
Question(s): 5/17

Geneva, 10-19 September 2003

TEMPORARY DOCUMENT

Source: Editor

Title: Working text for a proposed Amendment to draft new Recommendation X.84

Editor's note:

Enclosed is a working text for a proposed Amendment to draft new Recommendation X.84. It is based on the following documents D.97 and D.98 submitted to this meeting of SG 17.

This text requires further development to reflect the work being done in the IETF PWE3 Working Group on "Pseudo-wire setup and maintenance". Contributions on this topic are required in order to have the text ready for consent at the next meeting of SG 17 in March.

The proposed amendment will provide an overview of the control protocol signalling and will define procedures for PVC status monitoring over MPLS networks.

© ITU 2003

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

12 PVC Status monitoring

This clause specifies how the status of a frame relay PVC is reported between two PEs when the PVC is created, deleted, or when its state changes between inactive and active.

12.1 Frame Relay/MPLS Port Mode Signalling

Frame relay port mode provides transport of all frame relay VCs (including DLCI 0) in the same VC LSP. PVC status monitoring is performed on DLCI 0 by the frame relay devices connected to PEs.

The FR devices perform keep alive polling which is transferred across the MPLS network. When a frame relay device detects that a MPLS LSP is inoperative, each FR device notifies users of the PVCs associated with the MPLS port that the PVCs are inactive. The PVC status changes are propagated through the adjacent network(s) to the remote users.

12.2 Frame Relay/MPLS One-to-One Mode Signalling

Relationship with the frame relay PVC Status Signaling

The PVC Status or PVC attribute information received from the management plane or the attachment frame relay UNI or NNI is propagated to the corresponding LDP session of the LSP carrying the frame relay traffic. The status or information is then mapped onto the PW Status TLV and the *Notification* message for the pseudo-wire LSP is refreshed with the TLV containing the new information.

This is done each time the status or information for a PVC changes.

12.3 LDP Messages and TLVs used for PVC Status and Attribute signaling

Editor's note: Messages need to be included

12.4 Procedures

12.4.1 New PVCs

A new PVC is reported via signaling of new pseudo wire LSPs. This is done once for each direction. . The PE serving as the egress PE, with regard to the direction of the pseudo-wire LSP, sends a *Label Mapping* message per Appendix A. Once both LSPs for the pseudo-wire are established, the PVC state is assumed inactive and *Notification* messages must be exchanged between the PW endpoints to determine the PVC state.

12.4.2 Deleted PVCs

A PVC is deleted when its provisioned information has been removed from the network providing the frame relay attachment connection. The deletion of a PVC results in the Pseudo wire LSPs for

both directions of the PVC, to be cleared per the procedures in Appendix II, and a deleted indication to be propagated across both frame relay interfaces.

12.4.3 Status Change - Active to Inactive

Should one of the frame relay interfaces reflect that the PVC state has transitioned from active to inactive, this indication must be propagated to the remote interface. A *Notification* message is sent to the remote peer with a PW Status TLV indicating that the PVC corresponding to the pseudo-wire LSP is inactive. The remote peer propagates the inactive status for the PVC across its corresponding frame relay interface using frame relay PVC status signalling procedures.

12.4.4 Status Change - Inactive to Active

Should one of the frame relay interfaces reflect that the PVC state has transitioned from inactive to active, this indication must be propagated to the remote interface. A *Notification* message is sent to the remote peer with a PW Status TLV indicating that the PVC corresponding to the pseudo-wire LSP is active. If the local PVC state is active, the remote peer propagates the active status for the PVC across its corresponding frame relay interface using frame relay PVC status signaling procedures.

12.4.5 Status Change - Frame Relay UNI/NNI Service Affecting Conditions

Any PE that detects a service affecting condition for a given UNI or NNI, must follow the Status Change- Active to Inactive procedures for each PVC on the affected interface.

Once the service affecting condition has cleared, the PE must follow the Status Change - Inactive to Active procedures for each LSP on the restored interface.

APPENDIX II

PVC Pseudo-Wire Establishment for Frame Relay

This Appendix describes the procedures for establishing VC LSP label for transport of one-to-one mode frame relay connection. These procedures are based on LDP and provide extensions for PVC status signaling for one-to-one mode.

II.1 General

A bidirectional tunnel LSP is established, to transport frame relay PDUs from ingress LSR PE1 to egress LSR PE2, across an intervening MPLS network.

