

Internet Engineering Task Force (IETF)
Request for Comments: 7054
Category: Informational
ISSN: 2070-1721

A. Farrel
Juniper Networks
H. Endo
Hitachi, Ltd.
R. Winter
NEC
Y. Koike
NTT
M. Paul
Deutsche Telekom
November 2013

Addressing Requirements and Design Considerations for
Per-Interface Maintenance Entity Group Intermediate Points (MIPs)

Abstract

The framework for Operations, Administration and Maintenance (OAM) within the MPLS Transport Profile (MPLS-TP) describes how the Maintenance Entity Group Intermediate Points (MIPs) may be situated within network nodes at incoming and outgoing interfaces.

This document elaborates on important considerations for internal MIP addressing. More precisely, it describes important restrictions for any mechanism that specifies a way of forming OAM messages so that they can be targeted at MIPs on either incoming or outgoing interfaces and forwarded correctly through the forwarding engine. Furthermore, the document includes considerations for node implementations where there is no distinction between the incoming and outgoing MIP.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

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). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7054>.

Copyright Notice

Copyright (c) 2013 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 (http://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. Introduction2
- 2. Terminology3
- 3. Summary of the Problem Statement3
- 4. Requirements and Design Considerations for Internal-MIP Addressing6
- 5. Security Considerations10
- 6. Acknowledgments10
- 7. References10
 - 7.1. Normative References10
 - 7.2. Informative References11

1. Introduction

The framework for Operations, Administration and Maintenance (OAM) within the MPLS Transport Profile (MPLS-TP)(the MPLS-TP OAM Framework, [RFC6371]) distinguishes two configurations for the Maintenance Entity Group Intermediate Points (MIPs) on a node. It defines per-node MIPs and per-interface MIPs, where a per-node MIP is a single MIP per node in an unspecified location within the node and per-interface MIPs are two (or more) MIPs per node on each side of the forwarding engine.

In-band OAM messages are sent using the Generic Associated Channel (G-ACh) [RFC5586]. OAM messages for the transit points of pseudowires (PWs) or Label Switched Paths (LSPs) are delivered using the expiration of the MPLS shim header Time-to-Live (TTL) field. OAM messages for the end points of PWs and LSPs are simply delivered as normal.

OAM messages delivered to end points or transit points are distinguished from other (data) packets so that they can be processed as OAM. In LSPs, the mechanism used is the presence of the Generic Associated Channel Label (GAL) in the Label Stack Entry (LSE) under the top LSE [RFC5586]. In PWs, the mechanism used is the presence of the PW Associated Channel Header (PWACH) [RFC4385] or the presence of a GAL [RFC6423].

If multiple MIPs are present on a single node, these mechanisms alone provide no way to address one particular MIP out of the set of MIPs. A mechanism that addresses this shortcoming has to obey a few important design considerations, which are discussed in this document.

2. Terminology

In this document, we use the term in-MIP (incoming MIP) to refer to the MIP that processes OAM messages before they pass through the forwarding engine of a node. An out-MIP (outgoing MIP) processes OAM messages after they have passed the forwarding engine of the node. The two together are referred to as internal MIPs. The term "forwarding engine" is used as defined in [RFC6371].

Note that the acronym "OAM" is used in conformance with [RFC6291].

3. Summary of the Problem Statement

Figure 1 shows an abstract functional representation of an MPLS-TP node. It is decomposed as an incoming interface (labeled "In"), a forwarding engine (FW), and an outgoing interface (labeled "Out"). As per the discussion in [RFC6371], MIPs may be placed in each of the functional interface components.




Figure 1: Abstract Functional Representation of an MPLS-TP Node

Several distinct OAM functions are required within this architectural model for both PWs and LSPs, such as:

- o Connectivity Verification (CV) between a Maintenance Entity Group End Point (MEP) and a MIP
- o traceroute over an MPLS-TP LSP and/or an MPLS-TP PW containing MIPs
- o data-plane loopback configuration at a MIP
- o diagnostic tests

The MIPs in these OAM functions may also be the MIPs at the incoming or outgoing interfaces.

Per-interface MIPs have the advantage that they enable a more accurate localization and identification of faults and diagnostic tests. In particular, the identification of whether a problem is located between nodes or on a particular node and where on that node is greatly enhanced. For obvious reasons, it is important to narrow down the cause of a fault quickly to initiate a timely and well-directed maintenance action to resume normal network operation.

The following two figures illustrate the fundamental difference between using per-node and per-interface MEPs and MIPs for OAM. Figure 2 depicts OAM using per-node MIPs and MEPs. For reasons of exposition, we pick a location for the MIPs on the nodes but the standard does not mandate the exact location for the per-node model. In the figure, a bidirectional LSP is depicted where in the forward (FWD) direction MEP1, MIP1, and MEP2 are located on the ingress interface. In the backward (BWD) direction MEP1', MIP1' and MEP2' are located on the egress interface, i.e., the same interfaces. S1 in the figure denotes the segment from PE1 In to P1 In and S2 denotes the segment from PE1 In to P2 In. Figure 3, on the other hand, shows the same basic network, but per-interface maintenance points are configured for OAM operations. Note that these figures are merely examples. It is important to note that per-interface MEPs or per-interface MIPs must logically be placed at a point before (for in-MIP) or after (for out-MIP) passing the forwarding engine as defined in [RFC6371]. All traffic associated with the MEP/MIP must pass through or be terminated at that point.

