

PWE3 Working Group
Internet Draft
Expires: October 2002

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

April 2002

Frame Relay Transport and Encapsulation over Pseudo-Wires
draft-kamapabhava-fr-pwe3-01.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/1id-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 specific aspects. A frame relay pseudo wire exists between network edge nodes and support as faithfully as possible an instance of a frame relay service. A frame relay pseudo-wire is a mechanism that allows frame relay protocol control information and payload to be carried over IP or MPLS Packet Switched Networks.

Table of Contents

1. Introduction.	2
2. Terminology	2
3. Acronyms and Abbreviations	3
4. Requirements for Frame Relay Over PWE3	3
5. Reference model	4
6. General Encapsulation and FRoPW packet processing	5
7. Configuration prerequisites	10
8. Frame Relay over MPLS PSN	10
9. Frame Relay over IP PSN.	11
10. Use of L2TP with frame relay pseudo-wires over IP PSN	12
11. Frame relay port to port mode	13
12. Security Considerations.	13
13. References	13
14. Acknowledgements	14
15. Author's Addresses	14

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 IP or MPLS Packet Switched Networks (PSNs) using MPLS or L2TP for multiplexing as specified in the framework draft [[PWE3FW](#)]. A frame relay PW is a mechanism that emulates the essential attributes of frame relay service over a PSN. The required functions to support frame relay PW by a network edge node include:

- Encapsulation of frame relay specific information in a suitable frame relay over pseudo wire packet,
- Transfer of this information across a PSN for delivery to a peer edge, node,
- Extraction of frame relay specific information by the remote edge node,
- Generation of native frame relay frames for forwarding across an egress port of the remote PE device,
- Execution 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
FCS	Frame Check Sequence
FR	Frame Relay
FRoPW	Frame Relay over Pseudo Wire
LSP	Label Switched Path
LSR	Label Switching Router
MPLS	Multiprotocol 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 edge nodes (known as provider edges (PEs)):

- 1. Frame Transport: It must be possible to carry between two PEs both user and maintenance traffic on the same pseudo-wire. It must be possible also to distinguish between the two types of traffic.**

- 2. Frame Length:** Must transport variable length FR frames without being limited by the PSN MTU.
- 3. Bi-directional traffic:** Must support bi-directional traffic.
- 4. Frame ordering:** Must preserve FR frames order.
- 5. Control bits:** Must support the transport of FR 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. Note PVC status maintenance will be addressed in another IETF draft.
- 7. Traffic Management:** Should be able to map the following FR 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.
- 8. Frame Priority and QoS:** Should support the ability to map different FR priorities or QoS classes as defined in [[X36](#),[X76](#)] to appropriately engineered PWS and tunnel PSNs.
- 9. Frame relay VC type:** Must support PVC, support of SVC and SPVC is optional and will be addressed in the future.

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] and specific information related to frame relay.



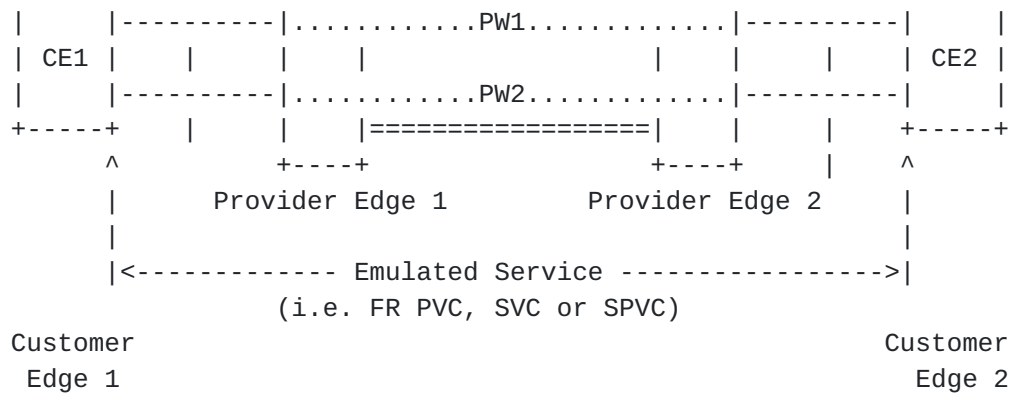


Figure 1 - PWE3 Reference Model

6. General Encapsulation and FRoPW packet processing

6.1. General Encapsulation

The general Frame Relay over Pseudo Wires (FRoPW) packet format to carry frame relay information (user's payload and frame relay control information) from one PE to another is shown in Figure 2.




