

Internet Draft

Document: [draft-ietf-pwe3-requirements-00.txt](#)

Expires: November 17, 2001

XiPeng Xiao

Photuris Inc.

Danny McPherson

Amber Networks

Prayson Pate

Overture Networks, Inc.

Craig White

Level 3 Communications, LLC.

Kireeti Kompella

Juniper Networks, Inc.

Requirements for  
Pseudo Wire Emulation Edge-to-Edge (PWE3)  
[draft-ietf-pwe3-requirements-00.txt](#)

## Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC 2026](#). 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>.

## Abstract

This document describes generic requirements for Pseudo-Wire Emulation Edge to Edge. It provides guidelines for other working group documents that will define mechanisms for providing pseudo-wire emulation of specific services such as Ethernet, PPP, (Cisco) HDLC, ATM, Frame Relay and TDM. It should be noted that the PWE3 Working Group (WG) standardizes mechanisms that can be used to provide pseudo-wire emulation services, but not the services themselves.

## Copyright Notice

Internet Draft

[draft-ietf-pwe3-requirements-00](#)

May 17, 2001

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

Internet Draft

[draft-ietf-pwe3-requirements-00](#)

May 17, 2001

## Table of Contents

<a href="#">1</a>	Introduction .....	<a href="#">4</a>
<a href="#">1.1</a>	Reference Model of Pseudo-Wire Emulation Edge to Edge (PWE3) .....	<a href="#">4</a>
<a href="#">1.2</a>	Terminology .....	<a href="#">5</a>
<a href="#">2</a>	PDU Encapsulation .....	<a href="#">6</a>
<a href="#">2.1</a>	Conveyance of Necessary L2/L1 Header Information .....	<a href="#">6</a>
<a href="#">2.2</a>	Location of the Pseudo-Wire Header .....	<a href="#">6</a>
<a href="#">2.3</a>	PDU Length .....	<a href="#">6</a>
<a href="#">2.4</a>	Encapsulation of Signaling Messages of the Native Services .....	<a href="#">7</a>
<a href="#">2.5</a>	Timing Information .....	<a href="#">7</a>
<a href="#">2.6</a>	Support for Circuit Multiplexing .....	<a href="#">7</a>
<a href="#">2.7</a>	Packet Ordering .....	<a href="#">7</a>
<a href="#">3</a>	Packet Transport .....	<a href="#">7</a>
<a href="#">3.1</a>	Setup of Pseudo-Wires .....	<a href="#">7</a>
<a href="#">3.2</a>	Pseudo-Wire MTU .....	<a href="#">8</a>
<a href="#">3.3</a>	Transport of Signaling Messages of the Native Services .....	<a href="#">8</a>
<a href="#">3.4</a>	Transport Efficiency .....	<a href="#">9</a>
<a href="#">4</a>	Faithfulness of Emulated Services .....	<a href="#">9</a>
<a href="#">4.1</a>	Characteristics of an Emulated Service .....	<a href="#">9</a>
<a href="#">4.2</a>	Service Quality of Emulated Services .....	<a href="#">10</a>
<a href="#">4.3</a>	Signaling of Native Services .....	<a href="#">10</a>
<a href="#">5</a>	Maintenance of Emulated Services .....	<a href="#">10</a>
<a href="#">5.1</a>	Signaling of Status Changes of an Emulated Service .....	<a href="#">10</a>
<a href="#">5.1.1</a>	Up/Down Notification .....	<a href="#">10</a>
<a href="#">5.1.2</a>	Packet Loss and/or Out-of-order Delivery .....	<a href="#">11</a>
<a href="#">5.1.3</a>	Other Status Signaling .....	<a href="#">11</a>
<a href="#">5.1.4</a>	Collective Status Signaling .....	<a href="#">11</a>
<a href="#">5.2</a>	Clock Recovery .....	<a href="#">12</a>
<a href="#">6</a>	Management of Emulated Services .....	<a href="#">12</a>
<a href="#">6.1</a>	MIB .....	<a href="#">12</a>
<a href="#">6.2</a>	Pseudo-Wire Traceroute .....	<a href="#">12</a>
<a href="#">7</a>	Scalability .....	<a href="#">12</a>
<a href="#">8</a>	Other Generic Requirements .....	<a href="#">13</a>
<a href="#">9</a>	Non-Requirements .....	<a href="#">14</a>

