

Network Working Group
Request for Comments: 4906
Category: Historic

L. Martini, Ed.
E. Rosen, Ed.
Cisco Systems, Inc.
N. El-Aawar, Ed.
Level 3 Communications, LLC
June 2007

Transport of Layer 2 Frames Over MPLS

Status of This Memo

This memo defines a Historic Document for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document describes methods for transporting the Protocol Data Units (PDUs) of layer 2 protocols such as Frame Relay, Asynchronous Transfer Mode (ATM) Adaption Layer 5 (AAL5), and Ethernet, and for providing a Synchronized Optical Network (SONET) circuit emulation service across an MPLS network. This document describes the so-called "draft-martini" protocol, which has since been superseded by the Pseudowire Emulation Edge to Edge Working Group specifications described in RFC 4447 and related documents.

Table of Contents

1. Introduction	3
2. Specification of Requirements	3
3. Special Note	3
4. Tunnel Labels and Virtual Circuit (VC) Labels	4
5. Protocol-Specific Details	5
5.1. Frame Relay	5
5.2. ATM	6
5.2.1. ATM AAL5 VCC Transport	6
5.2.2. ATM Transparent Cell Transport	6
5.2.3. ATM VCC and VPC Cell Transport	6
5.2.4. OAM Cell Support	6
5.2.5. ILMI Support	7
5.3. Ethernet VLAN	7
5.4. Ethernet	8
5.5. HDLC	8
5.6. PPP	8
6. LDP	8
6.1. Interface Parameters Field	10
6.2. C Bit Handling Procedures	12
6.2.1. VC Types for Which the Control Word is REQUIRED	12
6.2.2. VC Types for Which the Control Word is NOT Mandatory	12
6.2.3. Status Codes	15
6.3. LDP Label Withdrawal Procedures	15
6.4. Sequencing Considerations	15
6.4.1. Label Mapping Advertisements	15
6.4.2. Label Mapping Release	16
7. IANA Considerations	16
8. Security Considerations	16
9. Normative References	17
10. Informative References	18
11. Co-Authors	18

1. Introduction

In an MPLS network, it is possible to carry the Protocol Data Units (PDUs) of layer 2 protocols by prepending an MPLS label stack to these PDUs. This document specifies the necessary label distribution procedures for accomplishing this using the encapsulation methods in [RFC4905]. We restrict discussion to the case of point-to-point transport. Quality of service (QoS)-related issues are not discussed in this document. This document describes methods for transporting a number of protocols; in some cases, transporting a particular protocol may have several modes of operation. Each of these protocols and/or modes may be implemented independently.

An accompanying document [CEM] also describes a method for transporting time division multiplexed (TDM) digital signals (TDM circuit emulation) over a packet-oriented MPLS network. The transmission system for circuit-oriented TDM signals is the Synchronous Optical Network (SONET) [ANSI.T1.105] / Synchronous Digital Hierarchy (SDH) [ITU.G.707]. To support TDM traffic, which includes voice, data, and private leased line service, the MPLS network must emulate the circuit characteristics of SONET/SDH payloads. MPLS labels and a new circuit emulation header are used to encapsulate TDM signals and provide the Circuit Emulation Service over MPLS (CEM).

2. Specification of Requirements

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

3. Special Note

This document describes the so-called "draft-martini" protocol, which is used in many deployed implementations. This document and its contents have since been superseded by the Pseudowire Emulation Edge to Edge Working Group specifications: [RFC4447], [RFC4385], [RFC4448], [RFC4717], [RFC4618], [RFC4619], [RFC4553], [RFC4842], and related documents. This document serves as a documentation of current implementations, and MUST NOT be used for new implementations. The PWE3 Label Distribution Protocol (LDP) control document [RFC4447], which is backward compatible with this document, MUST be used for all new implementations of this protocol.

4. Tunnel Labels and Virtual Circuit (VC) Labels

Suppose it is desired to transport layer 2 PDUs from ingress Label Switching Router (LSR) R1 to egress LSR R2, across an intervening MPLS network. We assume that there is a Label Switched Path (LSP) from R1 to R2. That is, we assume that R1 can cause a packet to be delivered to R2 by pushing some label onto the packet and sending the result to one of its adjacencies. Call this label the "tunnel label", and the corresponding LSP the "tunnel LSP".