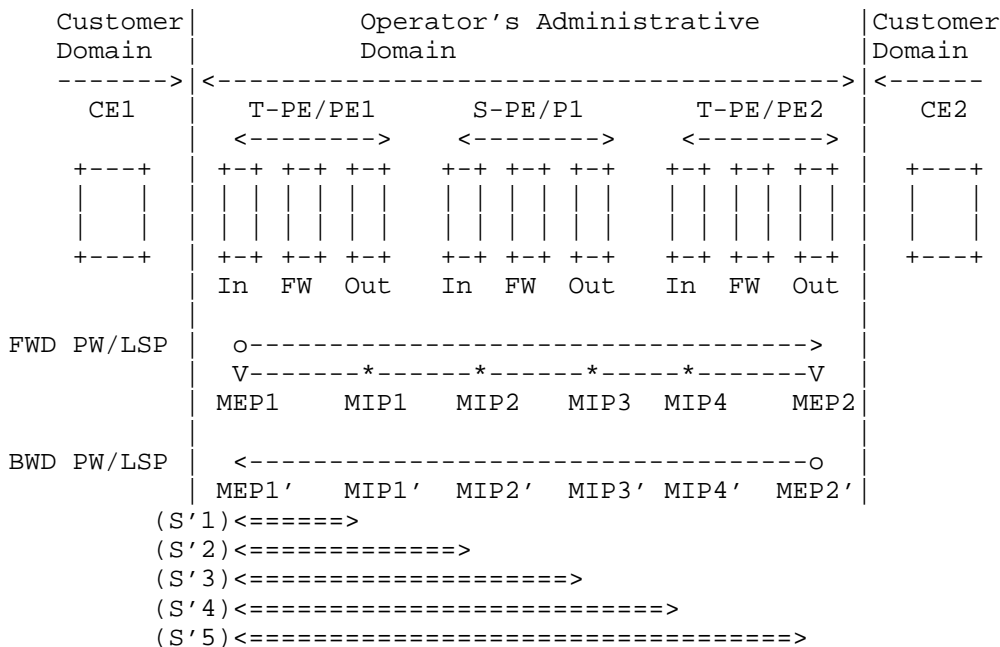


Figure 3: Example of OAM Relying on Per-Interface MIPs and MEPS

4. Requirements and Design Considerations for Internal-MIP Addressing

OAM messages for transit points of PWs or LSPs are delivered using the expiration of the TTL field in the top LSE of the MPLS packet header. OAM messages for the end points of PWs and LSPs are simply delivered as normal. These messages are distinguished from other (data) packets so that they can be processed as OAM. In LSPs, the mechanism used is the presence of the GAL in the LSE under the top LSE [RFC5586]. In PWs, the mechanism used is the presence of the PW Associated Channel Header [RFC4385] or the presence of a GAL [RFC6423]. In addition, two sets of identifiers exist that can be used to address MIPs, which are defined in [RFC6370] and [RFC6923]

Any solution for sending OAM messages to the in- and out-MIPs must fit within these existing models of handling OAM.

Additionally, many MPLS-TP nodes are implemented in a way that all queuing and the forwarding function are performed at the incoming interface. The abstract functional representation of such a node is shown in Figure 4. As shown in the figure, the outgoing interfaces are minimal and for this reason it may not be possible to include

MIP functions on those interfaces. This is the case for existing deployed implementations in particular.

Any solution that attempts to send OAM messages to the outgoing interface of an MPLS-TP node must not cause any problems when such implementations are present (such as leaking OAM packets with a TTL of 0).




Figure 4: Abstract Functional Representation of Some Existing MPLS-TP Nodes

OAM must operate on MPLS-TP nodes that are branch points on point-to-multipoint (P2MP) trees. This means that it must be possible to target individual outgoing MIPs as well as all outgoing MIPs in the abstract functional representation shown in Figure 5, and to handle the P2MP node implementations as shown in Figure 6 without causing problems.




