

**Application of PWE3 to MPLS Transport Networks
draft-bryant-pwe3-mpls-transport-00**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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>

This Internet-Draft will expire on April 13, 2007.

Abstract

A need has been identified to identify a strict subset of the PWE3 over MPLS pseudowire functionality suitable for the construction of transport networks. This draft describes the IETF recommended profile for such cases. This document describes the IETF-recommended profile for such a configuration. In particular the profile of an [RFC4448](#) PWE3 Ethernet pseudowire that can be used to address these requirements is described.

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](#) [[RFC2119](#)].

Table of Contents

- [1. Introduction.....](#)[2](#)
- [2. PWE3 Configuration.....](#)[3](#)
- [3. OAM.....](#)[4](#)
 - [3.1. VCCV profile 1: BFD without IP/UDP Headers.....](#)[4](#)
 - [3.2. VCCV profile 2: BFD with IP/UDP Headers.....](#)[4](#)
- [4. MPLS Layer.....](#)[4](#)
 - [4.1. External Configuration.....](#)[5](#)
 - [4.2. Control Plane Configuration.....](#)[5](#)
- [5. Congestion Considerations.....](#)[6](#)
- [6. Security Considerations.....](#)[6](#)
- [7. IANA Considerations.....](#)[6](#)
- [8. Acknowledgments.....](#)[6](#)
- [APPENDIX A: Requirements.....](#)[7](#)
- [9. References.....](#)[9](#)
 - [9.1. Normative References.....](#)[9](#)
 - [9.2. Informative References.....](#)[10](#)
- [Author's Addresses.....](#)[10](#)
- [Intellectual Property Statement.....](#)[12](#)
- [Disclaimer of Validity.....](#)[12](#)
- [Copyright Statement.....](#)[12](#)
- [Acknowledgment.....](#)[13](#)

1. Introduction

A requirement has been identified to identify a restricted subset of the Pseudowire Emulation Edge-to-Edge (PWE3) [[RFC 3985](#)] over Multi-Protocol Label Switching (MPLS) [[RFC3031](#)] pseudowire functionality suitable for the construction of transport networks. This document describes the IETF-recommended profile for such a configuration. In particular, it describes a profile of an [RFC4448](#) PWE3 Ethernet pseudowire [[RFC4448](#)] that can be used to address these requirements. The scope of the initial version of this profile is restricted to transporting client Ethernet over an MPLS Packet Switched Network (PSN).

The architecture required for this configuration is illustrated in Figure 1 below. It is based on requirements described in the Requirements below.

3. OAM

The OAM mechanism is VCCV [[VCCV](#)].

One of the following VCCV profiles MUST be used. Once one is provisioned and the operational status of the PW is UP, no other profile SHOULD be used until such time as the PW's operational status is set to DOWN.

3.1. VCCV profile 1: BFD without IP/UDP Headers

As specified in [[VCCV](#)], the first nibble is set to 0001b to indicate a channel associated with a pseudowire [[RFC4385](#)][RFC4446]. The Version and the Reserved fields are set to 0, the Version is 0, and the Channel Type is set to 0x07 to indicate that the payload carried in BFD without IP/UDP headers, as is defined in section 4.1.1 of [[VCCV](#)].

The connection verification method used by VCCV is BFD with diagnostics as defined in 4.1 of [[VCCV](#)].

3.2. VCCV profile 2: BFD with IP/UDP Headers

When PE1 and PE1 are IP capable and have been configured with IP addresses, the following VCCV mechanism MAY be used.

As specified in [[VCCV](#)], the first nibble is set to 0001b to indicate a channel associated with a pseudowire [[RFC4385](#)][RFC4446]. The Version and the Reserved fields are set to 0, the Version is 0, and the Channel Type is set to 0x21 for IPv4 and 0x56 for IPv6 payloads.

The connection verification method used by VCCV is BFD with diagnostics as defined in 4.1 of [[VCCV](#)].

4. MPLS Layer

This section describes the profile of the MPLS PSN [[RFC3031](#)]. The requirements are based on the requirements communicated to the IETF and included in Appendix 1. The profile considers two cases:

1. The case where external configuration is used
2. The case where a control plane is available

4.1. External Configuration

All MPLS labels [[RFC3032](#)] used by the transport LSPs utilized by the PWs described in sections [2](#) and [3](#) MUST be statically provisioned. Labels may be selected from the per-platform or per-interface label space.

