

Network Working Group  
Internet Draft  
Expiration Date: January 2005

S. Bryant  
Cisco Systems  
D. McPherson  
Arbor Networks

July 2004

PWE3 Control Word  
draft-bryant-mcpherson-pwe3-cw-00.txt

## Status of this Memo

By submitting this Internet-Draft, we certify that any applicable patent or other IPR claims of which we are aware have been disclosed, or will be disclosed, and any of which we become aware will be disclosed, in accordance with [RFC 3668](#).

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 a "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

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

## Abstract

This document describes the preferred designs of the PWE3 Control Word, and the PWE3 Payload Type Identifier. The design of these fields is chosen so that an MPLS LSR performing deep packet inspection will not confuse a PWE3 payload with an IP payload.

INTERNET DRAFT

PWE3 Control Word

July 2004

## 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](#)].

## 1. Introduction

Packets are carried in MPLS label stacks without any protocol identifier. In order for a pseudo wire (PW) [ARCH] to operate correctly over an MPLS PSN that performs deep packet inspection, a PW packet must not appear to the LSR as if it were an IP packet [BCP]. An example of an LSR that performs deep packet inspection is one that is performing equal-cost multiple-path load-balancing (ECMP) [[RFC2992](#)]. If ECMP were performed on PWE3 packets, the packets in the PW may not all follow the same path though the PSN. This may result in misordered packet deliver to the egress PE. The inability to ensure that all packets belonging to a PW follow the same path also prevents the PW OAM [VCCV] mechanism from correctly monitoring the PW.

This draft specifies how a PW Control Word distinguishes a PW MPLS payload from an IP MPLS payload.

## 2. PWE3 Packet Identification

All IP packets [[RFC791](#)] [RFC1883] start with a version number which is checked by LSRs performing packet inspection. Therefore, PWE3 packets carried over an MPLS PSN SHOULD NOT start with the value 4 or the value 6 in the first nibble [BCP].

A PW SHOULD employ either the generic PW Control Word described in [Section 3](#), or the PWE3 Payload Type Identifier (PWE3-PTI) described in [Section 4](#). These fields MUST immediately follow the bottom of the MPLS label stack.

If the first nibble of a PWE3 packet carried over an MPLS PSN has a value of 0, it starts with a Generic PW Control Word. If the first nibble of a packet carried over an MPLS PSN has a value of 1, it starts with a Payload Type Identifier. The use of any other first nibble value for a PWE3 packet is deprecated.

## 3. Generic PW Control Word

The Generic PW Control Word is shown in Figure 1.

INTERNET DRAFT

PWE3 Control Word

July 2004




Figure 1: Generic PW Control Word

The PW set-up protocol or configuration mechanism determines whether a PW uses a Control Word. Bits 0..3 differ from the first four bits of an IP packet [BCP] and hence provide the necessary MPLS payload discrimination.

When a Control Word is used, it SHOULD have the following preferred form:




Figure 2: MPLS Preferred Control Word

The meaning of the fields of the MPLS Preferred Control Word (Figure 2) are as follows:

Flags (bits 4 to 7):

These bits are available for per payload signalling. Their definition is encapsulation specific.

FRG (bits 8 and 9):

These bits are used when fragmenting a PW payload. Their use

is defined in [FRAG]. When the PW is of a type that will never need payload fragmentation, these bits may be used as general purpose flags.

Length (bits 10 to 15):

The length field is used to determine the size of a PW payload that might have been padded to the minimum Ethernet MAC frame size during its transit across the PSN. If the MPLS payload (defined as the CW + the PW payload + any additional PW headers) is less than 46 bytes, the length MUST be set to the length of the MPLS payload. If the MPLS payload is between 46 bytes and 63 bytes the implementation

MAY either set to the length to the length of the MPLS payload, or it MAY set it to 0. If the length of the MPLS payload is greater than 63 bytes the length MUST be set to 0.

[Editor's note: Both the MUSTs are needed to make the mechanism work, the MAY is for backwards compatibility with deployed systems]

Sequence number (Bit 16 to 31):

If the sequence number is not used, it is set to zero by the sender and ignored by the receiver. Otherwise it specifies the sequence number of a packet. A circular list of sequence numbers is used. A sequence number takes a value from 1 to 65535 ( $2^{16}-1$ ). The sequence number window size for packet acceptance is dependent on the parameters of the PSN, and SHOULD be configurable. The mechanism used by the decapsulating PE to (re)acquire the correct sequence number is implementation dependent.

