

Internet Draft  
Document: [draft-ietf-pwe3-requirements-05.txt](#)  
Expires: September 2003

XiPeng Xiao  
Riverstone Networks

Danny McPherson  
TCB

Prayson Pate  
Overture Networks

Editors

Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)  
[draft-ietf-pwe3-requirements-05.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/1id-abstracts.txt>

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

Abstract

This document describes base requirements for the Pseudo-Wire Emulation Edge to Edge Working Group (PWE3 WG). It provides guidelines for other working group documents that will define mechanisms for providing pseudo-wire emulation of Ethernet, ATM, and Frame Relay. Requirements for pseudo-wire emulation of TDM (i.e. "synchronous bit streams at rates defined by ITU G.702") are defined in another document. It should be noted that the PWE3 WG standardizes mechanisms that can be used to provide PWE3 services, but not the services themselves.

## Co-Authors

The following are co-authors of this document:

Vijay Gill	AOL Time Warner, Inc.
Kireeti Kompella	Juniper Networks, Inc.
Thomas D. Nadeau	Cisco Systems
Craig White	Level 3 Communications, LLC.

## Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Internet Draft

[draft-ietf-pwe3-requirements-05](#)

Mar. 2003

## Table of Contents

<a href="#">1</a>	<a href="#">Terminology</a>	<a href="#">4</a>
<a href="#">2</a>	<a href="#">Introduction</a>	<a href="#">4</a>
<a href="#">2.1</a>	<a href="#">What Are Pseudo Wires?</a>	<a href="#">4</a>
<a href="#">2.2</a>	<a href="#">Background and Motivation</a>	<a href="#">5</a>
<a href="#">2.3</a>	<a href="#">Current Network Architecture</a>	<a href="#">5</a>
<a href="#">2.4</a>	<a href="#">PWE3 as a Path to Convergence</a>	<a href="#">6</a>
<a href="#">2.5</a>	<a href="#">Suitable Applications for PWE3</a>	<a href="#">6</a>
<a href="#">2.6</a>	<a href="#">Summary</a>	<a href="#">7</a>
<a href="#">3</a>	<a href="#">Reference Model of PWE3</a>	<a href="#">7</a>
<a href="#">4</a>	<a href="#">Packet Processing</a>	<a href="#">8</a>
<a href="#">4.1</a>	<a href="#">Encapsulation</a>	<a href="#">8</a>
<a href="#">4.2</a>	<a href="#">Frame Ordering</a>	<a href="#">9</a>
<a href="#">4.3</a>	<a href="#">Frame Duplication</a>	<a href="#">9</a>
<a href="#">4.4</a>	<a href="#">Fragmentation</a>	<a href="#">9</a>
<a href="#">4.5</a>	<a href="#">Consideration of Per-PSN Packet Overhead</a>	<a href="#">9</a>
<a href="#">5</a>	<a href="#">Maintenance of Emulated Services</a>	<a href="#">10</a>
<a href="#">5.1</a>	<a href="#">Setup and Teardown of Pseudo-Wires</a>	<a href="#">10</a>
<a href="#">5.2</a>	<a href="#">Handling Maintenance Message of the Native Services</a>	<a href="#">10</a>
<a href="#">5.3</a>	<a href="#">PE-generated Maintenance Messages</a>	<a href="#">11</a>
<a href="#">6</a>	<a href="#">Management of Emulated Services</a>	<a href="#">13</a>
<a href="#">6.1</a>	<a href="#">MIBs</a>	<a href="#">13</a>
<a href="#">6.2</a>	<a href="#">General MIB Requirements</a>	<a href="#">13</a>
<a href="#">6.3</a>	<a href="#">Configuration and Provisioning</a>	<a href="#">13</a>
<a href="#">6.4</a>	<a href="#">Performance Monitoring</a>	<a href="#">13</a>
<a href="#">6.5</a>	<a href="#">Fault Management and Notifications</a>	<a href="#">14</a>
<a href="#">6.6</a>	<a href="#">Pseudo-Wire Connection Verification and Traceroute</a>	<a href="#">14</a>
<a href="#">7</a>	<a href="#">Faithfulness of Emulated Services</a>	<a href="#">14</a>
<a href="#">7.1</a>	<a href="#">Characteristics of an Emulated Service</a>	<a href="#">14</a>
<a href="#">7.2</a>	<a href="#">Service Quality of Emulated Services</a>	<a href="#">15</a>
<a href="#">8</a>	<a href="#">Non-Requirements</a>	<a href="#">15</a>
<a href="#">9</a>	<a href="#">Quality of Service (QoS) Considerations</a>	<a href="#">16</a>
<a href="#">10</a>	<a href="#">Inter-domain Issues</a>	<a href="#">16</a>
<a href="#">11</a>	<a href="#">Security Considerations</a>	<a href="#">17</a>
<a href="#">12</a>	<a href="#">Acknowledgments</a>	<a href="#">17</a>
<a href="#">13</a>	<a href="#">References</a>	<a href="#">17</a>
<a href="#">14</a>	<a href="#">Authors' Addresses</a>	<a href="#">18</a>
<a href="#">15</a>	<a href="#">Full Copyright Section</a>	<a href="#">20</a>