The payload of the MPLS packet is a frame relay PDU, so the VC LSP label will generally correspond to a particular frame relay VC or frame relay Port at PE2. That is, PE2 needs to be able to infer from the VC LSP label the outgoing interface and the DLCI value for the frame relay PDU. This process is uni-directional, and will be repeated independently for bi-directional operation.

This Appendix does not specify a method for distributing the PSN tunnel labels that appear above the VC LSP 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 LSP label. The VC LSP label shall be distributed from PE2 to PE1 using LDP in the downstream unsolicited mode; this

requires that an LDP session be created between PE1 and PE2. It should be noted that this LDP session is not necessarily transported along the same path as the frame relay PDUs. when using LDP to distribute the PW label, liberal label retention mode should be used.

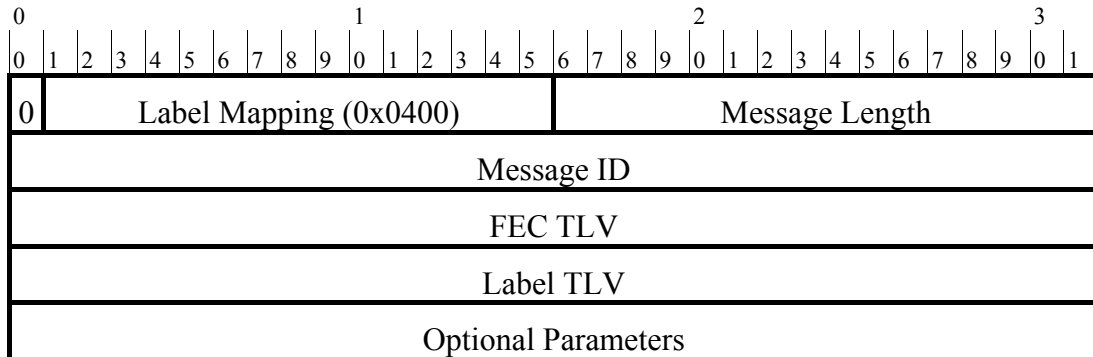
While this Appendix currently defines the emulation of Frame Relay PVC services, it specifically does not preclude future enhancements to support switched service (SVC and SPVC) emulation.

II.2 Protocol

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 X.36 sited in Implementation Agreement FRF1.2, it must withdraw the label that corresponds to the frame relay DLCI or frame relay port. The Egress LSR should generate the corresponding errors and alarms as defined in X.36 on the egress Frame relay VC.

The VC LSP label binding and status signaling procedures using LDP procedures are described below. These procedures are based on IETF PWE3 control protocol. Label bindings are distributed using the LDP downstream unsolicited mode described in to advertise the VC-FEC for a given Pseudowire [1]. The LSRs will establish an LDP session using the Extended Discovery mechanism described in [1], section 2.4.2 and 2.5.or this purpose a new type of FEC element is defined.

The LDP document section 3.5.7 defines the encoding of Label Mapping message as follows.



Message ID:

32-bit value used to identify this message.

FEC TLV:

Specifies the FEC component of the FEC-Label mapping being advertised. The [PWE3-Control] document specifies the FEC Element for the Pseudowire for the FEC TLV. This is commonly referred to as VC-FEC.

Label TLV:

Specifies the Label component of the FEC Label mapping.

Optional Parameters:

This variable length field contains 0 or more parameters, each encoded as a TLV.

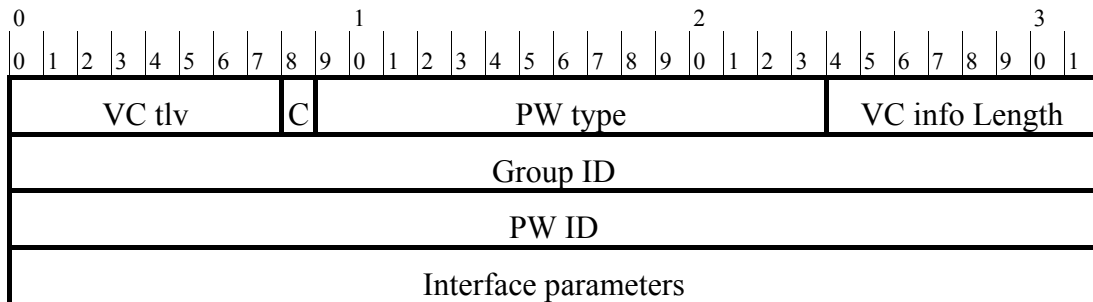
This Appendix specifies various TLVs that can be included in Optional Parameter field of Label Mapping Message as a mean to provide additional attributes for the Pseudowire.