If the payload is an OAM packet the sequence number MAY be used to mark the position in the sequence, in which case it has the same value as the last data PDU sent. The use of the sequence number is optional for OAM payloads.

#### [4.](#) PWE3 Payload Type Identifier

If technical considerations result in a PW Control Word that could be mistaken for an IP packet, the Control Word SHOULD be preceded by a PWE3 Payload Type Identifier (PWE3-PTI). The PWE3-PTI is defined

follows:




Figure 3: PWE3 Payload Type Identifier

The meaning of the fields of the PWE3-PTI (Figure 3) are as follows:

Payload Type:

The PW Type as defined in the IANA PW Type registry

Bits 4 to 15 inclusive are reserved for future use and must be zero. Bits 0..3 MUST be 0x01, and hence differ from the first four bits of an IP packet [BCP]. This provides the necessary MPLS payload discrimination.

5. IANA considerations

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to the PW-Type, in accordance with [BCP 26](#) [[RFC 2424](#)].

There is one namespace that requires allocation, the PW-Type value.

5.1 Definition of Terms

The following terms are used here with the meanings defined in [BCP 26](#): "name space", "assigned value", "registration".

The following policies are used here with the meanings defined in [BCP 26](#): "Private Use", "First Come First Served", "Expert Review", "Specification Required", "IETF Consensus", "Standards Action".

NOTE NEED TO UPDATE ABOVE ONCE SECTION IS COMPLETE

5.2 Recommended Registration Policies

For registration requests where a Designated Expert should be consulted, an IESG Area Director for the Internet Area should appoint the Designated Expert.

For registration requests requiring Expert Review, the PWE3 mailing list should be consulted.

PW-Type codes have a range from x to y. Because a new Packet Type has considerable impact on interoperability, a new PW-Type code requires Standards Action, and should be allocated starting at TBD.

PW-Types codes have a range from x to y, and are the scarcest resource in PWE3, thus they must be allocated with care.

PW codes k-j may be allocated following Expert Review, with Specification Required.

The values v to x are reserved for vendor specific or experimental use.

## 6. Security Considerations

No new security issues arise as a result of the work.

## Normative References

Internet-drafts are works in progress available from  
<http://www.ietf.org/internet-drafts/>

[RFC791] [RFC-791](#): DARPA Internet Program, Protocol Specification, ISI, September 1981.

[RFC1883] [RFC-1883](#): Internet Protocol, Version 6 (IPv6), S. Deering, et al, December 1995

[RFC2992] [RFC-2992](#): Analysis of an Equal-Cost Multi-Path Algorithm, C. Hopps, November 2000

[RFC2424] [RFC-2424](#): Guidelines for Writing an IANA Considerations Section in RFCs, Alvestrand and

Narten, October 1998.

## Informative References

Internet-drafts are works in progress available from  
<<http://www.ietf.org/internet-drafts/>>

ARCH Bryant, S., Pate, P., "PWE3 Architecture", Internet Draft, <[draft-ietf-pwe3-arch-07.txt](#)>, October 2003, Work in Progress.

BCP Swallow, G. et al, "Avoiding Equal Cost Multipath Treatment in MPLS Networks", Internet Draft <[draft-swallow-etc](#)>, To be published July 2004, Work in Progress.

FRAG Malism, A., Townsley, M., "PWE3 Fragmentation and Reassembly", Internet Draft, <[draft-ietf-pwe3-fragmentation-05.txt](#)>, February 2004, Work in Progress.

VCCV Nadeau, T., Aggarwal, T., "Pseudo Wire (PW) Virtual Circuit Connection Verification (VCCV)", Internet Draft, <[draft-ietf-pwe3-vccv-02.txt](#)>, February 2004, Work in Progress.

Bryant et al

Expires January 2005

[Page 6]

---

INTERNET DRAFT

PWE3 Control Word

July 2004

## Authors' Addresses

Stewart Bryant  
Cisco Systems,  
250, Longwater,  
Green Park,  
Reading, RG2 6GB,  
United Kingdom.

Email: [stbryant@cisco.com](mailto:stbryant@cisco.com)

Danny McPherson  
Arbor Networks

Email: [danny@arbor.net](mailto:danny@arbor.net)

## Full copyright statement

Copyright (C) The Internet Society (2004). 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.

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.