Internet Draft

[draft-ietf-pwe3-requirements-05](#)

Mar. 2003

## 1. 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](#).

Some terms used throughout this document are listed below.

### Attachment Circuit (AC)

The circuit or virtual circuit attaching a CE to a PE.

### Customer Edge

A device where one end of a service originates and/or terminates. The CE is not aware that it is using an emulated service rather than a native service.

### Packet Switched Network

A network using IP or MPLS as the mechanism for packet forwarding

### Provider Edge

A device that provides PWE3 to a CE.

### Pseudo Wire

A mechanism that carries the essential elements of an emulated circuit from one PE to one or more other PEs over a PSN.

### Pseudo Wire Emulation Edge to Edge

A mechanism that emulates the essential attributes of a service (such as a T1 leased line or Frame Relay) over a PSN.

### Pseudo Wire PDU

A PDU sent on the PW that contains all of the data and control information necessary to emulate the desired service.

### PSN Tunnel

A tunnel across a PSN inside which one or more PWs can be carried.

## 2. Introduction

### 2.1. What Are Pseudo Wires?

Pseudo Wire Emulation Edge-to-Edge (PWE3) is a mechanism that emulates the essential attributes of a service over a Packet Switched Network (PSN). The required functions of PWs include encapsulating service-specific PDUs arriving at an ingress port, and carrying them

across a path or tunnel, managing their timing and order, and any other operations required to emulate the behavior and characteristics of the service as faithfully as possible.

From the customer perspective, the PW is perceived as an unshared link or circuit of the chosen service. However, there may be deficiencies that impede some applications from being carried on a PW. These limitations should be fully described in the appropriate service-specific documents and Applicability Statements.

## [2.2.](#) Background and Motivation

The following sections give some background on where networks are today and why they are changing. It also talks about the motivation to provide converged networks while continuing to support existing services. Finally it discusses how PWs can be a solution for this dilemma.

## [2.3.](#) Current Network Architecture

### [2.3.1.](#) Multiple Networks

For any given service provider delivering multiple services, the current infrastructure usually consists of parallel or "overlay" networks. Each of these networks implements a specific service, such as Frame Relay, Internet access, etc. This is expensive, both in terms of capital expense and operational costs. Furthermore, the presence of multiple networks complicates planning. Service providers wind up asking themselves these questions:

- Which of my networks do I build out?
- How many fibers do I need for each network?
- How do I efficiently manage multiple networks?

A converged network helps service providers answer these questions in a consistent and economical fashion.

### [2.3.2.](#) Transition to a Packet-Optimized Converged Network

In order to maximize return on their assets and minimize their operating costs, service providers often look to consolidate the delivery of multiple service types onto a single networking technology.

As packet traffic takes up a larger and larger portion of the available network bandwidth, it becomes increasingly useful to optimize public networks for the Internet Protocol. However, many service providers are confronting several obstacles in engineering packet-optimized networks. Although Internet traffic is the fastest

growing traffic segment, it does not generate the highest revenue per bit. For example, Frame Relay traffic currently generates higher revenue per bit than native IP services do. Private line TDM services still generate even more revenue per bit than does Frame Relay. In addition, there is a tremendous amount of legacy equipment deployed within public networks that does not communicate using the Internet Protocol. Service providers continue to utilize non-IP equipment to deploy a variety of services, and see a need to interconnect this legacy equipment over their IP-optimized core networks.