The tunnel LSP merely gets packets from R1 to R2; the corresponding label doesn't tell R2 what to do with the payload. In fact, if penultimate hop popping is used, R2 may never even see the corresponding label. (If R1 itself is the penultimate hop, a tunnel label may not even get pushed on.) Thus, if the payload is not an IP packet, there must be a label, which becomes visible to R2, that tells R2 how to treat the received packet. Call this label the "VC label".

So when R1 sends a layer 2 PDU to R2, it first pushes a VC label on its label stack, and then (if R1 is not adjacent to R2) pushes on a tunnel label. The tunnel label gets the MPLS packet from R1 to R2; the VC label is not visible until the MPLS packet reaches R2. R2's disposition of the packet is based on the VC label.

Note that the tunnel could be a Generic Routing Encapsulation (GRE)-encapsulated MPLS tunnel between R1 and R2. In this case, R1 would be adjacent to R2, and only the VC label would be used, and the intervening network need only carry IP packets.

If the payload of the MPLS packet is, for example, an ATM AAL5 PDU, the VC label will generally correspond to a particular ATM VC at R2. That is, R2 needs to be able to infer from the VC label the outgoing interface and the VPI/VCI (Virtual Path Identifier / Virtual Circuit Identifier) value for the AAL5 PDU. If the payload is a Frame Relay PDU, then R2 needs to be able to infer from the VC label the outgoing interface and the DLCI (Data Link Connection Identifier) value. If the payload is an Ethernet frame, then R2 needs to be able to infer from the VC label the outgoing interface, and perhaps the VLAN identifier. This process is unidirectional, and will be repeated independently for bidirectional operation. It is REQUIRED to assign the same VC ID, and VC type for a given circuit in both directions. The group ID (see below) MUST NOT be required to match in both directions. The transported frame MAY be modified when it reaches the egress router. If the header of the transported layer 2 frame is modified, this MUST be done at the egress LSR only.

Note that 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. Of course, as the packet is transported across the MPLS network, additional labels may be pushed on (and then popped off) as needed. Even R1 itself may push on additional labels above the tunnel label. If R1 and R2 are directly adjacent LSRs, then it may not be necessary to use a tunnel label at all.

This document does not specify a method for distributing the tunnel label or any other labels that may appear above the VC label on the stack. Any acceptable method of MPLS label distribution will do.

This document does specify a method for assigning and distributing the VC label. Static label assignment MAY be used, and implementations SHOULD provide support for this. When signaling is used, the VC label MUST be distributed from R2 to R1 using LDP in the downstream unsolicited mode; this requires that an LDP session be created between R1 and R2. It should be noted that this LDP session is not necessarily transported along the same path as the Layer 2 PDUs [RFC3036]. In addition, when using LDP to distribute the VC label, liberal label retention mode SHOULD be used. However, as required in [RFC3036], the label request operation (mainly used by conservative label retention mode) MUST be implemented. VC labels MUST be allocated from the per-platform label space.

Note that this technique allows an unbounded number of layer 2 "VCs" to be carried together in a single "tunnel". Thus, it scales quite well in the network backbone.

While this document currently defines the emulation of Frame Relay and ATM Permanent Virtual Circuit (PVC) services, it specifically does not preclude future enhancements to support switched service (Switched Virtual Circuit (SVC) and Switched Permanent Virtual Circuit (SPVC)) emulation.

5. Protocol-Specific Details

5.1. Frame Relay

The Frame Relay PDUs are encapsulated according to the procedures defined in [RFC4905]. The MPLS edge LSR MUST provide Frame Relay PVC status signaling to the Frame Relay network. If the MPLS edge LSR detects a service affecting condition, as defined in [Q.933] Annex A.5 cited in Implementation Agreement FRF.1.1, it MUST withdraw the label that corresponds to the frame relay DLCI. The egress LSR SHOULD generate the corresponding errors and alarms as defined in [Q.933] on the egress Frame relay VC.

5.2. ATM

5.2.1. ATM AAL5 VCC Transport

ATM AAL5 Common Part Convergence Sublayer - Service Data Units (CPCS-SDUs) are encapsulated according to [RFC4905] ATM AAL5 CPCS-SDU mode. This mode allows the transport of ATM AAL5 CPCS-SDUs traveling on a particular ATM PVC across the MPLS network to another ATM PVC.