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 used by the PEs.

B - BECN (bit 4):

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

F - FECN (bit 5):

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

D - DE (bit 6)

FR DE bit indicates the discard eligibility [[Q922](#)].

C/R - Command/Response (bit 7)

FR frame C/R (command/response) bit [[Q922](#)].

I, L - Initial and Last fragment indicators (bits 8 and 9):

0 1: First fragment of a fragmented FRoPW packet

1 0: Last fragment of a fragmented FRoPW packet

0 0: Complete FRoPW packet

1 1: Neither the first nor the last fragment of a fragmented FRoPW packet

I and L bits are used for fragmentation and re-assembly of FRoPW packet 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 when Ethernet links are used in the PSN. The length field is used as follows: If the length of a frame relay frame (consisting of control information and 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 it is 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 value of zero indicates that the sequence number field is not used.

The sequence number must be used when fragmentation 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

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 as follows:
 - C/R bit is copied unchanged in the FRoPW header.
 - DE bit is copied as follows into the FRoPW header: If it was set to one in the incoming frame relay frame, it must be copied unchanged in the FRoPW header. Otherwise if it was set to zero, it may be set to zero or one, depending on the traffic policing performed by the

PE device.

- FECN is copied as follows into the FRFoPW header: If it was set to one in the incoming frame relay frame, it must be copied unchanged in the FROPW header. Otherwise if it was set to zero, it may be set to zero or one, depending on the congestion state of the PE device in the forward direction.
- BECN follows the same processing rules as FECN, except that it applies to the backward direction.

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

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 the sequence number 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.

6.2.2. Receiving FRoPW packets

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 across one of its frame relay UNI or NNI. The PE performs the following actions (not necessarily in the order shown):

- The PE checks the P bit to ensure that it is set to "user's data",
- It generates the following frame relay header fields from the corresponding fields of the FRoPW packet as follows:
 - C/R bit is copied unchanged in the frame relay header.
 - DE bit is copied as follows into the frame relay header: If it was set to one in the incoming FRoPW packet, it must be copied unchanged in the frame relay header. Otherwise if it was set to zero, it may be set to zero or one, depending on the traffic policing performed by the PE device.

- FECN is copied as follows in the frame relay header: If it was set to one in the incoming FRoPW packet, it must be copied unchanged in the frame relay header. Otherwise if it was set to zero, it may be set to zero or one, depending on the congestion state of the PE device in the forward direction.
- BECN follows the same processing rules as FECN, except that it applies to the backward direction.

- It processes the length and sequence field, the details are in the subsequent sub-sections.
- It generates the frame relay information field from the contents of the FROPW packet payload and retrieves the appropriate DLCI.

Once the frame relay frame has been generated, it is queued for transmission across the selected frame relay UNI or NNI after the FCS and HDLC flags have been added and any bit or byte stuffing has been performed.

6.2.2.1 Checking the sequence number

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

When a FROPW 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 FROPW packet passes the sequence number check, it is further processed in order to be forwarded to its next destination.

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

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

FROPW packets, which are received out of order are dropped. Packet re-ordering by the receiver is not supported in this version of the draft.

If the receiving PE does not support sequence number processing, then the sequence number field is ignored. Note support of the sequence number capability is agreed upon by the two PEs either by configuration or by signalling prior to any packet exchange.

6.2.2.3 Processing of the length field by the receiver

(To be completed)

Kawa/Malis/Pate/Bhat/Vasavada

[Page 8]

6.3. Fragmentation and Re-assembly procedures

6.3.1. Fragmentation procedure

Fragmentation is performed when a FRoPW packet is too long for transmission across the PSN and when the configuration of the PSN tunnel and/or PW allows fragmentation to be performed. I and L bits of a FRoPW fragment are used to indicate the type of fragment:

First fragment (I = 0 and L = 1), last fragment (I = 1 and L = 0) neither the first nor the last fragment (I = L = 1). The sequence number is used to number the fragments. If packet sequencing is used for every FRoPW packet, the numbering of FRoPW fragments follows the same procedure as the numbering of complete packets: For every fragment, the sequence number is the sequence number of the previous FRoPW packet/fragment incremented by one modulo 2^{16} .

