

Internet Engineering Task Force (IETF)
Request for Comments: 8339
Category: Standards Track
ISSN: 2070-1721

P. Jain, Ed.
Cisco Systems, Inc.
S. Boutros
VMWare, Inc.
S. Aldrin
Google Inc.
March 2018

Definition of P2MP PW TLV for Label Switched Path (LSP) Ping Mechanisms

Abstract

Label Switched Path (LSP) Ping is a widely deployed Operation, Administration, and Maintenance (OAM) mechanism in MPLS networks. This document describes a mechanism to verify connectivity of Point-to-Multipoint (P2MP) Pseudowires (PWs) using LSP Ping.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8339>.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	3
2. Terminology	3
2.1. Specification of Requirements	3
2.2. Abbreviations	4
3. Identifying a P2MP PW	5
3.1. P2MP Pseudowire Sub-TLV	5
4. Encapsulation of OAM Ping Packets	6
5. Operations	6
6. Controlling Echo Responses	7
7. Security Considerations	7
8. IANA Considerations	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Acknowledgments	10
Authors' Addresses	10

1. Introduction

A Point-to-Multipoint (P2MP) Pseudowire (PW) emulates the essential attributes of a unidirectional P2MP Telecommunications service such as P2MP ATM over a Public Switched Network (PSN). Requirements for P2MP PWs are described in [RFC7338]. P2MP PWs are carried over a P2MP MPLS LSP. The procedures for P2MP PW signaling using BGP are described in [RFC7117]; LDP for single segment P2MP PWs is described in [RFC8338]. Many P2MP PWs can share the same P2MP MPLS LSP; this arrangement is called an "Aggregate P2MP Tree". An Aggregate P2MP Tree requires an upstream-assigned label so that on the Leaf PE (L-PE), the traffic can be associated with a Virtual Private Network (VPN) or a Virtual Private LAN Service (VPLS) instance. When a P2MP MPLS LSP carries only one VPN or VPLS service instance, the arrangement is called an "Inclusive P2MP Tree". For an Inclusive P2MP Tree, the P2MP MPLS LSP label itself can uniquely identify the VPN or VPLS service being carried over the P2MP MPLS LSP. The P2MP MPLS LSP can also be used in the Selective P2MP Tree arrangement to carry multicast traffic. In a Selective P2MP Tree arrangement, traffic to each multicast group in a VPN or VPLS instance is carried by a separate unique P2MP LSP. In an Aggregate Selective P2MP Tree arrangement, traffic to a set of multicast groups from different VPN or VPLS instances is carried over the same shared P2MP LSP.

The P2MP MPLS LSPs are setup using either P2MP RSVP-TE [RFC4875] or Multipoint LDP (mDLP) [RFC6388]. Mechanisms for fault detection and isolation for data-plane failures for P2MP MPLS LSPs are specified in [RFC6425]. This document describes a mechanism to detect data-plane failures for P2MP PW carried over P2MP MPLS LSPs.

This document defines a new P2MP Pseudowire sub-TLV for the Target Forwarding Equivalence Class (FEC) Stack for P2MP PWs. The P2MP Pseudowire sub-TLV is added in the Target FEC Stack TLV by the originator of the echo request at the Root PE (R-PE) to inform the receiver at the Leaf PE (L-PE) of the P2MP PW being tested.

Support for multi-segment PWs is out of scope of this document.

2. Terminology

2.1. Specification of Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

ACH:	Associated Channel Header
AGI:	Attachment Group Identifier
ATM:	Asynchronous Transfer Mode
CE:	Customer Edge
FEC:	Forwarding Equivalence Class
GAL:	Generic Associated Channel Label
LDP:	Label Distribution Protocol
L-PE:	Leaf PE (one of many destinations of the P2MP MPLS LSP, i.e., egress PE)
LSP:	Label Switched Path
LSR:	Label Switching Router
mLDP:	Multipoint LDP
MPLS-OAM:	MPLS Operations, Administration, and Maintenance
P2MP:	Point-to-Multipoint
P2MP-PW:	Point-to-Multipoint Pseudowire
PE:	Provider Edge
PSN:	Public Switched Network
PW:	Pseudowire
R-PE:	Root PE (ingress PE, PE initiating P2MP PW setup)
RSVP:	Resource Reservation Protocol
TE:	Traffic Engineering
TLV:	Type, Length, Value
VPLS:	Virtual Private LAN Service

3. Identifying a P2MP PW

This document introduces a new LSP Ping Target FEC Stack sub-TLV, the P2MP Pseudowire sub-TLV, to identify the P2MP PW under test at the P2MP Leaf PE (L-PE).