5.2.2. ATM Transparent Cell Transport

This mode is similar to the Ethernet port mode. Every cell that is received at the ingress ATM port on the ingress LSR, R1, is encapsulated according to [RFC4905], ATM cell mode, and sent across the LSP to the egress LSR, R2. This mode allows an ATM port to be connected to only one other ATM port. [RFC4905] allows for grouping of multiple cells into a single MPLS frame. Grouping of ATM cells is OPTIONAL for transmission at the ingress LSR, R1. If the Egress LSR R2 supports cell concatenation, the ingress LSR, R1, should only concatenate cells up to the "Maximum Number of concatenated ATM cells" parameter received as part of the FEC element.

5.2.3. ATM VCC and VPC Cell Transport

This mode is similar to the ATM AAL5 Virtual Channel Connection (VCC) transport except that cells are transported. Every cell that is received on a pre-defined ATM PVC or ATM Permanent Virtual Path (PVP), at the ingress ATM port on the ingress LSR, R1, is encapsulated according to [RFC4905], ATM cell mode, and sent across the LSP to the egress LSR R2. Grouping of ATM cells is OPTIONAL for transmission at the ingress LSR, R1. If the egress LSR R2 supports cell concatenation, the ingress LSR, R1, MUST only concatenate cells up to the "Maximum Number of concatenated ATM cells in a frame" parameter received as part of the FEC element.

5.2.4. OAM Cell Support

Operations and Management (OAM) cells MAY be transported on the VC LSP. When the LSR is operating in AAL5 CPCS-SDU transport mode, if it does not support transport of ATM cells, the LSR MUST discard incoming MPLS frames on an ATM VC LSP that contain a VC label with the T bit set [RFC4905]. When operating in AAL5 SDU transport mode, an LSR that supports transport of OAM cells using the T bit defined in [RFC4905], or an LSR operating in any of the three cell transport modes, MUST follow the procedures outlined in [FAST] Section 8 for mode 0 only, in addition to the applicable procedures specified in [ITU.G.707].

5.2.4.1. OAM Cell Emulation Mode

AN LSR that does not support transport of OAM cells across an LSP MAY provide OAM support on ATM PVCs using the following procedures:

A pair of LSRs MAY emulate a bidirectional ATM VC by two unidirectional LSPs. If an F5 end-to-end OAM cell is received from a ATM VC, by either LSR that is transporting this ATM VC, with a loopback indication value of 1, and the LSR has a label mapping for the ATM VC, then the LSR MUST decrement the loopback indication value and loop back the cell on the ATM VC. Otherwise, the loopback cell MUST be discarded by the LSR.

The ingress LSR, R1, may also optionally be configured to periodically generate F5 end-to-end loopback OAM cells on a VC. If the LSR fails to receive a response to an F5 end-to-end loopback OAM cell for a pre-defined period of time it MUST withdraw the label mapping for the VC.

If an ingress LSR, R1, receives an AIS (Alarm Indication Signal) F5 OAM cell, or R1 fails to receive a pre-defined number of the End-to-End loop OAM cells, or a physical interface goes down, it MUST withdraw the label mappings for all VCs associated with the failure. When a VC label mapping is withdrawn, the egress LSR, R2, MUST generate AIS F5 OAM cells on the VC associated with the withdrawn label mapping. In this mode it is very useful to apply a unique group ID to each interface. In the case where a physical interface goes down, a wild card label withdraw can be sent to all LDP neighbors, greatly reducing the signaling response time.

5.2.5. ILMI Support

An MPLS edge LSR MAY provide an ATM Integrated Local Management Interface (ILMI) to the ATM edge switch. If an ingress LSR receives an ILMI message indicating that the ATM edge switch has deleted a VC, or if the physical interface goes down, it MUST withdraw the label mappings for all VCs associated with the failure. When a VC label mapping is withdrawn, the egress LSR SHOULD notify its client of this failure by deleting the VC using ILMI.

5.3. Ethernet VLAN

The Ethernet frame will be encapsulated according to the procedures in [RFC4905]. It should be noted that if the VLAN identifier is modified by the egress LSR, according to the procedures outlined above, the Ethernet spanning tree protocol might fail to work properly. If the LSR detects a failure on the Ethernet physical

port, or the port is administratively disabled, it MUST withdraw the label mappings for all VCs associated with the port.

5.4. Ethernet

The Ethernet frame will be encapsulated according to the procedures in [RFC4905]. If the LSR detects a failure on the Ethernet physical port, or the port is administratively disabled, the corresponding VC label mapping MUST be withdrawn.