#### 2.4. PWE3 as a Path to Convergence

How do service providers realize the capital and operational benefits of a new packet-based infrastructure, while leveraging the existing equipment and also protecting the large revenue stream associated with this equipment? How do they move from mature Frame Relay or ATM networks, while still being able to provide these lucrative services?

One possibility is the emulation of circuits or services via PWS. Circuit emulation over ATM and interworking of Frame Relay and ATM have already been standardized. Emulation allows existing services to be carried across the new infrastructure, and thus enables the interworking of disparate networks.

Implemented correctly, PWE3 can provide a means for supporting today's services over a new network.

#### 2.5. Suitable Applications for PWE3

What makes an application suitable (or not) for PWE3 emulation? When considering PWS as a means of providing an application, the following questions must be considered:

- Is the application sufficiently deployed to warrant emulation?
- Is there interest on the part of service providers in providing an emulation for the given application?
- Is there interest on the part of equipment manufacturers in providing products for the emulation of a given application?
- Are the complexities and limitations of providing an emulation worth the savings in capital and operational expenses?

If the answer to all four questions is "yes", then the application is likely to be a good candidate for PWE3. Otherwise, there may not be sufficient overlap between the customers, service providers, equipment manufacturers and technology to warrant providing such an emulation.

## 2.6. Summary

To maximize the return on their assets and minimize their operational costs, many service providers are looking to consolidate the delivery of multiple service offerings and traffic types onto a single IP-optimized network.

In order to create this next-generation converged network, standard methods must be developed to emulate existing telecommunications formats such as Ethernet, Frame Relay, and ATM over IP-optimized core networks. This document describes requirements for accomplishing this goal.

## 3. Reference Model of PWE3

A pseudo-wire (PW) is a connection between two provider edge (PE) devices which connects two attachment circuits (ACs). In this document, An AC is either:

- an Ethernet link or a 802.1Q link between two ports, or
- an ATM VCC or VPC, or
- a Frame Relay VC

between a customer edge (CE) device and a PE (See Figure 1).

