

PWE3 Working Group
Internet Draft
Expires: March 2002

Claude Kawa
Nortel Networks
Andrew G. Malis
Vivace Networks, Inc.
Prayson Pate
Overture Networks, Inc.
Ravi Bhat
Nishit Vasavada
Amber Networks

September 2001

Frame relay over Pseudo-Wire Emulation Edge-to-Edge
draft-kamapabhava-fr-pwe3-00.txt

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are 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/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright(C) The Internet Society (2001). All Rights Reserved.

Abstract

This document defines frame relay pseudo-wire (PW) specific aspects. A frame relay PW allows frame relay protocol control information and payload to be carried over Packet Switched Networks (PSNs) using IP, L2TP, GRE or MPLS for transport.

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

Table of Contents

1.	Introduction.	2
2.	Terminology	3
3.	Acronyms and Abbreviations	4
4.	Requirements for Frame Relay Over PWE3	4
5.	Reference model	5
6.	General Encapsulation	5
7.	PVC maintenance protocol.	10
8.	Configuration prerequisites	17
9.	Frame Relay over MPLS.	18
10.	Frame Relay over IP	19
11.	Frame Relay over GRE.	20
12.	Frame Relay over L2TP.	21
13.	Security Considerations.	21
14.	References	21
15.	Acknowledgements	22
16.	Author's Addresses	22

Conventions used in this document

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](#).

[1.](#) Introduction

This document defines frame relay pseudo-wire (PW) specific aspects. A frame relay PW allows frame relay protocol control information and payload to be carried over Packet Switched Networks (PSNs) using IP, L2TP, GRE or MPLS for transport. A frame relay-pseudo wire is a mechanism that emulates the essential attributes of frame relay service over a PSN. The required functions of frame relay PWs include

- Encapsulating frame relay specific information receives at an ingress port of a Provider Edge (PE) device,

- Carrying them across a PSN for delivery to another PE device,
- Extracting or decapsulating frame relay specific information,
- Generating of native frame relay frames for forwarding through an egress port and,
- Performing any other operations required to support frame relay service.

2. Terminology

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](#). Below are the definitions for the terms used throughout the document.

Packet Switched Network:

A Packet Switched Network (PSN) is a network using IP, MPLS or L2TP as the unit of switching.

Pseudo Wire Emulation Edge to Edge:

Pseudo Wire Emulation Edge to Edge (PWE3) is a mechanism that emulates the essential attributes of a service (such as a T1 leased line or Frame Relay) over a PSN.

Customer Edge:

A Customer Edge (CE) is a device where one end of an emulated service originates and terminates. The CE is not aware that it is using an emulated service rather than a "real" service.

Provider Edge:

A Provider Edge (PE) is a device that provides PWE3 to a CE.

Pseudo Wire:

A Pseudo Wire (PW) is a connection between two PEs carried over a PSN. The PE provides the adaptation between the CE and the PW.

PW End Service:	A Pseudo Wire End Service (PWES) is the interface between a PE and a CE. This can be a physical interface like a T1 or Ethernet, or a virtual interface like a VC or VLAN.
Pseudo Wire PDU:	A Pseudo Wire PDU is a PDU sent on the PW that contains all of the data and control information necessary to provide the desired service.
PSN Tunnel:	A PSN Tunnel is a tunnel inside which multiple PWs can be nested so that they are transparent to core network devices.

3. Acronyms and Abbreviations

CE	Customer Edge
FR	Frame relay
FRoPW	Frame Relay over Pseudo Wire
LSP	Label switched path
LSR	Label switching router
MPLS	Multi-protocol label switching
MTU	Maximum Transfer Unit
NNI	Network-Network Interface
PE	Provider Edge
PSN	Packet Switched Network
PW	Pseudo-wire
POS	Packet over Sonet/SDH
PVC	Permanent virtual circuit
SPVC	Switched/Soft permanent virtual circuit
SVC	Switched virtual circuit
UNI	User-Network Interface
VC	Virtual circuit

4. Requirements for Frame Relay Over Pseudo-Wire Emulation Edge to edge

The section lists the main frame relay pseudo-wires requirements to be

met between two PEs:

- [1](#). Frame Transport: It must be possible to carry between two PEs both user and maintenance PSN traffic on the same pseudo-wire. It should be possible also to distinguish between the two types of traffic.
- [2](#). Frame Length: Must transport variable length frame relay frames without being limited by the PSN MTU.
- [3](#). Bi-directional traffic: Must support bi-directional traffic.
- [4](#). Frame ordering: Must preserve frame relay frames order.
- [5](#). Control bits: Must support the transport of frame Discard Eligibility (DE), Forward Explicit Congestion Notification (FECN), Backward Explicit Congestion Notification (BECN) and Command/Response (C/R) bits [[Q922](#)]
- [6](#). PVC Status Maintenance: Must support the mapping and transport of frame relay PVC status indications defined in Q933 Annex A [[Q933](#)]. The support of PVC continuity check should be provided.
- [7](#). Traffic Management: Should be able to map the following frame relay traffic management parameters to PWs traffic parameters:
 - a) CIR (Committed Information Rate) or throughput,
 - b) Bc (Committed burst size),
 - c) Be (Excess Burst Size),
 - e) Maximum frame size.

Kawa/Malis/Pate/Bhat/Vasavada

[Page 4]

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

[8](#). Frame Priority and QoS: Should support the ability to map different frame relay priorities or QoS classes as defined in [[X36](#), [X76](#)] [[X36](#), [X76](#)] to appropriately engineered PWs and tunnel PSNs.

[9](#). Frame relay VC type: Must support PVC, support of SVC and SPVC is optional.

[5](#). Reference model

Figure 1 shows PWE3 reference model with additional frame relay concepts. This section will be completed with additional details imported from [PWE3REQ and PWE3FW].

(To be completed)




Figure 2 - General format of FROPW packet

The delivery header is a header specific to the PSN, it is discussed in the following sections addressing each type of PSN. FROPW header contains protocol control information, its structure is shown in Figure 3. The packet payload field is the content of the information field of frame relay frame defined in ITU-T Recommendation Q.922 Annex A [Q922].

FROPW header structure is shown in Figure 3.




Figure 3 - FROPW header structure

The meaning of the fields is as follows:

Res (bits 0 to 2):

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

P - Payload Type (bit 3):

If set to zero then the payload field contains user's data else it contains network maintenance data.

F - FECN (bit 4):

FR FECN (Forward Explicit Congestion Notification) bit [Q922].

B - BECN (bit 5):

FR BECN (Backward Explicit Congestion Notification) bit [Q.922].

D - DE (bit 6)

FR DE bit indicates the discard eligibility [Q922].

C/R (bit 7)

FR frame C/R (command/response) bit [[0922](#)].

X, Y (bits 8 and 9):

1 0: First segment of a segmented FRoPW packet

0 1: Last segment of a segmented FRoPW packet

0 0: Neither the first nor the last segment of a segmented FRoPW packet

1 1: Complete FRoPW packet

X and Y bits are used for segmentation and re-assembly of FRoPW packet fragments when these capabilities are enabled in the two PEs terminating a PW.

Length (bits 10 to 15):

The length field is used to pad short FRoPW packets over Ethernet PSN. The length field is used as follows: If the length of the layer 2 frame (consisting of layer two control information and layer 2 payload) is less than 64 bytes, the length field MUST be set to the FRoPW packet length. Otherwise the length field MUST be set to zero. The value of the length field, if non-zero, is used to remove any padding.

Sequence number (Bit 16 to 31):

If not used it is set to zero by the sender and ignored by the receiver. Otherwise it specifies the sequence number of a complete FRoPW packet or a fragment. A circular list of sequence numbers is used. A sequence number takes a value from 1 to 65535 ($2^{16}-1$). Arithmetic modulo 2^{16} is used to generate a new sequence number.

The sequence number must be used when segmentation and re-assembly are enabled between two peer PEs terminating a PSN tunnel. The sequence number may be used to detect out-of-order delivery of FRoPW packets when the PSN does not guarantee in-order delivery.

[6.2](#). FRoPW packet processing

(Section and subsections to be completed)

[6.2.1](#). Generating FRoPW packets

The generation process of a FRoPW 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 FRoPW header from the corresponding fields of the frame relay frame: C/R, DE, FECN and BECN,
- The length field will be addressed in the next iteration of this draft.
- The sequence field is set according to the configuration parameters.

Details are provided below.

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

- The payload field is the contents of the frame relay information field stripped from any escape bit or character.

Once the FRoPW packet has been generated, additional processing specific to the type of PW used and the underlying PSN is performed.