When packet sequencing is not used for every packet but only with fragmentation and re-assembly, the sequence number is re-initialized for every long frame relay frame that requires fragmentation. The first FRoPW packet fragment has a sequence number equal to 1 and for each subsequent fragment, the sequence number is incremented by **1 modulo 2^{16}** .

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

6.3.2. Re-assembly procedure

When a frame relay frame has been fragmented into multiple FRoPW fragments by the sending PE and fragmentation/re-assembly is supported by the receiving PE, the receiving PE will re-assemble the FRoPW fragments to synthesize the original frame relay frame.

6.4. Handling of error conditions

To be completed.

6.5. FR SVC and SPVC Signalling between PEs

Not supported in this document.

6.6. FR PVC provisioning

Provisioning of frame relay VCs requires the following actions: Frame relay VC segments (see Figure 1) are provisioned between each CE and the corresponding PE. A FR PW is established between the two PEs to complete the Frame Relay VC.

The two PEs terminating a frame relay PW executed a PVC maintenance

protocol to exchange PVC status information and for "keep-alive" purposes. The PVC maintenance protocol will be defined in another document.

7. 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 exchanging 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 fragmentation and re-assembly will be supported or not,**
- 3. Use of the sequence number: With fragmentation only, always or never,**
- 4. Additional configuration parameters specific to each PW and PSN are listed in the section covering each PSN.**

8. Frame Relay over MPLS PSN

8.1. 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.

8.2. 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 exists: 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 possible different values identify the pair of VC LSPs, one label value for each LSP.

In PE1 to PE2 direction a tunnel LSP transmits 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.

8.3. FRoPW packet format over MPLS PSN

In the case of frame relay over MPLS PSN, the generic FRoPW packet format of Figure 2 becomes as shown in Figure 4.

+-----+	
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 4 - Frame relay over MPLS packet format

(Section to be completed)

9. Frame relay over IP PSN

{Section and subsections to be completed}

Note both IP v4 and IP v6 are supported.

9.1. General aspects of frame relay over IP PSN

9.2 Frame relay over IP PSN FRoPW packet format

Two FRoPW packet format used over IP (v4 or v6) PSN are defined. One when MPLS is not used for switching and the other, when MPLS is used. When MPLS is not used, FRoPW packet format is shown in Figure 5.

+-----+	
IP v4 or v6 header	m words
+-----+	
FR PW identifier	
(See Figure 6)	1 word
+-----+	
FRoPW Header	
(See Figure 3)	1 word
+-----+	
Payload packet	
(Q.922 frame information field)	Maximum N bytes
	(to be specified)
+-----+	

Figure 5 -FRoPW packet format over IP PSN without MPLS

9.4. Additional configuration objects for frame relay over IP PSN

(To be completed as required).

10. Use of L2TP with frame relay pseudo-wires over IP PSN

When the PSN is an IP PSN, L2TP may be used as the multiplexing protocol as per the framework document [[PWE3FW](#)].

(To be completed).

11. Frame relay port to port mode

Frame relay port to port mode of operation applies to both IP and MPLS PSN. This mode provides transport of frame relay frames encapsulated in their entirety, including the address (with the control bits and DLCI) and information fields, but excluding the flags and the FCS. Bit or byte stuffing is undone.

Frame relay control bits (F, B, D and C bits) carried in FRoPW header are not used and must be set to 0 when transmitting ignored upon receipt. It must be noted that this mode is transparent to the FECN, BECN and DE bits; these bits are carried unchanged by the sending PE.

Frame relay port mode is an OPTIONAL capability.

12. Security Considerations

(To be completed).

13. 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

- [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.

- [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 services, Geneva, 2000

14. Acknowledgements

To be completed.

15. Author's Addresses

Claude Kawa
Nortel Networks
3500 Carling Ave.
Ottawa, Ontario, Canada
email: kawa@nortelnetworks.com

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@nokia.com

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