<a href="#">10</a>	Quality of Service (QoS) Considerations .....	<a href="#">14</a>
<a href="#">11</a>	Inter-domain Pseudo-Wires .....	<a href="#">15</a>
<a href="#">12</a>	Security Considerations .....	<a href="#">15</a>
12.1	Security Considerations for the Signaling/Control Channel .....	<a href="#">15</a>
<a href="#">12.2</a>	Security Considerations for the Encapsulated PDUs .....	<a href="#">15</a>
<a href="#">13</a>	Deployment Considerations .....	<a href="#">15</a>
<a href="#">14</a>	Acknowledgments .....	<a href="#">15</a>
<a href="#">15</a>	References .....	<a href="#">15</a>
<a href="#">16</a>	Authors' Addresses .....	<a href="#">16</a>
<a href="#">17</a>	Full Copyright Section .....	<a href="#">18</a>

Expires 11/2001      Xiao/McPherson/Pate/White/Kompella      [Page 2]

---

Internet Draft      [draft-ietf-pwe3-requirements-00](#)      May 17, 2001

## [1.](#) Introduction

### [1.1.](#) Reference Model of Pseudo-Wire Emulation Edge to Edge (PWE3)

To provide pseudo-wire emulation, two end-services of the same type are first configured between each service endpoint (SE) and the corresponding PWE3 Edge (PE) device (See Figure 1). An end-service is either:

- an Ethernet link between two ports or two VLANs, or
- a PPP link, or
- a (Cisco) HDLC link, or
- an ATM VC or VP, or
- a Frame Relay VC, or
- a TDM circuit.

A connection is then set up between the two PEs to connect these two end-services. This connection is called a pseudo-wire in the PWE3 context. During the setup of a pseudo-wire, the two pseudo-wire endpoints will exchange information about the service to be emulated so that later they know how to handle packets coming from the other end of the pseudo-wire. After the pseudo-wire is set up, layer-2 (e.g. Ethernet, ATM and Frame Relay) or layer-1 (e.g. SONET/SDH) PDUs of frames from an end-service are encapsulated at the pseudo-wire ingress. The encapsulated PDUs are then transported over the pseudo-wire to the egress, where L2 or L1 header information is re-constructed and the resulted frames are forwarded in their native format to the other end-service.

```

end-      |<----- pseudo-wire ----->|
          |                               | end-

```




Figure 1: Pseudo-Wire Reference Model

Different carriers may choose different types of pseudo-wires. For example, carriers may choose to use GRE tunnels [GRE], L2TP tunnels [L2TP], or MPLS LSPs [MPLS] as pseudo-wires. This document does not assume that a particular type of pseudo-wires is used. Instead, it describes generic requirements that apply to all pseudo-wires, for all services including Ethernet, PPP, (Cisco) HDLC, ATM, Frame Relay and TDM.

This document is not an introductory document. Readers are assumed to be familiar with the PWE3 framework document [PATE], especially the architecture assumptions stated in that document.

## [1.2.](#) 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.

End-Service	An End-Service is either an Ethernet link between two ports or two VLANs, or a PPP link, or a (Cisco) HDLC link, or an ATM VC or VP, or a Frame Relay VC, or a TDM circuit.
Pseudo-Wire Emulation	Pseudo-Wire Emulation is an approach to emulate the essential attributes of a service over a packet network so that two remote end-services appear directly connected by a wire.
Emulated Service	An Emulated Service is a service that is provided via pseudo-wire emulation.

Service Endpoint	A Service Endpoint (SE) is a device where one end of an emulated service terminates.
Pseudo-Wire	A Pseudo-Wire is a connection between two end-services.
Pseudo-Wire Endpoint	A Pseudo-Wire Endpoint is a device where one end of a pseudo-wire terminates.
Path-oriented Pseudo-Wire	A Path-oriented Pseudo-Wire is a pseudo-wire for which core network devices must maintain state information, for example, an MPLS LSP.
Non-path-oriented Pseudo-Wire	A Non-path-oriented Pseudo-Wire is a pseudo-wire for which core network devices need not maintain state information, for example, a L2TP tunnel.
Transport Tunnel	A Transport Tunnel is a tunnel inside which multiple pseudo-wires can be nested so that

Expires 11/2001      Xiao/McPherson/Pate/White/Kompella      [Page 5]

---

Internet Draft      [draft-ietf-pwe3-requirements-00](#)      May 17, 2001

they are transparent to core network devices.