3.1. P2MP Pseudowire Sub-TLV

The P2MP Pseudowire sub-TLV has the format shown in Figure 1. This TLV is included in the echo request sent over P2MP PW by the originator of the request.

The Attachment Group Identifier (AGI), as described in Section 3.4.2 of [RFC4446], in P2MP Pseudowire sub-TLV identifies the VPLS instance. The Originating Router's IP address is the IPv4 or IPv6 address of the P2MP PW root. The address family of the IP address is determined by the IP Addr Len field.




Figure 1: P2MP Pseudowire Sub-TLV Format

For Inclusive and Selective P2MP Trees, the echo request is sent using the P2MP MPLS LSP label.

For Aggregate Inclusive and Aggregate Selective P2MP Trees, the echo request is sent using a label stack of [P2MP MPLS LSP label, upstream assigned P2MP PW label]. The P2MP MPLS LSP label is the outer label and the upstream assigned P2MP PW label is the inner label.

4. Encapsulation of OAM Ping Packets

The LSP Ping echo request packet is encapsulated with the MPLS label stack as described in previous sections, followed by one of the two encapsulation options:

- o GAL [RFC6426] followed by an IPv4 (0x0021) or IPv6 (0x0057) type Associated Channel Header (ACH) [RFC4385]
- o PW ACH [RFC4385]

To ensure interoperability, implementations of this document MUST support both encapsulations.

5. Operations

In this section, we explain the operation of the LSP Ping over a P2MP PW. Figure 2 shows a P2MP PW PW1 setup from Root PE R-PE1, to Leaf PEs (L-PE2, L-PE3, and L-PE4). The transport LSP associated with the P2MP PW1 can be mLDP P2MP MPLS LSP or P2MP TE tunnel.

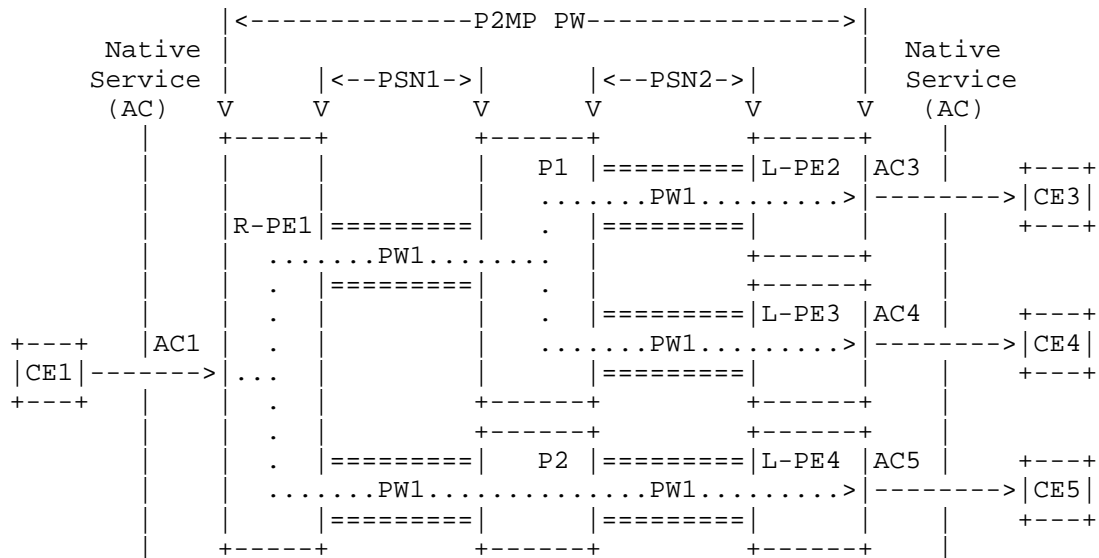


Figure 2: P2MP PW

When an operator wants to perform a connectivity check for the P2MP PW1, the operator initiates an LSP Ping echo request from Root PE R-PE1, with the Target FEC Stack TLV containing the P2MP Pseudowire sub-TLV in the echo request packet. For an Inclusive P2MP Tree arrangement, the echo request packet is sent over the P2MP MPLS LSP with one of the following two encapsulation options:

- o {P2MP LSP label, GAL} MPLS label stack and IPv4 or IPv6 ACH.
- o {P2MP LSP label} MPLS label stack and PW ACH.