Figure 5: Abstract Functional Representation of an MPLS-TP Node Supporting P2MP

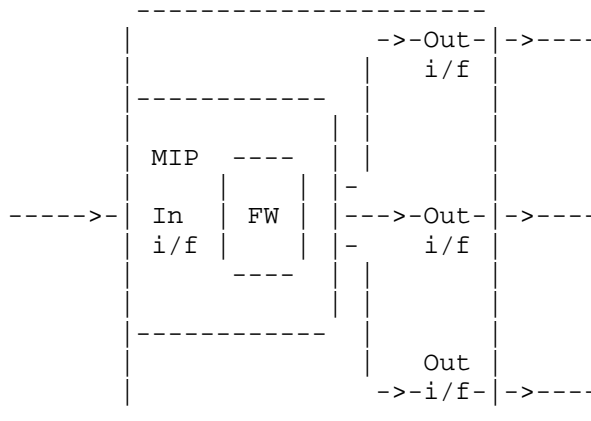


Figure 6: Abstract Functional Representation of Some Existing MPLS-TP Nodes Supporting P2MP

In summary, the solution for OAM message delivery must behave as follows:

- o Delivery of OAM messages to the correct MPLS-TP node.
- o Delivery of OAM instructions to the correct MIP within an MPLS-TP node.
- o Forwarding of OAM packets exactly as data packets.
- o Packet inspection at the incoming and outgoing interfaces must be minimized.

The first and second bullet points are obvious. However, the third bullet point is also vital. To illustrate the importance, a rejected solution is depicted in Figure 7. In the figure, all data and non-local OAM is handled as normal. Local OAM is intercepted at the incoming interface and delivered to the MIP at the incoming interface. If the OAM is intended for the incoming MIP, it is handled there with no issue. If the OAM is intended for the outgoing MIP, it is forwarded to that MIP using some internal messaging system that is implementation specific.




Figure 7: OAM Control Message Delivery Bypassing the Forwarding Engine

This solution is fully functional for the incoming MIP. It also supports control of data loopback for the outgoing MIP and can adequately perform some OAM features such as interface identity reporting at the outgoing MIP.

However, because the OAM message is not forwarded through the forwarding engine, this solution cannot correctly perform OAM loopback, connectivity verification, LSP tracing, or performance measurement.

The last bullet point is also an important requirement for any solution to the internal-MIP addressing problem. Since OAM packets that target an out-MIP need to be sent through the forwarding engine and treated exactly as regular data packets, the determination of whether to forward the packet or process it at the incoming MIP needs to be fast; therefore, the processing overhead must be kept to a minimum. In addition, there are a few OAM procedures that operate at line rate such as OAM loopback. This adds to the requirement of minimal processing overhead for both the in-MIP and out-MIP.

Most of the above superficially appears to be an implementation matter local to an individual node; however, the format of the message needs to be standardized so that:

- o A MEP can correctly target the outgoing MIP of a specific MPLS-TP node.
- o A node can correctly filter out any OAM messages that were intended for its upstream neighbor's outgoing MIP, but which were not handled there because the upstream neighbor is an implementation as shown in Figures 4 and 6.

Note that the last bullet point describes a safety net in case OAM messages leak beyond their intended scope, but implementations should take care that messages do not leak so that the safety net does not need to be used.

5. Security Considerations

OAM security is discussed in [RFC6371] and security aspects specific to MPLS-TP in general are outlined in [RFC6941].

OAM can provide useful information for detecting and tracing security attacks.

OAM can also be used to illicitly gather information or for denial-of-service attacks and other types of attack. Implementations therefore are required to offer security mechanisms for OAM. Deployments are therefore strongly advised to follow the security advice provided in RFCs 6371 and 6941.

Mixing of per-node and per-interface OAM on a single node is not advised as OAM message leakage could be the result.

6. Acknowledgments

The authors gratefully acknowledge the substantial contributions of Zhenlong Cui. We would also like to thank Eric Gray, Sami Boutros, and Shahram Davari for interesting input to this document through discussions.

7. References

7.1. Normative References

- [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, February 2006.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [RFC6370] Bocci, M., Swallow, G., and E. Gray, "MPLS Transport Profile (MPLS-TP) Identifiers", RFC 6370, September 2011.
- [RFC6371] Busi, I., Ed., and D. Allan, Ed., "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011.
- [RFC6423] Li, H., Martini, L., He, J., and F. Huang, "Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP)", RFC 6423, November 2011.

- [RFC6923] Winter, R., Gray, E., van Helvoort, H., and M. Betts,
"MPLS Transport Profile (MPLS-TP) Identifiers Following
ITU-T Conventions", RFC 6923, May 2013.

7.2. Informative References

- [RFC6291] Andersson, L., van Helvoort, H., Bonica, R., Romascanu,
D., and S. Mansfield, "Guidelines for the Use of the "OAM"
Acronym in the IETF", BCP 161, RFC 6291, June 2011.
- [RFC6941] Fang, L., Ed., Niven-Jenkins, B., Ed., Mansfield, S., Ed.,
and R. Graveman, Ed., "MPLS Transport Profile (MPLS-TP)
Security Framework", RFC 6941, April 2013.

Authors' Addresses

Adrian Farrel
Juniper Networks

EEmail: adrian@olddog.co.uk

Hideki Endo
Hitachi, Ltd.

EEmail: hideki.endo.es@hitachi.com

Rolf Winter
NEC

EEmail: rolf.winter@neclab.eu

Yoshinori Koike
NTT

EEmail: koike.yoshinori@lab.ntt.co.jp

Manuel Paul
Deutsche Telekom

EEmail: Manuel.Paul@telekom.de