[6.2.1.1](#). Setting the sequence number

If the two PEs terminating a PSN tunnel support packet sequencing capability then the following procedure to number FRoPW packets is used:

- The initial packet transmitted on the frame relay PW MUST use sequence number 1,
- For a subsequent packet, the sequence number corresponds to the sequence number of the preceding packet incremented by 1 modulo 2^{16} ,
- when the sequence number reaches the maximum 16 bit value (65535) the next sequence number wrap to 1.

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

[6.2.1.2](#). Sequence number and segmentation of long FRoPW packets (To be completed)

[6.2.2](#). Receiving FRoPW packets

(Section and subsections to be completed)

When a PE receives a FRoPW packet, it processes the different FRoPW header fields in order to synthesize a new frame relay frame for transmission through one of its frame relay UNI or NNI. The PE performs the following actions(not necessarily in the order shown):

Note - PSN and PW specific processing are specified in the section on each type of PSN.

- The PE checks the P bit to ensure that it is set to "user's data",
- It generates the following frame relay frame fields from the FRoPW header: C/R, DE, FECN and BECN,

- It processes the length and sequence field,
- It generates the frame relay information field from the contents of the FRoPW packet payload.

Once the frame relay frame has been generated, it is queued for transmission across the selected frame relay UNI or NNI.

[6.2.2.1](#) Checking the sequence number

If the receiving PE support packet sequencing, the processing is done as follows:

The "expected sequence number" of the first FRoPW packet is 1.

When a 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
- 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 a packet passes the sequence number check, it is kept for further processing in order to be forwarded to its next destination.

If the packet is in order, then the expected sequence number should be set using the following assignment:

```
expected_sequence_number := packet_sequence_number + 1 mod 2**16
if (expected_sequence_number = 0) then expected_sequence_number := 1;
```

Packets, which are received out of order, MAY be dropped or reordered at the discretion of the receiver.

If the receiving PE does not support sequence number processing, then

the sequence number field MAY be ignored.

[6.2.2.2](#) Sequence number and segmentation of long FRoPW packets

(To be completed)

[6.3.](#) Segmentation and Re-assembly procedures

(Section and subsections to be completed)

[6.3.1.](#) Segmentation procedure

Segmentation is performed when the information field of the frame relay frame received is too long for transmission across the PSN and when the configuration of the PSN tunnel and/or PW allows segmentation to be performed. X and Y bits are used to indicate the type of segment: First

Kawa/Malis/Pate/Bhat/Vasavada

[Page 9]

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

segment (X bit set to 1), last segment (Y bit set to 1) neither the first nor the last segment (X and Y bits set to 0). The sequence number is used to number the fragments. If packet sequencing is used for every packet, numbering FRoPW packet fragments follows the same procedure as numbering complete packets: For every packet fragment, the sequence number is the sequence number of the previous packet/fragment incremented by one modulo $2^{*}16$. When packet sequencing is not used for every packet but only with segmentation and re-assembly, the sequence number is re-initialized for every long packet that requires segmentation. The first FRoPW packet fragment has sequence number and for each subsequent fragment, the sequence number is incremented by 1 modulo $2^{*}16$.

If segmentation is not allowed and the frame relay frame is too long for transmission in a single FRoPW packet, it shall be discarded.

(Detailed procedures will be provided in the next version of this draft)

[6.3.2.](#) Re-assembly procedure

When a frame relay frame has been segmented by the sending PE and segmentation/re-assembly is allowed by the receiving PE, the receiving PE will re-assemble the frame relay segments received in FRoPW packets to regenerate the original frame relay information field.

(Detailed procedures will be provided in the next version of this draft)

[6.4.](#) Handling of error conditions

Figure 4 - PVC maintenance message format

Meaning of each field:

Bits 0 to 3: Version number (V)

The version number of the PVC maintenance protocol.

This version is version 0001.

Bits 4 to 11: Message type

This field identifies a PVC maintenance protocol message. The following messages are defined:

0000 0000: PVC status indication

0000 0001: PVC status response

0000 0010: Keep alive indication

0000 0011: Keep alive response

Other values are reserved for future use.

Bits 12 to 15: PVC state

This field indicates the operational state of a PVC as follows:

A (active) bit

This bit indicates whether the PVC is active (value 1) or inactive (value 0).

N (new) bit

This bit indicates whether the status reporting is for a new PVC (value 1) or an existing PVC (value 0).

D (delete) bit

This bit indicates that a PVC will be deleted when it is set to 1. D bit set to 0 is used with other PVC status announcement.

Bits 16 to 23: Message correlation number (MSN)

This field allows the correlation of a PVC status indication message with the corresponding PVC status response or a keep alive indication message with the corresponding keep-alive response message.