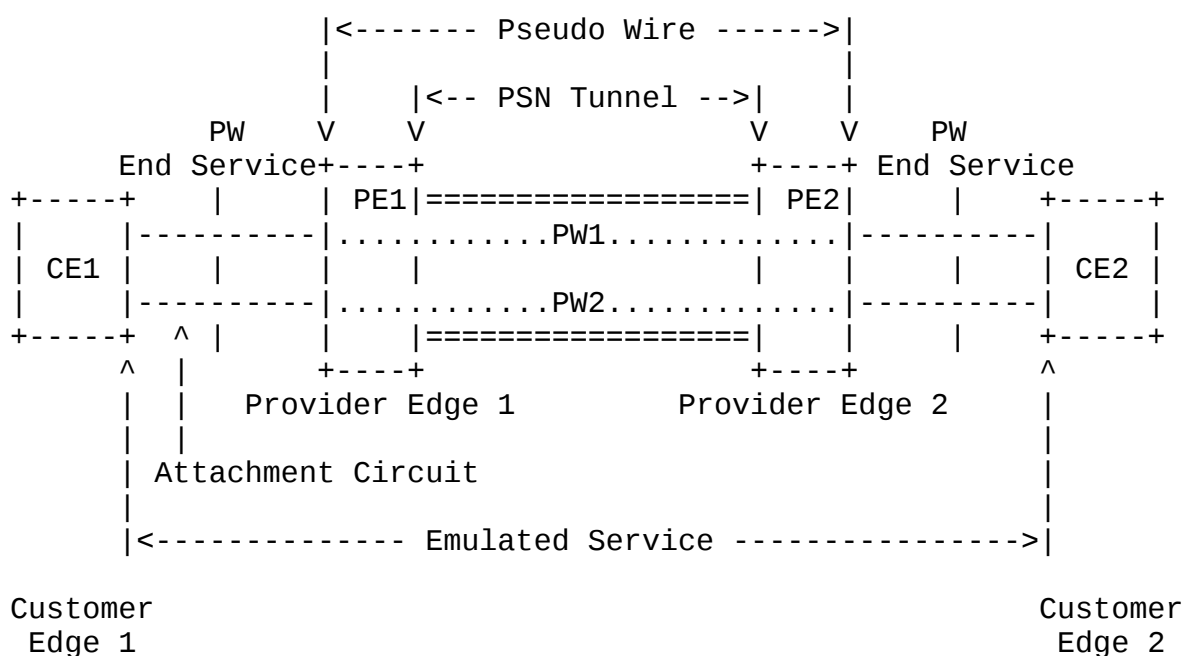


Figure 1: PWE3 Reference Model

During the setup of a PW, the two PEs will be configured or will automatically exchange information about the service to be emulated

so that later they know how to process packets coming from the other

end. After a PW is set up between two PEs, frames received by one PE from an AC are encapsulated and sent over the PW to the remote PE, where native frames are re-constructed and forwarded to the other CE. For a detailed PWE3 architecture overview, readers should refer to the PWE3 architecture document [[PWE3\\_ARCH](#)].

This document does not assume that a particular type of PWs (e.g. [[L2TPv3](#)] sessions or [[MPLS](#)] LSPs) is used. Instead, it describes generic requirements that apply to all PWs, for all services including Ethernet, ATM, and Frame Relay.

## [4.](#) Packet Processing

This section describes data plane requirements for PWE3.

### [4.1.](#) Encapsulation

Every PE MUST provide encapsulation mechanism for PDUs from a PWES. It should be noted that the PDUs to be encapsulated may or may not contain L2 header information. This is service specific. Every PWE3 service MUST specify what the PDU is.

A PW header consists of all the header fields in a PW PDU that are used by the PW egress to determine how to process the PDU. The PSN tunnel header is not considered as part of the PW header.

Specific requirements on PDU encapsulation are listed below.

#### [4.1.1.](#) Conveyance of Necessary L2 Header Information

The egress of a PW needs some information, e.g. which native service the PW PDUs belong to, and possibly some L2 header information, in order to know how to process the PDUs received. A PWE3 encapsulation approach MUST provide some mechanism for conveying such information from the PW ingress to the egress. It should be noted that not all such information must be carried in the PW header of the PW PDUs. Some information (e.g. service type of a PW) can be stored as state information at the egress during PW setup.

#### [4.1.2.](#) Support of Variable Length PDUs

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.

#### [4.1.3.](#) Support of Multiplexing and Demultiplexing

Ethernet 802.1Q interfaces in a port, some mechanism SHOULD be provided so that a single PW can be used to connect two end-trunks. From encapsulation perspective, sufficient information MUST be carried so that the egress of the PW can demultiplex individual circuits from the PW.

#### 4.2. Frame Ordering

When packets carrying the PW PDUs traverse a PW, they may arrive at the egress out of order. For some services, the frames (either control frames only or both control and data frames) must be delivered in order. For such services, some mechanism MUST be provided for ensuring in-order delivery. Providing a sequence number in the PW header for each packet is one possible approach to detect out-of-order frames. Mechanisms for re-ordering frames may be provided by Native Service Processing (NSP) [[PWE3\\_ARCH](#)] but are out of scope of PWE3.

#### 4.3. Frame Duplication

In rare cases, packets traversing a PW may be duplicated. For some services, frame duplication is not allowed. For such services some mechanism MUST be provided to ensure that duplicated frames will not be delivered. The mechanism may or may not be the same as the mechanism used to ensure in-order frame delivery.

#### 4.4. Fragmentation

If the combined size of the L2 payload and its associated PWE3 and PSN headers exceeds the PSN path MTU, the L2 payload may need to be fragmented (Alternatively the L2 frame may be dropped). For certain native service, fragmentation may also be needed to maintain a control frame's relative position to the data frames (e.g. an ATM PM cell's relative position). In general, fragmentation has a performance impact. It is therefore desirable to avoid fragmentation if possible. However, for different services, the need for fragmentation can be different. When there is potential need for fragmentation, each service-specific PWE3 document MUST specify whether to fragment the frame in question or to drop it. If an emulated service chooses to drop the frame, the consequence MUST be specified in its applicability statement.