5.5. HDLC


HDLC (High-Level Data Link Control) frames are encapsulated according to the procedures in [RFC4905]. If the MPLS edge LSR detects that the physical link has failed, or the port is administratively disabled, it MUST withdraw the label mapping that corresponds to the HDLC link.

5.6. PPP

PPP frames are encapsulated according to the procedures in [RFC4905]. If the MPLS edge LSR detects that the physical link has failed, or the port is administratively disabled, it MUST withdraw the label mapping that corresponds to the PPP link.

6. LDP

The VC label bindings are distributed using the LDP downstream unsolicited mode described in [RFC3036]. The LSRs will establish an LDP session using the Extended Discovery mechanism described in sections 2.4.2 and 2.5 of [RFC3036]; for this purpose, a new type of FEC element is defined. The FEC element type is 128. Only a single VC FEC element MUST be advertised per LDP VC label. The Virtual Circuit FEC element is defined as follows:



- VC Type

A 15-bit quantity containing a value that represents the type of VC. Assigned values are:

VC Type	Description
0x0001	Frame Relay DLCI
0x0002	ATM AAL5 VCC transport
0x0003	ATM transparent cell transport
0x0004	Ethernet VLAN
0x0005	Ethernet
0x0006	HDLC
0x0007	PPP
0x8008	CEM [CEM]
0x0009	ATM VCC cell transport
0x000A	ATM VPC cell transport

- Control word bit (C)

The highest order bit (C) of the VC type is used to flag the presence of a control word (defined in [RFC4905]) as follows:

bit 15 = 1 control word present on this VC.
bit 15 = 0 no control word present on this VC.

Please see Section 6.2, "C Bit Handling Procedures", for further explanation.

- VC information length

Length of the VC ID field and the interface parameters field in octets. If this value is 0, then it references all VCs using the specified group ID, and there is no VC ID present, nor any interface parameters.

- Group ID

An arbitrary 32-bit value, which represents a group of VCs that is used to create groups in the VC space. The group ID is intended to be used as a port index, or a virtual tunnel index. To simplify configuration, a particular VC ID at ingress could be part of the virtual tunnel for transport to the egress router. The group ID is very useful to send wild card label withdrawals to remote LSRs upon physical port failure.

- VC ID


A non-zero 32-bit connection ID that, together with the VC type, identifies a particular VC.

- Interface parameters

This variable length field is used to provide interface-specific parameters, such as interface MTU.

6.1. Interface Parameters Field

This field specifies interface-specific parameters. When applicable, it MUST be used to validate that the LSRs, and the ingress and egress ports at the edges of the circuit, have the necessary capabilities to interoperate with each other. The field structure is defined as follows:



The parameter ID is defined as follows:

Parameter	ID	Length	Description
0x01		4	Interface MTU in octets.
0x02		4	Maximum Number of concatenated ATM cells.
0x03	up to	82	Optional Interface Description string.
0x04		4	CEM [CEM] Payload Bytes.
0x05		4	CEM options.

The Length field is defined as the length of the interface parameter including the Parameter ID and Length field itself. Processing of the interface parameters should continue when encountering unknown interface parameters, and they MUST be silently ignored.

- Interface MTU

A 2-octet value indicating the MTU in octets. This is the Maximum Transmission Unit, excluding encapsulation overhead, of the egress packet interface that will be transmitting the decapsulated PDU that is received from the MPLS network. This

parameter is applicable only to VC types 1, 2, 4, 5, 6, and 7, and is REQUIRED for these VC types. If this parameter does not match in both directions of a specific VC, that VC MUST NOT be enabled.

- Maximum Number of concatenated ATM cells

A 2-octet value specifying the maximum number of concatenated ATM cells that can be processed as a single PDU by the egress LSR. An ingress LSR transmitting concatenated cells on this VC can concatenate a number of cells up to the value of this parameter, but MUST NOT exceed it. This parameter is applicable only to VC types 3, 9, and 0x0a, and is REQUIRED for these VC types. This parameter does not need to match in both directions of a specific VC.

- Optional Interface Description string

This arbitrary, OPTIONAL interface description string can be used to send an administrative description text string to the remote LSR. This parameter is OPTIONAL, and is applicable to all VC types. The interface description parameter string length is variable, and can be from 0 to 80 octets.