Kawa/Malis/Pate/Bhat/Vasavada

[Page 11]

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

Note on the use of the Message correlation number field:

The field value is an integer between 0 and 255. A simple way to generate Message correlation numbers is to use a counter modulo 256 per frame relay PW. Before transmitting a new PVC status /keep-alive indication message, a PE inserts a new value of the counter in the Message correlation number field. A receiving PE returns in a PVC status /keep-alive response message the value of the message correlation number received in the corresponding PVC status/keep-alive indication message.

Bits 24 to 31: Reserved

These bits are reserved for future use.

[7.2.](#) Execution of PVC maintenance protocol

PVC maintenance functions may be executed to report a PVC state change or for keep-alive purposes in the absence of user's traffic. Each side of a related pair of PSN tunnel has a timer T1 to indicate when PVC maintenance protocol functions are executed. However, any time before T1 expires a PE may send a PVC status indication message to inform the other side of a PVC state change in a timely manner. But keep-alive functions are executed only when T1 expires. Note timer T1 applies to a pair of bi-directional PSN tunnel and not to individual nested frame relay PW.

[7.3.](#) Notification of a newly created PVC

A PE shall send a PVC status indication message to notify its peer that a PVC has been created. The N bit is set to 1 (new PVC). If the PE is ready to send and receive user's frames it shall set the A bit to 1 (active) otherwise it shall set it to 0. The PE shall insert a Message correlation number as described in [Section 7.1](#). The other fields are set as shown in Figure 1 and described in [Section 7.1](#).

If the PE does not receive a reply when timer T1 expires it shall retransmit the PVC status message with the N bit set to 0 (new) and the A bit set according to the availability of the PVC with a new Message correlation number. At any time there is only one outstanding PVC status indication message per frame relay PW. The PVC status message may be retransmitted up to N1 time (the default value of N1 is 255 retransmissions) but only after T1 expires. After N1 retransmissions the management system is notified of the lack of reply and the PE may continue with the retransmission of the PVC status indication message until further notification by the management system.

When a PVC status response message is received, the PE receiving the PVC status response message checks the various fields to ensure that they are correct. In particular, the Message correlation number must correspond to the Message correlation number of the last transmitted PVC status indication message awaiting a reply, otherwise the PVC status reply message is silently discarded.

The receiver checks also the N bit and A bit. The N bit must be set to 1

(new PVC) otherwise the management system must be notified and the PVC must be unavailable for use until further notification by the management system. If the A bit is set to 0 (inactive) no user frames may be

transmitted or received. If the A bit is set to 1 (active) transmission of user frame is allowed only if the PVC is active at this side too.

Transmission in the two directions is allowed only when the PVC is active at the two ends of the VC LSPs. Any user's frame received when the PVC is not active at the two PEs shall be discarded.

[7.4.](#) Replying to the notification of a newly created PVC

When a PE receives from its peer a PVC status indication message indicating that a PVC has been created, it shall reply by returning a PVC status response message as follows:

- [1.](#) If the PVC has been created at this side, the N bit shall be set to 1
- [2.](#) (new PVC).
- [3.](#) If the A bit received in the PVC status indication message was set to [1](#) (active) and the PVC has been activated at this side, the A bit shall be set to 1 otherwise it shall be set to zero (inactive).

[7.5.](#) Notification of PVC activation

A PE shall send a PVC status indication message to notify its peer that a PVC is active at its side and ready to send and receive user's data. The PE shall set the N bit to 0 (existing PVC) and the A bit to 1 (active). The other fields are set as explained in [Section 7.1](#).

If the PE does not receive a reply when timer T1 expires it shall retransmit the PVC status message with the N bit set to 0 and the A bit set to according to the activation state of the PVC at this side. The PVC status message may be retransmitted up to N1 time. After N1 retransmission the management system is notified of the lack of a reply and the PE may continue to retransmit the PVC status message to notify the other side that the PVC is active until further notification by the management system.

When a reply is received indicating that the PVC is active, user frames may be sent in both directions. Otherwise, any user's frame received it will be discarded.

[7.6.](#) Replying to the notification of PVC availability

When a PE receives from its peer a PVC status indication message indicating that a PVC is available (active), it shall reply by returning a PVC status response message as follows:

- [1.](#) If the PVC has been created at this side and is active, the N bit

shall be set to 0 and the A bit to 1 (active)

[2.](#) If the PVC has been created and is not active it shall set the N bit to 0 and the A bit to 0 (inactive).