All LSPs for the transport LSPs utilized by the PWs described in sections [2](#) MUST support both unidirectional and bi-directional point-to-point connections.

The transport LSPs SHOULD support unidirectional point-to-multipoint connections.

The forward and backward directions of a bi-directional connection should follow the same path along the transport LSP.

Equal cost multi-path (ECMP) load balancing MUST NOT be configured on the transport LSPs utilized by the PWs described in sections [2](#) and [3](#).

The Merging of label switched paths is prohibited and MUST NOT be configured the transport LSPs utilized by the PWs described in sections [2](#) and [3](#).

Penultimate hop popping by the LSRs MUST be disabled on LSPs providing PWE3 transport network functionality.

Both E-LSP and L-LSP MUST be supported as defined in [RFC3270](#) [[RFC3270](#)].

The MPLS EXP field is supported according to [RFC3270](#) for only when the pipe and short-pipe models are utilized.

4.2. Control Plane Configuration

In this section we describe the profile when RSVP-TE [] or the bi-directional support in GMPLS [] are used to configure the MPLS PSN used to provide the transport LSPs. When these protocols are used to provide the control plane the following are automatically provided:

1. There is no label merging unless it is deliberately enabled to support Fast Re-route (FRR)[].
2. A single path is provided end-to-end (There is no ECMP).
3. Paths may be unidirectional or bidirectional as required.

Additionally the following configurations restrictions required to support external configuration MUST be applied:

Penultimate hop popping by the LSRs MUST be disabled on LSPs providing PWE3 transport network functionality.

Both E-LSP and L-LSP MUST be supported as defined in [RFC3270](#) [[RFC3270](#)].

The MPLS EXP field is supported according to [RFC3270](#) for only when the pipe and short-pipe models are utilized.

5. Congestion Considerations

This draft is a profile of an [RFC4448](#) PWE3 Ethernet pseudowire. The security considerations associated with that pseudowire and all subsequent work on congestion considerations regarding Ethernet pseudowires is applicable to this draft.

6. Security Considerations

PWE3 security considerations are described in [RFC3985](#) [[RFC3985](#)].

This draft is a profile of existing IETF proposed standards and raises no new security issues.

7. IANA Considerations

There are no IANA actions required by this draft.

8. Acknowledgments

APPENDIX A: Requirements

This appendix is NOT normative.

The following are the requirements for a transport pseudowire and are based on inputs from ITU SG15 [[ITU1](#)]. These in turn are the result of ITU studies on Transport MPLS [[ITU2](#)].

1. A transport pseudowire shall be based on a connection-oriented packet switched technology. Transport networks can also support circuit switched transport technologies (e.g. SDH, OTH or WDM)
2. A transport pseudowire shall operate under a common operation, control and management paradigm with respect to other transport technologies (e.g. SDH, OTN or WDM)
3. A transport pseudowire network shall support multiple layer network instances; e.g. a TRANSPORT PSEUDOWIRE transport network "service layer" instance in which the connections carry the customer's (individual or bundled) service, a TRANSPORT PSEUDOWIRE transport network "trunk layer" instance in which the connections carry aggregates of the network "service layer" connections and a transmission media layer instance in which the connections carry the aggregate of network trunk or network service connections.
4. The identification of each connection within its aggregate shall be based on labels. Connections can be aggregated into trunks by pushing and de-aggregated by popping labels. TRANSPORT PSEUDOWIRE labels may be swapped within a connection in a layer network instance when the traffic is forwarded from one TRANSPORT PSEUDOWIRE link connection to another TRANSPORT PSEUDOWIRE link connection. TRANSPORT PSEUDOWIRE pop, push, and swap label operations should be performed as defined by the relevant IETF RFCs.
5. Labeling shall make use of the MPLS label and label stack entry as defined in [RFC3032](#).
6. A transport pseudowire layer network shall support TRANSPORT PSEUDOWIRE and non TRANSPORT PSEUDOWIRE client layer networks and shall be able to be carried over TRANSPORT PSEUDOWIRE and non TRANSPORT PSEUDOWIRE server layer networks.
7. A transport pseudowire transport plane shall be able to operate without any IP functionality present.