For an Aggregate Inclusive Tree arrangement, the echo request packet is sent over the P2MP MPLS LSP with one of the following two encapsulation options:

- o {P2MP LSP label, P2MP PW upstream assigned label, GAL} MPLS label stack and IPv4 or IPv6 ACH.
- o {P2MP LSP label, P2MP PW upstream assigned label} MPLS label stack and PW ACH.

The intermediate P routers do MPLS label swap and replication based on the incoming MPLS LSP label. Once the echo request packet reaches L-PEs, L-PEs use the GAL and the IPv4/IPv6 ACH Channel header or PW ACH as the case may be, to determine that the packet is an OAM Packet. The L-PEs process the packet and perform checks for the P2MP Pseudowire sub-TLV present in the Target FEC Stack TLV as described in Section 4.4 in [RFC8029] and respond according to the processing rules in that document.

6. Controlling Echo Responses

The procedures described in [RFC6425] for preventing congestion of Echo Responses (Echo Jitter TLV in Section 3.3 of [RFC6425]) and limiting the echo reply to a single L-PE (Node Address P2MP Responder Identifier TLV in Section 3.2 of [RFC6425]) should be applied to P2MP PW LSP Ping.

7. Security Considerations

The proposal introduced in this document does not introduce any new security considerations beyond those that already apply to [RFC6425].

8. IANA Considerations

This document defines a new sub-TLV type included in the Target FEC Stack TLV (TLV Type 1) [RFC8029] in LSP Ping.

IANA has assigned the following sub-TLV type value from the "Sub-TLVs for TLV Types 1, 16, and 21" sub-registry within the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry:

37 P2MP Pseudowire

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC4385] Bryant, S., Swallow, G., Martini, L., and D. McPherson, "Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN", RFC 4385, DOI 10.17487/RFC4385, February 2006, <<https://www.rfc-editor.org/info/rfc4385>>.
- [RFC4446] Martini, L., "IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)", BCP 116, RFC 4446, DOI 10.17487/RFC4446, April 2006, <<https://www.rfc-editor.org/info/rfc4446>>.
- [RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", RFC 6425, DOI 10.17487/RFC6425, November 2011, <<https://www.rfc-editor.org/info/rfc6425>>.
- [RFC6426] Gray, E., Bahadur, N., Boutros, S., and R. Aggarwal, "MPLS On-Demand Connectivity Verification and Route Tracing", RFC 6426, DOI 10.17487/RFC6426, November 2011, <<https://www.rfc-editor.org/info/rfc6426>>.
- [RFC7117] Aggarwal, R., Ed., Kamite, Y., Fang, L., Rekhter, Y., and C. Kodeboniya, "Multicast in Virtual Private LAN Service (VPLS)", RFC 7117, DOI 10.17487/RFC7117, February 2014, <<https://www.rfc-editor.org/info/rfc7117>>.

- [RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8338] Boutros, S., Ed. and S. Sivabalan, Ed., "Signaling Root-Initiated Point-to-Multipoint Pseudowire Using LDP", RFC 8338, DOI 10.17487/RFC8338, March 2018, <<https://www.rfc-editor.org/info/rfc8338>>.

9.2. Informative References

- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.
- [RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <<https://www.rfc-editor.org/info/rfc6388>>.
- [RFC7338] Jounay, F., Ed., Kamite, Y., Ed., Heron, G., and M. Bocci, "Requirements and Framework for Point-to-Multipoint Pseudowires over MPLS Packet Switched Networks", RFC 7338, DOI 10.17487/RFC7338, September 2014, <<https://www.rfc-editor.org/info/rfc7338>>.

Acknowledgments

The authors would like to thank Shaleen Saxena, Greg Mirsky, Andrew G. Malis, and Danny Prairie for their valuable input and comments.

Authors' Addresses

Parag Jain (editor)
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, ON K2K-3E8
Canada

Email: paragj@cisco.com

Sami Boutros
VMWare, Inc.
United States of America

Email: sboutros@vmware.com

Sam Aldrin
Google Inc.
United States of America

Email: aldrin.ietf@gmail.com