#### 4.5. Consideration of Per-PSN Packet Overhead

When the L2 PDU size is small, in order to reduce PSN tunnel header overhead, multiple PDUs MAY be concatenated before a PSN tunnel header is added. Each encapsulated PDU still carries its own PW header so that the egress PE knows how to process it. However, the

benefit of concatenating multiple PDUs for header efficiency should be weighed against the resulting increase in delay, jitter and the

larger penalty incurred by packet loss.

## 5. Maintenance of Emulated Services

This section describes maintenance requirements for PWE3.

### 5.1. Setup and Teardown of Pseudo-Wires

A PW must be set up before an emulated circuit can be established, and must be torn down when an emulated circuit is no longer needed. Setup and teardown of a PW can be triggered by a CLI command from the management plane of a PE, or by Setup/Teardown of an AC (e.g., an ATM SVC), or by an auto-discovery mechanism.

Every PWE3 approach MUST define some setup mechanism for establishing the PWs. During the setup process, the PEs need to exchange some information (e.g. to learn each other's capability). The setup mechanism MUST enable the PEs to exchange all necessary information. For example, both endpoints must agree on methods for encapsulating PDUs and handling frame ordering. Which signaling protocol to use and what information to exchange are service specific. Every PWE3 approach MUST specify them. Manual configuration of PWs can be considered as a special kind of signaling and is allowed.

If a native circuit is bi-directional, the corresponding emulated circuit can be signaled "Up" only when the associated PW and PSN tunnels in both directions are functional.

### 5.2. Handling Maintenance Message of the Native Services

Some native services have mechanisms for maintenance purpose, e.g. ATM OAM and FR LMI. Such maintenance messages can be in-band (i.e. mixed with data messages in the same AC) or out-of-band (i.e. sent in a dedicated control circuit). For such services, all in-band maintenance messages related to a circuit SHOULD be transported in-band just like data messages through the corresponding PW to the remote CE. In other words, no translation is needed at the PEs for in-band maintenance messages. In addition, it MAY be desirable to provide higher reliability for maintenance messages. The mechanisms for providing high reliability NEED NOT be defined in the PWE3 WG.

Out-of-band maintenance messages between a CE and a PE may relate to multiple ACs between the CE and the PE. They need to be processed at the local PE and possibly at the remote PE as well. If a native service has some out-of-band maintenance messages, the corresponding emulated service MUST specify how to process such messages at the PEs.

new maintenance mechanisms.

### 5.3. PE-generated Maintenance Messages

A PE needs to generate some maintenance messages under two circumstances for an emulated service. Such circumstances are triggered either:

1. by maintenance messages of the native service; or
2. by some events such as an "Up/Down" status change of the PW or the local AC.

In circumstance #1, when a PE receives an out-of-band maintenance message from the local CE, for each relevant emulated circuit, the PE may need to translate the out-of-band maintenance message into an appropriate in-band maintenance message of the native service and send it via the PW to the remote CE. For example, if the ACs between a CE and a PE are some ATM VCCs, and a F4 AIS is received by the PE from the CE, the PE may need to translate that F4 AIS into a F5 AIS for each VCC and send it to the remote CE. Alternatively, for each relevant emulated circuit, the PE may generate a PWE-specific maintenance message to the remote PE, e.g. a label withdrawal message. Sometimes, multiple such maintenance messages can be grouped together; this is further discussed in the "collective status notification" paragraph later. When the remote PE receives such a PWE-specific maintenance message, it may need to generate a maintenance message of the native service and send it to the attached CE.

In circumstance #2, the reason the PEs need to generate some maintenance messages under some events is because the existence of a PW between two CEs reduces the CEs' maintenance capability if the PEs do nothing. This is illustrated in the following example. If two CEs are directly connected by a physical wire, a native service (e.g. ATM) can use notifications of the lower layer (e.g. the physical link layer) to assist its maintenance. For example, the ATM PVC can be signaled "Down" if the physical wire fails. However, consider the following scenario.

```
+-----+ Phy-link +-----+           +-----+ Phy-link +-----+
| CE1 |-----| PE1|.....PW.....| PE2 |-----| CE2 |
+-----+           +-----+           +-----+           +-----+
```

If the PW between PE1 and PE2 fails, CE1 and CE2 will not receive physical link-failure notification. As a result, they cannot declare failure of the emulated circuit in a timely fashion, which will in turn affect higher layer applications. Therefore, when the PW fails, PE1 and PE2 need to generate some maintenance messages to notify the client layer on CE1 and CE2 that use the PW as a server layer. (In this case, the client layer is the emulated service). Similarly, if

the physical link between PE1-CE1 fails, PE1 needs to generate some

maintenance message(s) so that the client layer at CE2 will be notified. PE2 may be involved in this process.