## [2.](#) PDU Encapsulation

Every pseudo-wire edge device MUST provide service interfaces to use common service-specific techniques for encapsulating PDUs. It should be noted that the PDUs to be encapsulated may or may not contain L2 or L1 header information. This is service specific. Every PWE3 service MUST specify what a PDU is.

A pseudo-wire header must contain sufficient but not excessive information so that the other end of the pseudo-wire, aided with information obtained during the pseudo-wire setup process, knows exactly how to handle the encapsulated PDUs. A pseudo-wire header consists of all the header fields in an encapsulated PDU that are used by the pseudo-wire egress to determine how to handle the PDU. The header that is used for transporting the encapsulated PDUs from the pseudo-wire ingress to the egress (e.g. the tunnel label in [\[MART2\]](#)) is not considered as part of the pseudo-wire header. It is called transport header instead. It should be noted that transport headers are totally different from transport-layer headers (e.g. TCP/UDP headers).

Specific requirements on PDU encapsulation for a PWE3 approach are

listed below.

### [2.1.](#) Conveyance of Necessary L2/L1 Header Information

Sometimes the egress of a pseudo-wire needs some L2 or L1 header information in order to know how to process the PDUs received. A PWE3 approach SHOULD provide some mechanism (e.g. using one or more control words) for conveying such L2/L1 header information from the pseudo-wire ingress to the egress. The use of these mechanisms can be REQUIRED or OPTIONAL, depending on services.

### [2.2.](#) Location of the Pseudo-Wire Header

A pseudo-wire header MUST be between the PDU and the transport header. This is necessary for the pseudo-wire egress to locate the pseudo-wire header, and to process the PDU.

### [2.3.](#) PDU Length

A PWE3 approach MUST accommodate variable length PDUs, if variable length PDUs are allowed by the native service. For example, a PWE3 approach for Frame Relay MUST accommodate variable length frames.

### [2.4.](#) Encapsulation of Signaling Messages of the Native Services

PDUs that contain signaling information of the native service SHOULD be encapsulated in the same way as data PDUs. This simplifies handling at both the ingress and the egress of a pseudo-wire.

### [2.5.](#) Timing Information

Timing information can be essential for the endpoints of a service. If the original frames contain timing information, the encapsulation mechanism SHOULD not cause loss of such timing information. Requirements on clock recovery between pseudo-wire endpoints are further discussed in the "Maintenance of Emulated Services" section.

### [2.6.](#) Support for Circuit Multiplexing

If a service in its native form is capable of grouping multiple circuits into a "trunk", e.g. an ATM VP, some mechanism SHOULD be provided so that a single pseudo-wire can be used to connect two

end-trunks. From encapsulation perspective, the encapsulation header MUST carry sufficient information, e.g. VCI of each cell, so that the egress of the pseudo-wire can demultiplex individual circuits from the pseudo-wire.

### [2.7.](#) Packet Ordering

When packets carrying the encapsulated PDUs traverse a pseudo-wire, they may arrive at the egress out of order. For some services, the encapsulated PDUs must be delivered in order. For such services, some mechanism MUST be provided to ensure that. Providing a sequence number in the pseudo-wire header for each packet is one possible approach. Every specific PWE3 approach MUST define as a mechanism if in-order PDU delivery is required.

## [3.](#) Packet Transport

This section describes requirements on how to transport packets carrying the encapsulated PDUs over the network that provides PWE3 services.

### [3.1.](#) Setup of Pseudo-Wires

A pseudo-wire is a tunnel that connects two end-services. After the L2/L1 PDUs of a service are encapsulated, they must be transported over the pseudo-wire to the other end-service.

