PALS Workgroup
Internet-Draft
Intended status: Standards Track
Expires: February 22, 2018

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

August 21, 2017

Definition of P2MP PW TLV for LSP-Ping Mechanisms
draft-ietf-pals-p2mp-pw-lsp-ping-05

Abstract

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 (PW) using LSP Ping.

Table of Contents

1. Introduction

A Point-to-Multipoint (P2MP) Pseudowire (PW) emulates the essential attributes of a unidirectional P2MP Telecommunications service such as P2MP ATM over Public Switched Network (PSN). Requirements for P2MP PW are described in [RFC7338]. P2MP PWs are carried over P2MP MPLS LSP. The Procedures for P2MP PW signaling using BGP are described in [RFC7117] and LDP for single segment P2MP PWs are described in [I-D.ietf-pals-p2mp-pw]. Many P2MP PWs can share the same P2MP MPLS LSP and this arrangement is called 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 Inclusive P2MP Tree. For Inclusive P2MP Tree, P2MP MPLS LSP label itself can uniquely identify the VPN or VPLS service being carried over P2MP MPLS LSP. The P2MP MPLS LSP can also be used in Selective P2MP Tree arrangement for carrying 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 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 LSP are setup either using 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 Target FEC Stack for P2MP PW. The P2MP Pseudowire sub-TLV is added in Target FEC Stack TLV by the originator of the Echo Request at Root PE(R-PE) to inform the receiver at Leaf PE(L-PE) of the P2MP PW being tested.

Multi-segment Pseudowires support is out of scope of this document.

2. Specification of Requirements

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 [RFC2119].

3. Terminology

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

4. Identifying a P2MP PW

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

4.1. 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 request.

The Attachment Group Identifier (AGI) in P2MP Pseudowire Sub-TLV as described in Section 3.4.2 in [RFC4446], 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.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| AGI Type     |   AGI Length  |                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                                |
~                          AGI Value                            ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IP Addr Len |                                                 |
+-+-+-+-+-+-+-+                                                 |
~               Originating Routers IP Addr                     ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

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 upstream assigned P2MP PW label is inner label.

5. 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:

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

- PW ACH [RFC4385]

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

6. Operations

In this section, we explain the operation of the LSP Ping over 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.

```
                 |<--------------P2MP PW---------------->|
          Native |                                       |  Native
         Service |     |<--PSN1->|      |<--PSN2->|      |  Service
          (AC)   V     V         V      V         V      V   (AC)
            |    +-----+         +------+         +------+    |
            |    |     |         |   P1 |=========|L-PE2 |AC3 |    +---+
            |    |     |         |   .......PW1.........>|-------->|CE3|
            |    |R-PE1|=========|   .  |=========|      |    |    +---+
            |    |  .......PW1........  |         +------+    |
            |    |  .  |=========|   .  |         +------+    |
            |    |  .  |         |   .  |=========|L-PE3 |AC4 |    +---+
    +---+   |AC1 |  .  |         |   .......PW1.........>|-------->|CE4|
    |CE1|------->|...  |         |      |=========|      |    |    +---+
    +---+   |    |  .  |         +------+         +------+    |
            |    |  .  |         +------+         +------+    |
            |    |  .  |=========|   P2 |=========|L-PE4 |AC5 |    +---+
            |    |  .......PW1..............PW1.........>|-------->|CE5|
            |    |     |=========|      |=========|      |    |    +---+
            |    +-----+         +------+         +------+    |
```

Figure 2: P2MP PW

When an operator wants to perform a connectivity check for the P2MP PW1, the operator initiate a LSP-Ping request from Root PE R-PE1, with the Target FEC Stack TLV containing 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 GAL label 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 [RFC8029] processing rules.

7. 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 [RFC6425]) should be applied to P2MP PW LSP Ping.

8. Security Considerations

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

9. IANA Considerations

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

IANA is requested to assign a sub-TLV type value to the following sub-TLV from the "Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs) Parameters - TLVs" registry, "TLVs and sub- TLVs" sub-registry:

- o P2MP Pseudowire sub-TLV

10. Acknowledgments

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

11. References

11.1. Normative References

[I-D.ietf-pals-p2mp-pw]
Boutros, S. and S. Sivabalan, "Signaling Root-Initiated Point-to-Multipoint Pseudowire using LDP", draft-ietf-pals-p2mp-pw-03 (work in progress), June 2017.

[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>.

11.2. Informative 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>.

[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>.

Authors' Addresses

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

Email: paragj@cisco.com

Sami Boutros
VMWare, Inc.
USA

Email: sboutros@vmware.com

Sam Aldrin
Google Inc.
USA

Email: aldrin.ietf@gmail.com