- Payload Bytes

A 2-octet value indicating the number of TDM payload octets contained in all packets on the CEM stream from 48 to 1,023 octets. All of the packets in a given CEM stream have the same number of payload bytes. Note that there is a possibility that the packet size may exceed the Synchronous Payload Envelope (SPE) size in the case of an STS-1 SPE, which could cause two pointers to be needed in the CEM header, since the payload may contain two J1 bytes for consecutive SPEs. For this reason, the number of payload bytes must be less than or equal to 783 for STS-1 SPEs.

- CEM Options

An optional 16-bit value of CEM flags. See [CEM] for the definition of the bit values.

6.2. C Bit Handling Procedures

6.2.1. VC Types for Which the Control Word is REQUIRED

The Label Mapping messages which are sent in order to set up these VCs MUST have c=1. When a Label Mapping message for a VC of one of these types is received, and c=0, a Label Release MUST be sent, with an "Illegal C-bit" status code. In this case, the VC will not come up.

6.2.2. VC Types for Which the Control Word is NOT Mandatory

If a system is capable of sending and receiving the control word on VC types for which the control word is not mandatory, then each such VC endpoint MUST be configurable with a parameter that specifies whether the use of the control word is PREFERRED or NOT PREFERRED. For each VC, there MUST be a default value of this parameter. This specification does NOT state what the default value should be.

If a system is NOT capable of sending and receiving the control word on VC types for which the control word is not mandatory, then it behaves exactly as if it were configured for the use of the control word to be NOT PREFERRED.

If a Label Mapping message for the VC has already been received, but no Label Mapping message for the VC has yet been sent, then the procedure is the following:

- i. If the received Label Mapping message has c=0, send a Label Mapping message with c=0, and the control word is not used.
- ii. If the received Label Mapping message has c=1, and the VC is locally configured such that the use of the control word is preferred, then send a Label Mapping message with c=1, and the control word is used.
- iii. If the received Label Mapping message has c=1, and the VC is locally configured such that the use of the control word is not preferred or the control word is not supported, then act as if no Label Mapping message for the VC had been received (i.e., proceed to the next paragraph).

If a Label Mapping message for the VC has not already been received (or if the received Label Mapping message had c=1, and either local configuration says that the use of the control word is not preferred or the control word is not supported), then send a Label Mapping message in which the c bit is set to correspond to the locally configured preference for use of the control word. (That is, set c=1

if locally configured to prefer the control word, set c=0 if locally configured to prefer not to use the control word or if the control word is not supported).