Every PWE3 approach MUST define some signaling mechanism for setting up the pseudo-wires. During the setup process, the pseudo-wire

Expires 11/2001      Xiao/McPherson/Pate/White/Kompella      [Page 7]

---

Internet Draft      [draft-ietf-pwe3-requirements-00](#)      May 17, 2001

endpoints need to exchange some information (e.g. learn each other's capability). The signaling mechanism MUST enable the pseudo-wire endpoints to exchange all necessary information. For example, both endpoints must agree on an encapsulation method and the MTU size for the pseudo-wire. As another example, in order to support circuit multiplexing using ATM VPs, both pseudo-wire endpoints must agree on using the VCIs in the encapsulated PDUs to demultiplex individual VCs from the VP at the pseudo-wire egress. Which signaling protocol to use and what information to exchange are service specific. Every PWE3 approach MUST specify these. Manual configuration can be considered as a special kind of signaling and is explicitly allowed.

This document does not assume a particular type of pseudo-wires. GRE tunnels, L2TP tunnels, or MPLS LSPs can all be used. Selection of a particular type of pseudo-wires is carrier-dependent and is outside



scope of the PWE3 WG.

A pseudo-wire can be either path-oriented or non-path-oriented. The difference is core network devices need to maintain state information for path-oriented pseudo-wires (e.g. LSPs) but not for non-path-oriented ones (e.g. GRE/L2TP tunnels). This has scalability implication that is further discussed in the "Scalability" section.

### [3.2.](#) Pseudo-Wire MTU

A pseudo-wire MUST be able to be configured with an MTU that is sufficient to transport packets whose size equals to the largest PDU size of the native service plus the pseudo-wire header and transport header. At a pseudo-wire ingress, if a packet's length exceeds the pseudo-wire MTU, it MUST be dropped. At a pseudo-wire egress, if the length of a frame that is restored from a PDU exceeds the MTU of destination end-service, it MUST be dropped.

### [3.3.](#) Transport of Signaling Messages of the Native Services

Some native services use signaling messages for maintaining the circuits. These signaling messages may be in-band, e.g. Ethernet flow control or ATM performance management, or out-of-band, e.g. the signaling VC of an ATM VP (from the perspective of other VCs in that VP). If such signaling messages are lost, functionality of the services will be significantly affected. Therefore, it can be desirable to provide higher reliability for transporting signaling messages.

Each PWE3 approach MAY state what approach can be used to ensure that signaling messages of the native service will be delivered over the network with high probability. Differentiating signaling packets from data packets and giving them preferable forwarding treatment in the network is one possible approach. Such approaches NEED not be

defined in a PWE3 approach itself.

### [3.4.](#) Transport Efficiency

In order to reduce transport header overhead and increase bandwidth efficiency, multiple PDUs MAY be concatenated before a transport header is added. Each PDU still carries its own pseudo-wire header so that the egress of the pseudo-wire knows how to handle it.

## [4.](#) Faithfulness of Emulated Services

An emulated service SHOULD be as similar to the native service as possible, but it is not REQUIRED that they should be identical. Each difference between an emulated service and the corresponding native service, if any, can be classified as either functional or non-functional. Functional differences are those that cause some features of the native service to become disabled in the emulated one. For example, if an emulated Ethernet service introduces some difference that can cause the Spanning Tree Protocol (STP) to malfunction, such a difference will be classified as a functional difference. Other differences are non-functional. For examples, differences in service quality between an emulated service and the native one are non-functional. In every PWE3 approach:

- All functional differences and the features disabled MUST be pointed out;
- Non-functional differences SHOULD be pointed out.

Some basic requirements on faithfulness of an emulated service are described below.

#### 4.1. Characteristics of an Emulated Service

Functionally, two service endpoints that are connected by an emulated service MUST appear directly connected by a native service, although service quality of the emulated service may be different from that of a native one. Specifically, this implies (but is not limited to) the followings:

- 1) It MUST be possible to define type (e.g. Ethernet or PPP, which is inherited from the native service), speed (e.g. 100Mbps), and MTU size for an emulated service.
- 2) From the service endpoints' perspective, each emulated service MUST appear as unshared.

- 3) If an emulated service fails (either at one of the end-services or in the middle of the pseudo-wire), both service endpoints MUST be notified reasonably timely. The definition of "reasonable timeliness" is service-dependent.
- 4) Two service endpoints connected by an emulated service MUST be able to establish routing protocol (e.g. IGP) adjacency over the

emulated service. To be clear, the Spanning Tree Protocol (STP) is not considered as a routing protocol and requirements on support/non-support of STP is outside scope of this document.

- 5) When desired, an emulated service MUST be able to appear as a normal link in the service endpoints' IGP routing table.
- 6) The TTL fields of the encapsulated PDUs, if exist, MUST not be changed inside an emulated service.