In the rare case when a physical wire between two CEs incurs many bit errors, the physical link can be declared "Down" and the client layer at the CEs be notified. Similarly, a PW can incur packet loss, corruption, and out-of-order delivery. These can be considered as "generalized bit error". Upon detection of excessive "generalized bit error", a PE may need to take some maintenance actions so that the client layer at the CEs is notified.

To summarize the requirements for circumstance #2: PWE MUST provide to an emulated circuit the server-layer (i.e. lower-layer) maintenance assistance that a native circuit would receive from a physical wire.

Overall, the need for PE-generated maintenance messages is different for different services. Every emulated service MUST specify:

- \* what PE-generated maintenance messages are needed,
- \* when they are needed, and
- \* how to generate and process them at the PEs.

Furthermore, if a PE needs to generate and send a maintenance message to a CE, the PE MUST use a maintenance message of the native service. This is essential in keeping the emulated service transparent to the CEs.

In specifying "when PE-generated maintenance messages are needed", some monitoring mechanisms are needed for detecting the triggering events. (Some of such events are briefly discussed above). Such mechanisms NEED NOT be defined in the PWE3 WG. A service-specific-PWE-definition document MAY cite some status monitoring mechanisms defined elsewhere, e.g. [[LSPPING](#)].

Status of a group of emulated circuits may be affected identically by a single incidence. For example, when the physical link between a CE and a PE fails, all the emulated circuits that go through that link will fail. It is likely that there exists a group of emulated circuits which all terminate at a remote CE. (There can be multiple such CEs). Therefore, it is desirable that a single maintenance message be used to notify failure of the whole group of emulated circuits. A PWE3 approach MAY provide some mechanism for notifying status changes of a group of emulated circuits. One possible approach is to associate each emulated circuit with a group ID while setting up the PW for that emulated circuit. Multiple emulated circuits can then be grouped by associating them with an identical group ID. In the maintenance message, that group ID can be used to refer to all the emulated circuits in that group. This is called "collective status notification".

The requirements stated in this section comply with the ITU-T

maintenance philosophy (client layer/server layer concept) for telecommunications networks [[G805](#)].

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

Each PWE3 approach SHOULD provide some mechanisms for network operators to manage the emulated service. These mechanisms can be in the forms described below.

### [6.1.](#) MIBs

SNMP MIBs [[SMIV2](#)] MUST be provided for managing each emulated circuit as well as pseudo-wire in general. These MIBs SHOULD be created with the following requirements.

### [6.2.](#) General MIB Requirements

New MIBs MUST augment or extend where appropriate, existing tables as defined in other existing service-specific MIBs for existing services such as MPLS or L2TP. For example, the ifTable as defined in the Interface MIB [[IFMIB](#)] MUST be augmented to provide counts of out-of-order packets. A second example is the extension of the MPLS-TE-MIB [[TEMIB](#)] when emulating circuit services over MPLS. Rather than redefining the tunnelTable so that PWES can utilize MPLS tunnels, for example, entries in this table MUST instead be extended to add additional PWES-specific objects. Doing so facilitates a natural extension of those objects defined in the existing MIBs in terms of management, as well as leveraging existing agent implementations.

Interfaces implementing a PWES MUST appear as an interface in the ifTable.

### [6.3.](#) Configuration and Provisioning

MIB Tables MUST be designed to facilitate configuration and provisioning of the PWES.

The MIB(s) MUST facilitate intra-PSN configuration and monitoring of PWES connections.

### [6.4.](#) Performance Monitoring

MIBs MUST collect statistics for performance and fault management.

MIBs MUST provide a description of how existing counters are used for PW emulation and SHOULD not replicate existing MIB counters.