If the identification of the frame relay PW does not correspond to a configured PVC, the PVC status indication message received is discarded and no reply is returned.

[7.7.](#) Notification of PVC de-activation

A PE shall send a PVC status indication message to notify its peer that a PVC is inactive at its side. The PE shall set the N bit to 0 (existing PVC) and the A bit to 0 (inactive). The other fields are described in [Section 7.1](#).

If the PE does not receive a reply when timer T1 expires it shall consider the PVC inactive and any user's frame received shall be discarded. The PVC status message is retransmitted up to N1 time. After N1 retransmissions the management system is notified of a lack of a reply.

If a user frame is received when a PVC is inactive, it shall be discarded and the PE may retransmit the PVC status message to notify the other side that the PVC is inactive

When a reply is received indicating that the PVC is inactive, user frames must not be sent in any direction until the PVC has been activated at the two sides. Any user's frame received when the PVC is received will be discarded and the receiving side may return a PVC status message indicating that the PVC is still inactive.

[7.8.](#) Replying to the notification of PVC de-activation

When a PE receives from its peer a PVC status indication message indicating that a PVC is unavailable (inactive), if the PVC is configured, it shall reply by returning a PVC status message with the N bit set to 0 and the A bit to 0 (inactive).

If the identification of the frame relay PW does not correspond to a configured PVC, the PVC status indication message received is discarded and no reply is returned.

No user's frames may be sent when the PVC is inactive. Any user's frame received is discarded.

[7.9.](#) PVC status checking

At the expiration of timer T1 a PE may send a PVC status indication message to the other PE with the PVC state bits (D, N and A) set to appropriate values to elicit a response from the other PE about the PVC

state. This procedure is useful, for example, when a PVC has been de-activated for too long and one of the PE wants to check the PVC state at the other side.

[7.10.](#) Notification of the deletion of a PVC

When a PVC is in the process of being deleted, a PE notifies the other side by sending a PVC status indication message with the D bit set to 1, the N bit must be set to 0 (existing PVC) and the A bit to 0 (inactive). One or both sides may initiate the PVC deletion notification procedure. When both sides initiate the procedure, each side must reply to the PVC status indication message sent by the other side. If no reply is received, the PVC status indication message may be retransmitted up to N3 times. After N3 retransmissions the management system is notified of a lack of a reply and the PVC is considered deleted.

When the PE receives a PVC status response message replying to the notification of the deletion of a PVC, it shall ensure that the Message correlation number is valid and the D bit received is set to 1, otherwise it shall discard the PVC status response message. If the message received is correct and after performing any other book-keeping function at this side, the PVC is considered deleted at this side.

[7.11.](#) Reply to the notification of the deletion of a PVC

When a PE receives from its peer a PVC status indication message indicating that a PVC is in the process of being deleted, it shall return a PVC status response message with the D bit set to 1, the N bit set to 0 and the A bit to 0. The other fields are set according to Figure 1 and Table 1. After sending the PVC status response message and after performing any other book-keeping function, the PVC shall be considered deleted at this side.

[7.12.](#) Keep-alive function

The keep-alive procedure is optional to implement, it is performed for each frame relay PW in the absence of user's or PVC status reporting traffic in one or both directions and does not induce a PVC state change in a PE. It merely allows a PE to inform the other PE that it is

functioning.

The keep alive procedure is executed by the PEs when timer T1 expires. A PE does not have to execute the keep-alive procedure every time timer T1 expires, in particular under heavy load or for any other reason.

[7.12.1](#). Sending a keep alive request message

A PE may initiate the keep-alive procedure when one of the following conditions is met:

[1](#). There has not been any user's traffic in one or both directions and the PVC is in the active state when timer T1 expires,

Kawa/Malis/Pate/Bhat/Vasavada

[Page 15]

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

[2](#). There is no PVC status reporting message to send to the other PE applicable to the corresponding VC LSP,

[3](#). For other implementation-specific reasons.

When one of the above conditions is met, a PE initiates the execution of the keep-alive procedure by sending a keep-alive indication message to the other PE indicating the PVC state at its side. The PVC state bits are set according to the PVC state at the PE side.

When a keep-alive response message is received. The PVC state bits are checked to determine if there is a state mismatch. When a state mismatch exists, the PE receiving the keep-alive response message shall execute the PVC status reporting procedures to bring the PVC to the same state at the two PEs. The PVC state bits are set to the lowest common denominator (detailed will be provided in the next version).