#### [4.2.](#) Service Quality of Emulated Services

It is RECOMMENDED but not REQUIRED that an emulated service provide as high service quality as the native service. However, the PWE3 WG only defines mechanisms for providing pseudo-wire emulation, not the services themselves. What quality to provide for a specific emulated service is a matter between a service provider (SP) and its customers, and is outside scope of the PWE3 WG. In fact, different SPs can use the same PWE3 approach but different QoS approaches to provide the same emulated service with different service quality.

#### [4.3.](#) Signaling of Native Services

A PWE3 approach SHOULD support signaling of the native service. This is further discussed in the "Maintenance of Emulated Services" section.

### [5.](#) Maintenance of Emulated Services

Every PWE3 approach MUST provide mechanisms for maintaining the working condition of an emulated service and for signaling any status changes.

#### [5.1.](#) Signaling of Status Changes of an Emulated Service

##### [5.1.1.](#) Up/Down Notification

With pseudo-wire emulation, a failure may occur at a remote place such that one or both physical links between the SEs and PEs remains up. For example in Figure 1, if the physical link between SE1 and PE1 fails, the physical link between SE2 and PE2 will not be affected

and will remain up. Unless some signaling is done to notify SE2 of the remote failure, SE2 will continue to send traffic over the emulated service to SE1. Such traffic will be discarded at PE1.

Therefore, when an emulated service fails (either at an end-service or in the middle of the pseudo-wire), both service endpoints MUST be notified. Every PWE3 approach MUST provide such a signaling mechanism. This functionality is equivalent to failure notification of normal links.

Similarly, every PWE3 approach MUST provide some signaling mechanism so that when a network failure is fixed, all the affected emulated services will change status from "Down" to "Up".

#### 5.1.2. Packet Loss and/or Out-of-order Delivery

A pseudo-wire can incur packet loss and/or out-of-order delivery. This can significantly impact the working condition of an emulated service. Number of packets lost or delivered out of order in unit time can be considered as two new types of "generalized bit error rate" of a pseudo-wire. Every PWE3 approach SHOULD provide some mechanism so that if either type of "generalized bit error rate" exceeds some configured threshold, the pseudo-wire and the emulated service will be signaled down.

#### 5.1.3. Other Status Signaling

A PWE3 approach MAY provide mechanism for other status signaling, if any.

#### 5.1.4. Collective Status Signaling

Status of a group of emulated services may be affected identically by a single network incidence. For example, when the physical link between a SE and a PE fails, all the emulated services that go through that link will fail. It is likely that there exists a group of emulated services which all terminate at a remote SE. (There can be multiple such SEs). Therefore, it is desirable that a single signaling message can be used to signal the failure of the whole group of emulated services. A PWE3 approach MAY provide some mechanism for signaling status changes of a group of emulated circuits. One possible approach is to associate each emulated service with a group ID when the pseudo-wire for that emulated service is set up. Multiple emulated services can then be grouped by associating them with identical group ID. In status signaling, that group ID can then be used to refer all the emulated services in that group.

## [5.2.](#) Clock Recovery

For some services, the pseudo-wire endpoints need to maintain some kind of timing relationship (e.g. synchronization). A PWE3 approach for such services MUST provide some mechanism to support that.

## [6.](#) Management of Emulated Services

Each PWE3 approach SHOULD provide some mechanisms for network operators to manage the emulated service.

### [6.1.](#) MIB

A pseudo-wire MIB (PW-MIB) MAY be provided for managing each emulated service. The MIB SHOULD have the following requirements:

- The counters in the MIB SHOULD NOT replicate existing MIB counters.
- Each end-service SHOULD appear as an interface in the ifTable.
- The MIB SHOULD describe how the existing counters are used for pseudo-wire emulation.
- The MIB MAY support row creation to create new end-service pairs.
- The MIB MAY augment existing tables. For example, the ifTable SHOULD be augmented to provide counts of out-of-order packets.
- The MIB SHOULD define meanings for the standard linkUp and linkDown traps, as well as any protocol-specific traps that are needed.
- The MIB SHOULD be structured in a hierarchical fashion such that generic PW objects and tables are not duplicated for each service.