## 6.5. Fault Management and Notifications

Notifications SHOULD be defined where appropriate to notify the network operators of any interesting situations, including faults detected in the PWES.

Objects defined to augment existing protocol-specific notifications in order to add PWES functionality MUST explain how these notifications are to be emitted.

## 6.6. Pseudo-Wire Connection Verification and Traceroute

For network management purpose, a connection verification mechanism SHOULD be supported by PWs. Connection verification as well as other alarming mechanisms can alert network operators that a PW has lost its remote connection. It is sometimes desirable to know the exact functional path of a PW for troubleshooting purpose, thus a traceroute function capable of reporting the path taken by data packets over the PW SHOULD be provided.

## 7. Faithfulness of Emulated Services

An emulated service SHOULD be as similar to the native service as possible, but NOT REQUIRED to be identical. The applicability statement of a PWE3 service MUST report limitations of the emulated service.

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

### 7.1. Characteristics of an Emulated Service

From the perspective of a CE, an emulated circuit is characterized as an unshared link or circuit of the chosen service, although service quality of the emulated service may be different from that of a native one. Specifically, the following requirements MUST be met:

- 1) It MUST be possible to define type (e.g. Ethernet, which is inherited from the native service), speed (e.g. 100Mbps), and MTU size for an emulated circuit, if it is possible to do so for a native circuit.
- 2) If the two endpoints CE1 and CE2 of emulated circuit #1 are connected to PE1 and PE2, respectively, and CE3 and CE4 of emulated circuit #2 are also connected to PE1 and PE2, then the PWs of these two emulated circuits may share the same physical paths between PE1 and PE2. But from each CE's perspective, its emulated circuit MUST appear as unshared. For example, CE1/CE2 MUST NOT be aware of existence of emulated circuit #2 or CE3/CE4.

- 3) If an emulated circuit fails (either at one of the ACs or in the

middle of the PW), both CEs MUST be notified in a timely manner, if they will be notified in the native service. The definition of "timeliness" is service-dependent.

- 4) If a routing protocol (e.g. IGP) adjacency can be established over a native circuit, it MUST be possible to be established over an emulated circuit as well.

## 7.2. Service Quality of Emulated Services

It is NOT REQUIRED that an emulated service provide the same service quality as the native service. The PWE3 WG only defines mechanisms for providing PW 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.

## 8. Non-Requirements

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

- Service interworking;

In Service Interworking, the IWF (Interworking Function) between two dissimilar protocols (e.g., ATM & MPLS, Frame Relay & ATM, ATM & IP, ATM & L2TP, etc.) terminates the protocol used in one network and translates (i.e. maps) its Protocol Control Information (PCI) to the PCI of the protocol used in other network for User, Control and Management Plane functions to the extent possible.

- Selection of a particular type of PWs;
- To make the emulated services perfectly match their native services;
- Defining mechanisms for signaling the PSN tunnels;
- Defining how to perform traffic management on packets that carry PW PDUs;
- Providing security for the PW 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;

## 9. 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 Queuing (WFQ) or Random Early Detection (RED).

QoS is essential for ensuring that 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 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 packets are MPLS or IP packets, their EXP or DSCP fields can be used for marking and differentiating. A PWE3 approach MAY provide guidelines for marking and differentiating. But the exact procedure of how to perform marking and differentiating, e.g. specifying the mapping function from Ethernet 802.1p field to EXP field, is out of scope.

## 10. Inter-domain Issues

PWE is a matter between the PW end-points and is transparent to the network devices between the PW end-points. Therefore, inter-domain PWE is fundamentally similar to intra-domain PWE. As long as PW end-points use the same PWE approach, they can communicate effectively, regardless of whether they are in the same domain. Security may become more important in the inter-domain case and some security measure such as end-point authentication MAY be applied.

Inter-domain PSN tunnels are generally more difficult to set up, tear down and maintain than intra-domain ones. For example, inter-domain PSN tunnels cannot be set up using RSVP-TE. But that is an issue for PSN tunneling protocols such as MPLS and L2TPv3 and is outside the scope of PWE3.