The PE shall take no special action if no reply to the keep-alive indication message is received at or before the next expiration of timer T1. However if there is no activity on the frame relay PWs in the transmitting and receiving directions after N2 retransmissions of the keep-alive indication message and no reply to the keep-alive indication message has been received, the network management system shall be informed. The PE may continue executing the PVC maintenance protocol until it receives a notification from the management system.

[7.12.2](#). Replying to a keep alive indication message

When a PE receives a keep-alive indication message it returns a keep-alive response indicating the state of the PVC at its side, if the PVC state at its side is consistent with the PVC state received in the keep-

alive indication message. If the PVC state is inconsistent the PE shall execute the PVC status reporting procedure without returning the keep-alive response message in order to bring the PVC in the same state at the two sides. The PVC state bits are set to the lowest common denominator (details will be provided in the next version).

[7.13.](#) Resumption of PVC activities after a failure

After a failure (link layer, transmission facility or internal failure) a PE shall execute the PVC status activation to return the PVC to the active state between the two PEs.

[7.14.](#) Handling of error conditions

As a general rule, if a PVC receives a PVC maintenance message with erroneous information, it shall discard it. In particular:

- [1.](#) If a PE receives an unexpected PVC maintenance message,
- [2.](#) If a PE received a PVC maintenance message with a content not as specified in [Section 7.1](#),

Kawa/Malis/Pate/Bhat/Vasavada

[Page 16]

Internet Draft

[draft-kamapabhava-fr-pwe3-00.txt](#)

September 2001

- [3.](#) If a PE receives a PVC maintenance message with the N bit set to 0 (existing PVC) and a frame relay PW identification not recognized as belonging to a configured PVC,

- [4.](#) If a PE receives a PVC maintenance status response message with an invalid Message correlation number.

[7.15.](#) PVC maintenance protocol Parameters

Timers:

Timer T1 - This timer exists at each peer PE for a pair of bi-directional. tunnel LSP. It controls the periodical execution of the PVC maintenance protocol.

The default value of timer T1 is 1 second.

Counter Thresholds:

N1: This threshold is the maximum number of retransmissions of PVC status indication before informing the management system of any problem. The default value is 255 retransmissions.

N2: This threshold is the maximum number of retransmissions of the PVC keep-alive indication message before informing the management system of any problem. The default value is 255 retransmissions.

N3: This threshold is the maximum number of retransmissions of the PVC status indication message with the D bit set to 1 (PVC to be deleted) before considering the PVC to be deleted and before informing the management system. The default value is 15 retransmissions.

8. Configuration prerequisites

Since frame relay VCs are bi-directional and carry traffic in the two opposite directions, at least one pair of tunnel PSN must be available between two PEs with appropriate traffic and QoS capabilities before they can start sending FRoPW packets. Some PSN tunnels require to be created and maintained, other do not. Establishing, maintaining and releasing PSN tunnel are outside the scope of this document. Binding between a pair of frame relay interfaces/DLCI and a PW is done when the PW and FR VC segments are created.

In addition a PE must be configured with the following parameters per tunnel PSN. These parameters will apply to all Frame Relay PWs nested inside the tunnel PSN:

1. Maximum transfer unit (MTU) of the PSN,
2. Whether segmentation and re-assembly will be supported or not,
3. Use of the sequence number: With segmentation only, always or never,
4. Additional configuration parameters specific to each PW and PSN are listed in the section covering each PSN.

9. Frame Relay over MPLS

9.1. MPLS tunnel and VC LSPs

MPLS label switched paths (LSPs) called "Tunnel LSPs" establish an association between a pair of PEs. A tunnel LSP correspond to a "PSN tunnel" of Figure 1. Several "Virtual Circuit LSPs" (VC LSPs) may be nested inside one Tunnel LSP. Each VC LSP carries the traffic of a frame relay PVC or SVC in one direction. Since LSPs are uni-directional, a pair of VC LSPs (and corresponding tunnel LSPs) carrying traffic in opposite directions will be required usually for each frame relay PVC or SVC.

9.2. Relationship between FR VCs and MPLS VC LSPs

Frame Relay VCs are considered to be bi-directional objects mainly because of the way they are created and identified. A single frame relay identifier (DLCI) refers to the two directions of a frame relay VC and frame relay signalling establishes the two directions simultaneously with the same message flows. In general each direction of a frame relay VC may have different traffic and 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.