### [6.2.](#) Pseudo-Wire Traceroute

It can be desirable to know the exact path of a pseudo-wire, especially for troubleshooting purpose. A pseudo-wire emulation service MAY support pseudo-wire traceroute, even if the pseudo-wire used is non-path-oriented.

## [7.](#) Scalability

Scalability requirements for Pseudo-Wire Emulation Edge to Edge are described below.

- Pseudo-wire emulation services SHOULD be transparent to other packet services (e.g. best effort Internet service).
- Only pseudo-wire endpoints should be aware of existence of pseudo-wires and PWE3 services. Core network devices SHOULD NOT be exposed to a large number of pseudo-wires.

Pseudo-wires can be path-oriented or non-path-oriented. For path-oriented pseudo-wires, core network devices must maintain state information for them. If a large number of path-oriented pseudo-wires are used, core network devices will have to maintain a large amount of state information. In order to avoid scalability problem, transport tunnels between pseudo-wire endpoints can be introduced. Pseudo-wires from the same ingress to the same egress are nested inside the transport tunnel that is from that ingress to that egress. By creating such a tunnel hierarchy, individual pseudo-wires are transparent to the core devices. If a specific PWE3 approach uses path-oriented pseudo-wires, transport tunnels SHOULD be used to improve scalability. However, a transport tunnel is not part of any pseudo-wire. Signaling of transport tunnels NEED NOT be defined in the PWE3 approach itself. As an example, if LSPs are used as pseudo-wires in a PWE3 approach, they can be nested inside a tunnel LSP from the pseudo-wire ingress to the pseudo-wire egress. The tunnel LSPs can be signaled by any mechanism defined in [[MPLS](#)].

- Circuit multiplexing: circuit multiplexing SHOULD be supported.
- Collective status signaling: Collective status signaling SHOULD be supported.

## [8.](#) Other Generic Requirements

Other generic requirements for Pseudo-Wire Emulation Edge to Edge include:

- No New Protocol Operations:

No matter which protocol (e.g. GRE or L2TP or MPLS) is used for packet transport, PWE3 SHOULD just be an application of that protocol. In other words, no new protocol (GRE or L2TP or MPLS) operations SHOULD be introduced. For example, if an MPLS label switching operation is performed, a PWE3 approach SHOULD not require that the TTL field of the top label remains unchanged. (Otherwise, this label switching would be different from a normal MPLS label switching and would be considered as a new protocol operation). If any new operations are indeed introduced in a PWE3 approach, they MUST be clearly pointed out.

Internet Draft      [draft-ietf-pwe3-requirements-00](#)

May 17, 2001

- Minimum configuration work at the pseudo-wire endpoints;
- Easy to maintain and manage.

## 9. Non-Requirements

Some non-requirements are mentioned in various sections of this document. Those work items are outside scope of the PWE3 WG. The non-requirements are listed below:

- Selection of a particular type of pseudo-wires;
- Striving to make the emulated services perfectly match their native services;
- Defining mechanisms for signaling the transport tunnels;
- Defining how to perform traffic management on packets that carry encapsulated PDUs;
- Providing security for the encapsulated PDUs;
- Providing any multicast service that is not native to the emulated medium.

To illustrate this point, Ethernet transmission to a multicast IEEE-48 address is considered in scope, while multicast services like [\[MARS\]](#) that are implemented on top of the medium are out of scope;

## 10. Quality of Service (QoS) Considerations

In this document, QoS means satisfactory service quality. It should not be confused with QoS mechanisms such as Weighted Fair Queueing (WFQ) or Random Early Detection (RED).

QoS is essential for ensuring that the emulated services are similar (but not necessarily identical) to their native forms. It is up to network operators to decide how to provide QoS - They can choose to rely on over-provisioning and/or deploy some QoS mechanisms.

In order to take advantage of some QoS mechanisms defined in other working groups, e.g. the traffic management schemes defined in

DiffServ WG, it is desirable that some mechanisms exist for differentiating the packets resulted from PDU encapsulation. These mechanisms need not be defined in the PWE3 approaches themselves. For

example, if the resulted packets are MPLS or IP packets, their EXP or DSCP fields can be used for marking and differentiating, respectively. Every PWE3 approach SHOULD provide guidelines for marking and differentiating, e.g. what fields in the original L2 or L1 headers should be used for determining the EXP/DSCP value. But the exact procedure on how to perform marking and differentiating, e.g. the mapping from Ethernet .1P field to EXP field, is out of scope.