8. A transport pseudowire network shall be able to be operated with a centralized NMS system
9. A transport pseudowire network should be able to be operated by a centralized NMS system with the support of a distributed control plane
10. A transport pseudowire shall support both unidirectional and bi-directional point-to-point connections
11. A transport pseudowire should support unidirectional point-to-multipoint connections
12. The forward and backward directions of a bi-directional connection should follow the same path along the TRANSPORT PSEUDOWIRE network.
13. The intermediate nodes should be aware about the binding of the forward and the backward directions belonging to the same bi-directional connection.
14. A transport pseudowire shall support traffic-engineering capabilities with traffic- and/or resource-oriented performance objectives.
15. A transport pseudowire shall support a method to offer packet loss objectives comparable to those in TDM transport networks (only due to bit errors)
16. A transport pseudowire should support transport and QoS mechanisms that can deliver statistical multiplexing gain. Packets exceeding the agreed traffic profile are however not allowed to enter into the TRANSPORT PSEUDOWIRE network (i.e. they are discarded by the traffic conditioning at the ingress of the network)
17. A transport pseudowire layer network shall operate independently of other layer networks (either TRANSPORT PSEUDOWIRE or non TRANSPORT PSEUDOWIRE).
18. A transport pseudowire shall support connections through a single domain
19. A transport pseudowire should support connections through multiple domains

20. A transport pseudowire shall offer as much commonality as possible with the MPLS user plane as defined by IETF. When MPLS offers multiple options in this respect, TRANSPORT PSEUDOWIRE should select the minimum sub-set (necessary and sufficient subset) applicable to a transport network application.

OAM Requirements:

- 21. A transport pseudowire should support transport network connection and performance monitoring mechanisms
- 22. A transport pseudowire OAM shall support fault management, performance management and protection switching mechanisms
- 23. A transport pseudowire OAM shall be able to operate without any IP functionality present.
- 24. Survivability Requirements
- 25. A transport pseudowire should support transport network-style protection switching mechanisms (sub-network connection protection and trail protection)
- 26. A transport pseudowire should support transport network restoration mechanisms
- 27. A transport pseudowire protection switching shall be able to operate without any IP functionality present.

Control Plane Requirements:

Control plane requirements are for further study.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[RFC3031] E. Rosen, A. Viswanathan, R. Callon., "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", [RFC 3032](#), January 2001.
- [RFC3270] F. Le Faucheur, Editor "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", [RFC3270](#), May 2002
- [RFC4385] S. Bryant, G. Swallow, L. Martini, D. McPherson, "Control Word for Use over an MPLS PSN", [RFC 4385](#), February 2006.
- [RFC4446] L. Martini, "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)" [RFC4446](#), April 2006.
- [RFC4447] L. Martini, Ed. "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", [RFC4447](#), April 2006.
- [RFC4448] L. Martini, Ed., E. Rosen, N. El-Aawar, G. Heron, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", [RFC 4448](#), April 2006
- [VCCV] T. Nadeau (Ed), "Pseudo Wire Virtual Circuit Connectivity Verification (VCCV)", [draft-ietf-pwe3-vcv-11.txt](#) (work in progress), October 2006.

9.2. Informative References

- [ITU1]
https://datatracker.ietf.org/public/liaison_detail.cgi?detail_id=265
- [ITU2]
http://ties.itu.int/u/tsg15/sg15/xchange/wp3/q12/0609_sophia/wd/wd09r3_editor_g8110.1v1_0909.doc
- [RFC3985] Bryant, S. and P. Pate, "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", [RFC3985](#), March 2005.

Author's Addresses

Stewart Bryant
Cisco Systems
250, Longwater,
Green Park,
Reading, RG2 6GB, UK
Email: stbryant@cisco.com

Monique Morrow
Cisco Systems, Inc.
Glatt-com
CH-8301 Glattzentrum
Switzerland
Email: mmorrow@cisco.com

Rao Cherukuri
Juniper Networks
1194 N. Mathilda Ave
Sunnyvale, CA 94089

Thomas D. Nadeau
Cisco Systems
300 Beaver Brook Drive
Boxborough, MA USA
Email: tnadeau@cisco.com

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719
Email: swallow@cisco.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.