In the case of frame relay over MPLS, during the creation of a frame relay VC, a pair of VC LSPs will have to be established between two PEs. For one end-to-end frame relay VC two VC LSPs exist: One VC LSP to transport the traffic from PE1 to PE2 and another one to transport the traffic in the opposite direction. The VC LSP in each direction must comply with the characteristics of the corresponding direction of the frame relay VC.

The VC LSP from PE1 to PE2 is the transmit VC LSP for PE1 and the receive LSP for PE2. Likewise, the VC LSP from PE2 to PE1 is the transmit LSP for PE2 and the receive LSP for PE1.

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.

In PE1 to PE2 direction a tunnel LSP gets MPLS packets from PE1 to PE2, the corresponding label is not related to any frame relay VC. When PE1 sends a FROPW packet to PE2, it first pushes a VC label on its label stack, and then pushes on a tunnel LSP label. The VC label is not visible until the FROPW packet reaches PE2. PE2 forwards FROPW packet based on the LSP VC label. The VC label must be always at the bottom of the label stack. While the packet is transported across the MPLS network, additional labels may be pushed on (and then popped off) as needed. The processing is similar in the opposite direction from PE2 to PE1.

[9.3](#). Frame relay over MPLS packet format

In the case of frame relay over MPLS, the generic FROPW packet format of Figure 2 becomes as shown in Figure 5.

```
+-----+
|      Tunnel LSP label(s)      | n words (4n bytes per label)
```

+-----+		
	VC LSP label	1 word
+-----+		
	FRoPW Header (See Figure 3)	1 word
+-----+		
	Payload packet	
	(Q.922 frame information field)	Maximum N bytes
		(to be specified)
+-----+		

Figure 5 - Frame relay over MPLS packet format

{Section to be completed}

[10.](#) Frame Relay over IP

{Section and subsections to be completed}

Note both IP v4 and IP v6 are supported.

[10.1.](#) General aspects of frame relay over IP


[10.2](#) Frame relay over IP packet format

For frame relay over IP v4 or v6, the packet format is shown in Figure 6.

+-----+		
	IP v4 or v6 header	n words
+-----+		
	Frame relay PW identifier (See Figure 7)	1 word
+-----+		
	FRoPW Header (See Figure 3)	1 word
+-----+		
	Payload packet	
	(Q.922 frame information field)	Maximum N bytes
		(to be specified)
+-----+		

Figure 6 - Frame relay over IP packet format

Frame relay PW identifier field shown in Figure 6 is used to identify a frame relay VC between two PEs in the case of frame relay over IP. The format of this field is shown in Figure 7. Actually the format is similar to MPLS LSP label format. This format allows to have one size for identifying a frame relay VC between two PEs, instead of having two formats as allowed in FRF.1.2 and FRF.2.2 (see [Q922, FRF1 and FRF2]). The use of the other fields will be described in a subsequent version of this draft.



```
Identifier:  Frame relay PW identifier, 20 bits
Exp:         Experimental Use, 3 bits
TTL:         Time to Live, 8 bits
```

Figure 7 - frame relay PW identifier for frame relay over IP

10.3. Frame relay over IP specific packet processing

(To be completed).

10.4. Additional configuration objects for frame relay over IP

(To be completed).

11. Frame Relay over GRE

(Section and subsections to be completed).

11.1. General aspects of frame relay over GRE

11.2. Frame relay over GRE packet format

For frame relay over GRE [RFC2784, [RFC2890](#)], the packet format is similar to frame relay over IP packet format shown in Figure 6, except that there is one or more additional GRE header(s) as shown in Figure 8.