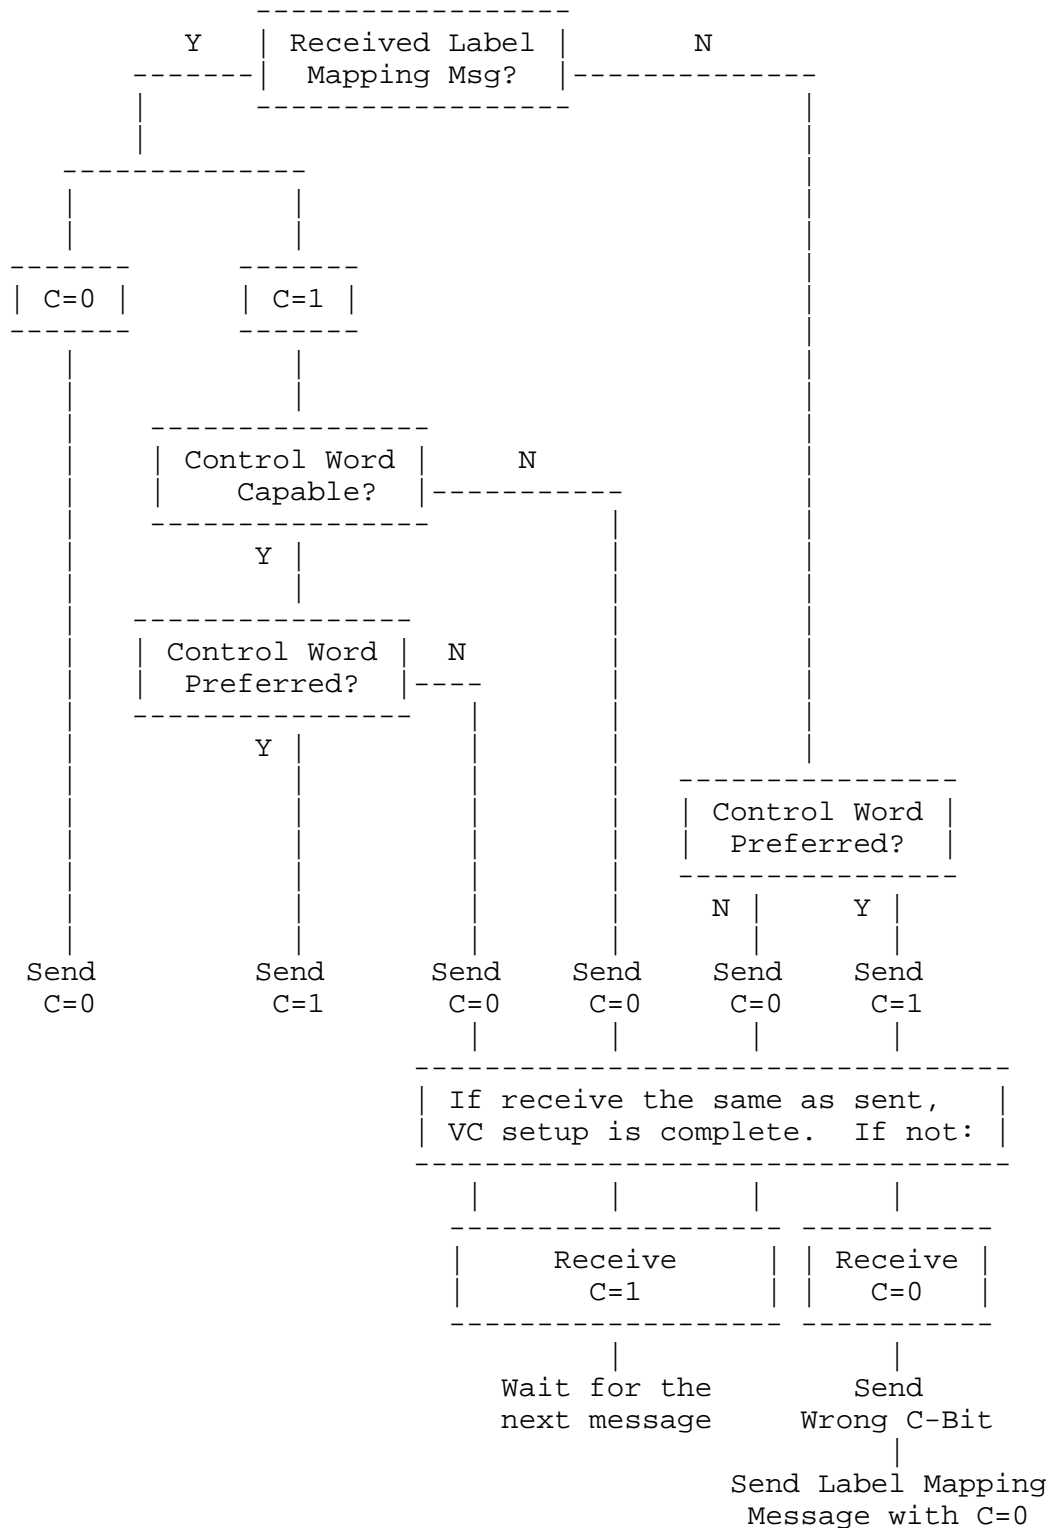
The next action depends on what control message is next received for that VC. The possibilities are:

- i. A Label Mapping message with the same c bit value as specified in the Label Mapping message that was sent. VC setup is now complete, and the control word is used if c=1 but not used if c=0.
- ii. A Label Mapping message with c=1, but the Label Mapping message that was sent has c=0. In this case, ignore the received Label Mapping message, and continue to wait for the next control message for the VC.
- iii. A Label Mapping message with c=0, but the Label Mapping message that was sent has c=1. In this case, send a Label Withdraw message with a "Wrong C-bit" status code, followed by a Label Mapping message that has c=0. VC setup is now complete, and the control word is not used.
- iv. A Label Withdraw message with the "Wrong C-bit" status code. Treat as a normal Label Withdraw, but do not respond. Continue to wait for the next control message for the VC.

If, at any time after a Label Mapping message has been received, a corresponding Label Withdraw or Release is received, the action taken is the same as for any Label Withdraw or Release that might be received at any time.

