value	Service Sub-type
----	-----
0	STS-1
1	STS-3c
2	STS-12c
3	STS-48c

St-Denis                      Informational-Expires May 2001  
Multi-service over MPLS

15  
November 2000

### 3) Format Indicators

Format indicators are required to request support for sequence ordering and/or support for fragmentation.

Sequencing Flag

```
Set to 0 to indicate no sequence ordering
Set to 1 to indicate sequence ordering
```

Fragmentation Flag

```
Set to 0 to indicate no fragmentation
Set to 1 to indicate fragmentation support
```

When either of the format indicators are set, the service endpoints must be able to support the MPLS extension header in the MSSL.

#### 4)Service Specific Information

Service specific information is optionally signaled when a service requires additional support for specific encoding of the PDUs. This information must be agreed by both service endpoints.

For ATM, the service specific information has been defined as follows:

##### ATM Encapsulation format

- 0 cell
- 1 frame
- 2 both

##### Cell Concatenation Indicator (only applicable for the cell encapsulation format)

- Set to 0, for single cell support
- Set to 1, for support of concatenated cells

##### PM cell Indicator

- Set to 0, for no support of PM cells
- Set to 1, for support of PM cells

For CES, the service specific information is to be defined in a subsequent draft.

#### 5.1 Examples of Signaling Options

This section provides some examples of the signaling information to be used for the different types of ATM services.

Serv	Sub- Type	Format Indicator		SSI			Service Description
		Seq. Flg	Frag.Fl g	Encap	Conc	PM	
ATM	VCC	yes	yes	frame	no	no	aal5 VCC w/ frag & sequencing
ATM	VCC	no	no	frame	no	no	aal5 VCC w/o frag
ATM	VCC	yes	yes	frame	no	yes	aal5 VCC w/ frag & PM support
ATM	VCC	no	no	cell	yes	no	aal2 VCC concat
ATM	VCC	no	no	cell	no	no	aal2 VCC no conca
ATM	VPC	no	no	both	no	no	VPC



## 5.2 ATM Switched Services

In the case of ATM switched services over MPLS, ATM signaling MAY be used to correlate the service information and to distribute the service label. Note that this is only applicable for the aggregated model, in which ATM switched services are tunneled across an MPLS transport network.

The ATM service information is signaled using standard information elements (e.g. AAL IE, ATM traffic descriptor IE, Connection Identifier IE). The label information MAY be distributed via the GIT IE in the Setup and Connect message. Procedures on how to handle label distribution via ATM signaling are out of the scope of this draft.

## 6. Security Considerations

This draft does not introduce any new security considerations to MPLS.

## 7. References

[1] \_Transport of Layer 2 Frames Over MPLS\_, [draft-martini-l2circuit-trans-mpls-02.txt](#), L. Martini et al( work in progress )

## 8. Acknowledgments

The authors would like to acknowledge the contributions to this work by Sandra Ballarte, John Davies, Tim Pearson and Greg Wright.

## 9. Author's Addresses

St-Denis	Informational-Expires May 2001	17
	Multi-service over MPLS	November 2000

Bernie St-Denis  
Nortel Networks  
[3500](#) Carling Avenue  
Nepean, Ontario  
Canada. K2H 8E9  
Email: [bernie@nortelnetworks.com](mailto:bernie@nortelnetworks.com)

Don Fedyk  
Nortel Networks  
600 Technology Park Drive  
Billerica, Massachusetts 01821  
USA  
Phone: (978) 288-3041  
Email: [dwfedyk@nortelnetworks.com](mailto:dwfedyk@nortelnetworks.com)

Aditya Bhatnagar  
Cable & Wireless  
[11700](#) Plaza America Drive,  
10th Floor  
Reston VA 20190  
USA

Adeel Siddiqui  
Cable & Wireless  
11700 Plaza America Drive,  
10th Floor  
Reston VA 20190  
USA

Aditya.Bhatnagar@cwusa.com

Adeel.Siddiqui@cwusa.com

Riad Hartani  
Caspian Networks  
Address: 170 Baytech Drive,  
San Jose, CA, U.S.A., 95134  
Phone: 408-382-5216  
Email: riad@caspiannetworks.com

Tricci So  
Caspian Networks  
Address: 170 Baytech Drive,  
San Jose, CA, U.S.A., 95134  
Phone: 408-382-5588  
Email: tso@caspiannetworks.com

St-Denis

Informational-Expires May 2001  
Multi-service over MPLS

18  
November 2000

#### 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 Standards process must be followed, or as required to translate it into 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 ASK 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.