When these TLVs are passed as an update to an existing VC-FEC, the Label Mapping message is sent with the same VC-FEC and the label that was advertised earlier.

Pseudo-Wire FEC

The FEC element type is 128. [note 1]

NOTE : Only a single PW FEC element must be advertised per LDP PW label. The pseudo wire FEC element, is defined as follows:



PW type

A 15 bit quantity containing a value which represents the type of VC. Assigned Values are:

PW type	Description
0x0001	Frame Relay DLCI
0x	Port mode
All other values	Reserved

Control word bit (C)

The highest order bit (C) of the PW type is used to flag the presence of a control word. For frame relay control word is always present. The bit is set 1.

VC information length

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

Group ID

An arbitrary 32 bit value which represents a group of PWs 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 PW 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.

PW ID


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

Interface parameters

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

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:



Length:

The Length field is defined as the length of the interface parameter including the parameter id and length field itself.


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. If this parameter does not match in both directions of a specific PW, that VC LSP must not be enabled.

Signalling PW status.

The PE devices use an LDP TLV to indicate status to their remote peers. This PW Status TLV contains more information than the alternative simple Label Withdraw message.

The format of the PW Status TLV is:

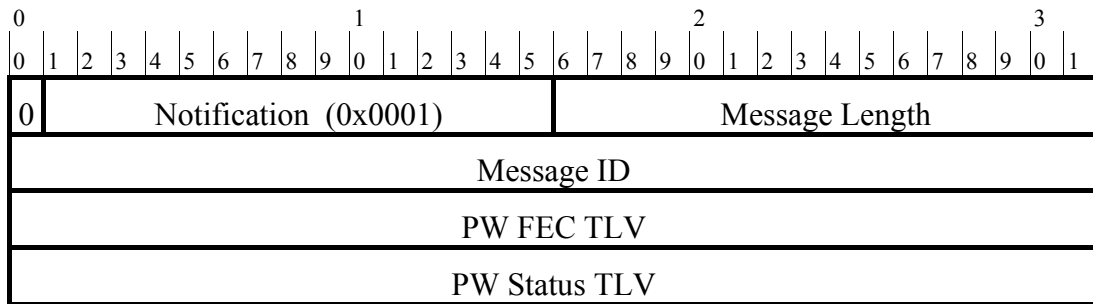


Where status is a 4 octet bit field is specified in the PW IANA Allocations document.

Each bit in the status code field can be set individually to indicate more than a single failure at once. Each fault can be cleared by sending an appropriate status message with the respective bit cleared. The presence of the lowest bit (PW Not Forwarding) acts only as a generic failure indication when there is a link-down event for which none of the other bits apply.

Notification message or status TLV

The Status TLV is transported to the remote PW peer via the LDP notification message. The format of the Notification Message is:



The PW FEC TLV should not include the interface parameters as they are ignored in the context of this message. When a PE's CE-facing interface encounters an error, use of the PW status message allows the PE to send a single status message, using a PW FEC TLV with only the group ID set, to denote this change in status for all affected PW connections.

As mentioned above the Group ID field can be used to send a status notification for all PWs associated with a particular group ID. This procedure is optional, and if it is implemented the LDP Notification message should be as follows: the PW information length field is set to 0, the PW 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 PW status notification procedure", and all PEs implementing this design are required to accept such a notification message, but are not required to send it.

LDP Label Withdrawal procedures

As mentioned above the Group ID field can be used to withdraw all PW 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 PW information length field is set to 0, the PW 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 PW label withdrawal message or release message. A wildcard release message must include only the group ID. A Label Release message initiated from the imposition router must always include the PW ID.

Label Mapping Release

In situations where the Imposition router (**Editor's note: terminology needs to aligned with egress and ingress PE**) wants to restart forwarding of packets with sequence number 1, the router shall 1) Send to disposition router a label mapping release, and 2) Send to disposition router a label mapping request. When sequencing is supported, advertisement of a PW label in response to a label mapping request must also consider the issues discussed in the section on Label Mapping Advertisements.