If both endpoints prefer the use of the control word, this procedure will cause it to be used. If either endpoint prefers not to use the control word, or does not support the control word, this procedure will cause it not to be used. If one endpoint prefers to use the control word but the other does not, the one that prefers not to use it is has no extra protocol to execute; it just waits for a Label Mapping message that has c=0.

The following diagram illustrates the above procedures:



6.2.3. Status Codes

RFC 3036 has a range of Status Code values, which are assigned by IANA on a First Come, First Served basis. These are in the range 0x20000000-0x3effffff. The following new status codes are defined:

```
0x20000001 "Illegal C-Bit"
0x20000002 "Wrong C-Bit"
```

6.3. LDP Label Withdrawal Procedures

As mentioned above, the Group ID field can be used to withdraw all VC labels associated with a particular group ID. This procedure is OPTIONAL, and if it is implemented, the LDP label withdraw message should be as follows: the VC information length field is set to 0, the VC ID field is not present, and the interface parameters field is not present. For the purpose of this document, this is called the "wild card withdraw procedure", and all LSRs implementing this design are REQUIRED to accept such a withdraw message, but are not required to send it.

The interface parameters field MUST NOT be present in any LDP VC label withdrawal message or release message. A wild card release message MUST include only the group ID. A Label Release message initiated from the imposition router must always include the VC ID.

6.4. Sequencing Considerations

In the case where the router considers the sequence number field in the control word, it is important to note the following when advertising labels.

6.4.1. Label Mapping Advertisements

After a label has been withdrawn by the disposition router and/or released by the imposition router, care must be taken to not re-advertise (reuse) the released label until the disposition router can be reasonably certain that old packets containing the released label no longer persist in the MPLS network.

This precaution is required to prevent the imposition router from restarting packet forwarding with sequence number of 1 when it receives the same label mapping if there are still older packets persisting in the network with sequence number between 1 and 32768. For example, if there is a packet with sequence number=n where n is in the interval[1,32768] traveling through the network, it would be possible for the disposition router to receive that packet after it re-advertises the label. Since the label has been released by the

imposition router, the disposition router SHOULD be expecting the next packet to arrive with sequence number of 1. Receipt of a packet with sequence number equal to n will result in n packets potentially being rejected by the disposition router until the imposition router imposes a sequence number of n+1 into a packet. Possible methods to avoid this are for the disposition router to always advertise a different VC label, or for the disposition router to wait for a sufficient time before attempting to re-advertise a recently released label. This is only an issue when sequence number processing at the disposition router is enabled.

6.4.2. Label Mapping Release

In situations where the imposition router wants to restart forwarding of packets with sequence number 1, the router shall 1) send a label mapping release to the disposition router, and 2) send a label mapping request to the disposition router. When sequencing is supported, advertisement of a VC label in response to a label mapping request MUST also consider the issues discussed in Section 6.4.1.

7. IANA Considerations

As specified in this document, a Virtual Circuit FEC element contains the VC Type field. VC Type value 0 is reserved. VC Type values 1 through 10 are defined in this document. VC Type values 11 through 63 are to be assigned by IANA using the "IETF Consensus" policy defined in RFC 2434. VC Type values 64 through 127 are to be assigned by IANA, using the "First Come First Served" policy defined in RFC 2434. VC Type values 128 through 32767 are vendor-specific, and values in this range are not to be assigned by IANA.

As specified in this document, a Virtual Circuit FEC element contains the Interface Parameters field, which is a list of one or more parameters, and each parameter is identified by the Parameter ID field. Parameter ID value 0 is reserved. Parameter ID values 1 through 5 are defined in this document. Parameter ID values 6 through 63 are to be assigned by IANA using the "IETF Consensus" policy defined in RFC 2434. Parameter ID values 64 through 127 are to be assigned by IANA, using the "First Come First Served" policy defined in RFC 2434. Parameter ID values 128 through 255 are vendor-specific, and values in this range are not to be assigned by IANA.

8. Security Considerations

This document does not affect the underlying security issues of MPLS, described in [RFC3032]. More detailed security considerations are also described in Section 8 of [RFC4447].

9. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, February 2006.
- [RFC4842] Malis, A., Pate, P., Cohen, R., Ed., and D. Zelig, "Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Circuit Emulation over Packet (CEP)", RFC 4842, April 2007.
- [RFC4553] Vainshtein, A., Ed., and YJ. Stein, Ed., "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, June 2006.
- [RFC4619] Martini, L., Ed., Kawa, C., Ed., and A. Malis, Ed., "Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks", RFC 4619, September 2006.
- [RFC4717] Martini, L., Jayakumar, J., Bocci, M., El-Aawar, N., Brayley, J., and G. Koleyni, "Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks", RFC 4717, December 2006.
- [RFC4618] Martini, L., Rosen, E., Heron, G., and A. Malis, "Encapsulation Methods for Transport of PPP/High-Level Data Link Control (HDLC) over MPLS Networks", RFC 4618, September 2006.
- [RFC4448] Martini, L., Ed., Rosen, E., El-Aawar, N., and G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, April 2006.
- [RFC3036] Andersson, L., Doolan, P., Feldman, N., Fredette, A., and B. Thomas, "LDP Specification", RFC 3036, January 2001.
- [Q.933] ITU-T Recommendation Q.933, and Q.922 Specification for Frame Mode Basic call control, ITU Geneva 1995.

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [ANSI.T1.105] American National Standards Institute, "Synchronous Optical Network Formats," ANSI T1.105-1995.
- [ITU.G.707] ITU Recommendation G.707, "Network Node Interface For The Synchronous Digital Hierarchy", 1996.
- [RFC4905] Martini, L., Ed., Rosen, E., Ed., and N. El-Aawar, Ed., "Encapsulation Methods for Transport of Layer 2 Frames over MPLS Networks", RFC 4905, June 2007.

10. Informative References

- [CEM] Malis, A., Brayley, J., Vogelsang, S., Shirron, J., and L. Martini, "SONET/SDH Circuit Emulation Service Over MPLS (CEM) Encapsulation", Work in Progress, June 2007.
- [FAST] ATM Forum, "Frame Based ATM over SONET/SDH Transport (FAST)", af-fbatm-0151.000, July 2000.

11. Co-Authors

Giles Heron
Tellabs
Abbey Place
24-28 Easton Street
High Wycombe
Bucks
HP11 1NT
UK
Email: giles.heron@tellabs.com

Dimitri Stratton Vlachos
Mazu Networks, Inc.
125 Cambridgepark Drive
Cambridge, MA 02140
Email: d@mazunetworks.com

Dan Tappan
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
EMail: tappan@cisco.com

Jayakumar Jayakumar,
Cisco Systems Inc.
225, E.Tasman, MS-SJ3/3,
San Jose, CA 95134
EMail: jjayakum@cisco.com

Alex Hamilton,
Cisco Systems Inc.
285 W. Tasman, MS-SJCI/3/4,
San Jose, CA 95134
EMail: tahamilt@cisco.com

Steve Vogelsang
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: sjv@laurelnetworks.com

John Shirron
Laurel Networks, Inc.
Omega Corporate Center
1300 Omega Drive
Pittsburgh, PA 15205
EMail: jshirron@laurelnetworks.com

Toby Smith
Network Appliance, Inc.
800 Cranberry Woods Drive
Suite 300
Cranberry Township, PA 16066
EMail: tob@netapp.com

Andrew G. Malis
Tellabs
90 Rio Robles Dr.
San Jose, CA 95134
EMail: Andy.Malis@tellabs.com

Vinai Sirkay
Reliance Infocomm
Dhirubai Ambani Knowledge City
Navi Mumbai 400 709
India
EMail: vinai@sirkay.com

Vasile Radoaca
Nortel Networks
600 Technology Park
Billerica MA 01821
EMail: vasile@nortelnetworks.com

Chris Liljenstolpe
Alcatel
11600 Sallie Mae Dr.
9th Floor
Reston, VA 20193
EMail: chris.liljenstolpe@alcatel.com

Dave Cooper
Global Crossing
960 Hamlin Court
Sunnyvale, CA 94089
EMail: dcooper@gblix.net

Kireeti Kompella
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089
EMail: kireeti@juniper.net

Authors' Addresses

Luca Martini
Cisco Systems, Inc.
9155 East Nichols Avenue, Suite 400
Englewood, CO 80112
EMail: lmartini@cisco.com

Nasser El-Aawar
Level 3 Communications, LLC.
1025 Eldorado Blvd.
Broomfield, CO 80021
EMail: nna@level3.net

Eric Rosen
Cisco Systems, Inc.
250 Apollo Drive
Chelmsford, MA 01824
EMail: erosen@cisco.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