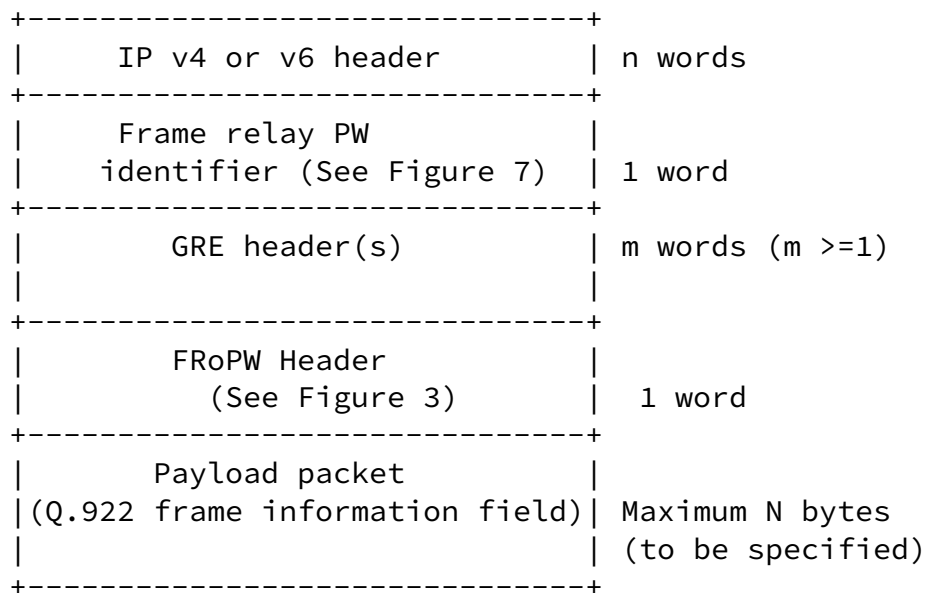


Figure 8 - Frame relay over GRE packet format

[11.3.](#) Frame relay over GRE specific packet processing[11.4.](#) Additional configuration objects for frame relay over GRE[12.](#) Frame Relay over L2TP

(To be completed).

[13.](#) Security Considerations

(To be completed).

[14.](#) References

- [I233] ITU-T Recommendation I.233.1, ISDN frame relay bearer service, Geneva, October 1991

- [FRF1] FRF.1.2, Frame relay PVC UNI Implementation Agreement, Frame Relay Forum, April 2000
- [FRF2] FRF.2.2, Frame relay PVC UNI Implementation Agreement, Frame Relay Forum, <month> 2001
- [FRF4] FRF.4.1, Frame relay SVC UNI Implementation Agreement, Frame Relay Forum, January 2000

Kawa/Malis/Pate/Bhat/Vasavada

[Page 21]

Internet Draft [draft-kamapabhava-fr-pwe3-00.txt](#) September 2001

- [FRF10] FRF.10.1, Frame relay SVC NNI Implementation Agreement, Frame Relay Forum, January 2000
- [PWE3REQ] XiPeng Xiao, et al., Internet draft, [draft-ietf-pwe3-requirements-01.txt](#), July 2001
- [PWE3FW] Prayson Pate, et al., Internet draft, Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3)
- [Q922] ITU-T Recommendation Q.922, ISDN Data Link Layer Specification for Frame Mode Bearer Services, ITU, Geneva, 1992.
- [Q933] ITU-T Recommendation Q.933, ISDN Signaling Specification for Frame Mode Bearer Services, Geneva, October 1995.
- [RFC3032] E. Rosen, et al., [RFC 3032](#) MPLS Label Stack encoding, January 2001.
- [RFC3031] E. Rosen, et al., [RFC 3031](#) MPLS Architecture, January 2001.
- [RFC2615] A. G. Malis, [RFC 2615](#) PPP over SONET/SDH, June 1999.
- [RFC2784] D. Farinacci, et. al, [RFC 2784](#) Generic routing encapsulation, March 2000
- [RFC2890] G. Dommety, [RFC 2890](#) Key and sequence number extensions to GRE, September 2000
- [X36] ITU-T Recommendation X.36, Interface between a DTE and DCE for public data networks providing frame relay, Geneva, 2000
- [X76] ITU-T Recommendation X.76, Network-to-network interface between public data networks providing frame relay

[15](#). Acknowledgements

To be completed.

[16](#). Author's Addresses

Claude Kawa
Nortel Networks
[3500](#) Carling Ave.
Ottawa, Ontario, Canada
email: kawa@nortelnetworks.com

Kawa/Malis/Pate/Bhat/Vasavada

[Page 22]

Internet Draft [draft-kamapabhava-fr-pwe3-00.txt](#) September 2001

Andrew G. Malis
Vivace Networks, Inc.
[2730](#) Orchard Parkway
San Jose, CA 95134
e-mail: Andy.Malis@vivacenetworks.com

Prayson Pate
Overture Networks
[P.](#) O. Box 14864
RTP, NC, USA 27709
email: prayson.pate@overturenetworks.com

Ravi Bhat
Amber Networks
[50](#) Vanivalas Road
Bangalore 560 004 India
email: rbhat@ambernetworks.com

Nishit Vasavada
Amber Networks
<address>
email: nishit@ambernetworks.com