## [11.](#) Inter-domain Pseudo-Wires

The PWE3 WG will determine the requirements for having a pseudo-wire pass across an administrative boundary. An edge could be a native hand-off or hand-off to a further pseudo-wire. The working group may analyze requirements for both security and performance for the inter-administration technology. This topic is for further study.

## [12.](#) Security Considerations

### [12.1.](#) Security Considerations for the Signaling/Control Channel

If a signaling/control channel is used in a PWE3 approach, some security mechanism SHOULD be provided to ensure integrity of the information transmitted across the signaling/control channel.

### [12.2.](#) Security Considerations for the Encapsulated PDUs

Providing security for the encapsulated PDUs is outside scope of the PWE3 WG.

## [13.](#) Deployment Considerations

Transition from using separate networks for TDM, ATM, Frame Relay and IP services to using a single network for multiple services requires careful planning and execution. This is an important topic for further study.

## [14.](#) Acknowledgments

Some requirements specified in this document were originally articulated in a number of documents including [[MART1](#)], [[MART2](#)], [[CES](#)], and [[SFBS](#)]. The authors would like to acknowledge authors of



those documents. The authors would also like to acknowledge the input and comments from T. So.

## 15. References

- [CES] A. Malis et al, "SONET/SDH Circuit Emulation Service Over MPLS (CEM) Encapsulation" <[draft-malis-sonet-ces-mpls-04.txt](#)>, work in progress, April 2001.

Expires 11/2001 Xiao/McPherson/Pate/White/Kompella

[Page 15]

---

Internet Draft [draft-ietf-pwe3-requirements-00](#) May 17, 2001

- [GRE] D. Farinacci, T. Li, S. Hanks, D. Meyer, P. Traina, "Generic Routing Encapsulation (GRE)", [RFC2784](#), March 2000.
- [L2TP] W.M. Townsley, A. Valencia, A. Rubens, G. Singh Pall, G. Zorn, B. Palter, "Layer Two Tunneling Protocol (L2TP)", [RFC 2661](#), August 1999.
- [MARS] G. Armitage, "Support for Multicast over UNI 3.0/3.1 based ATM Networks", [RFC2022](#), November 1996
- [MART1] L. Martini et al, "Transport of Layer 2 Frames Over MPLS", <[draft-martini-l2circuit-trans-MPLS-06.txt](#)>, work in progress, May 2001.
- [MART2] L. Martini et al, "Encapsulation Methods for Transport of Layer 2 Frames Over MPLS", <[draft-martini-l2circuit-encap-MPLS-02.txt](#)>, work in progress, May 2001.
- [MPLS] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture", [RFC3031](#), January 2001
- [PATE] P. Pate, X. Xiao, T. So, C. White, and K. Kompella, "Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3)", <[draft-pate-pwe3-framework-00.txt](#)>, work in progress, May 2001.
- [RTP] H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson "RTP: A Transport Protocol for Real-Time Applications", [RFC1889](#), January 1996.
- [SFBS] B. St-Denis, D. Fedyk, A. Bhatnagar, A. Siddiqui, R. Hartani, and T. So "Multi-service over MPLS", <[draft-stdenis-ms-over-mpls-00.txt](#)>, work in progress, Nov. 2000.

## 16. Authors' Addresses

XiPeng Xiao  
Photuris, Inc.  
2025 Stierlin Court  
Mountain View, CA 94043  
Email: xxiao@photuris.com

Danny McPherson  
Amber Networks, Inc.  
48664 Milmont Drive  
Fremont, CA 94538  
Email: danny@ambernetworks.com

Expires 11/2001      Xiao/McPherson/Pate/White/Kompella

[Page 16]

---

Internet Draft      [draft-ietf-pwe3-requirements-00](#)

May 17, 2001

Prayson Pate  
Overture Networks  
P. O. Box 14864  
RTP, NC, USA 27709  
Email: prayson.pate@overturenetworks.com

Craig White  
Level 3 Communications, LLC.  
1025 Eldorado Blvd.  
Broomfield, CO, 80021  
e-mail: Craig.White@Level3.com

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

[17.](#) Full Copyright Section

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF

MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.