## 11. Security Considerations

The PW end-point, PW demultiplexing mechanism, and the payloads of the native service can all be vulnerable to attack. PWE3 should leverage security mechanisms provided by the PW Demultiplexer or PSN Layers. Such mechanisms SHOULD protect PW end-point and PW

Demultiplexer mechanism from denial-of-service (DoS) attacks and spoofing of the native data units. Controlling PSN access to the PW end-point is generally effective against DoS attacks and spoofing, and can be part of protection mechanism. Protection mechanisms SHOULD also address the spoofing of tunneled PW data. The validation of traffic addressed to the PW Demultiplexer end-point is paramount in ensuring integrity of PW encapsulation. Security protocols such as IPSec [[RFC2401](#)] can be used.

## [12.](#) Acknowledgments

The authors would like to acknowledge input from M. Bocci, S. Bryant, R. Cohen, N. Harrison, G. Heron, T. Johnson, A. Malis, L. Martini, E. Rosen, J. Rutenmiller, T. So, Y. Stein and S. Vainshtein.

## [13.](#) References

- [G805] "Generic Functional Architecture of Transport Networks", ITU-T Recommendation G.805, 2000.
- [IFMIB] K. McCloghrie, and F. Kastenholtz, "The Interfaces Group MIB using SMIV2", [RFC 2233](#), Nov. 1997.
- [L2TPv3] J. Lau, M. Townsley, and I. Goyret, et. al., "Layer Two Tunneling Protocol (Version 3)", <[draft-ietf-l2tpext-l2tp-base-07.txt](#)>, work in progress, Feb. 2003.
- [LSPPING] K. Kompella, P. Pan, N. Sheth, D. Cooper, G. Swallow, S. Wadhwa, and R. Bonica, "Detecting Data Plane Liveliness in MPLS", <[draft-ietf-mpls-lsp-ping-02.txt](#)>, work in progress, Mar. 2003.
- [MARS] G. Armitage, "Support for Multicast over UNI 3.0/3.1 based ATM Networks", [RFC2022](#), November 1996
- [MPLS] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture", [RFC3031](#), January 2001
- [TEMIB] C. Srinivasan, A. Viswanathan, and T. Nadeau, "MPLS Traffic Engineering Management Information Base", <[draft-ietf-mpls-te-mib-09.txt](#)>, work in progress, Nov. 2002.

Expires Sept. 03

Xiao/McPherson/Pate

[Page 17]

---

Internet Draft      [draft-ietf-pwe3-requirements-05](#)      Mar. 2003

- [PWE3\_ARCH] S. Bryant and P. Pate, et. al., "PWE3 Architecture", <[draft-ietf-pwe3-arch-02.txt](#)>, work in progress, Feb. 2003.
- [RFC2401] S. Kent, R. Atkinson, "Security Architecture for the Internet Protocol", [RFC 2401](#), Nov. 1998.
- [RTP] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", [RFC1889](#), January 1996.

[SMIV2] J. Case, K. McCloghrie, M. Rose, and S. Waldbusser,  
"Structure of Management Information for Version 2 of the  
Simple Network Management Protocol (SNMPv2)", [RFC 1902](#),  
January 1996.

#### [14.](#) Authors' Addresses

XiPeng Xiao  
Riverstone Networks  
5200 Great America Parkway  
Santa Clara, CA 95054  
Email: [xxiao@riverstonenet.com](mailto:xxiao@riverstonenet.com)

Danny McPherson  
TCB.net  
Email: [danny@tcb.net](mailto:danny@tcb.net)

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

Vijay Gill  
AOL Time Warner  
EMail: [vijaygill19@aol.com](mailto:vijaygill19@aol.com)

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

Thomas D. Nadeau  
Cisco Systems, Inc.  
250 Apollo Drive  
Chelmsford, MA 01824  
Email: [tnadeau@cisco.com](mailto:tnadeau@cisco.com)

Expires Sept. 03

Xiao/McPherson/Pate

[Page 18]

---

Internet Draft

[draft-ietf-pwe3-requirements-05](#)

Mar. 2003

Craig White  
Level 3 Communications, LLC.  
1025 Eldorado Blvd.  
Broomfield, CO, 80021  
Email: [Craig.White@Level3.com](mailto:Craig.White@Level3.com)

## [15.](#) 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.