

Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 3, 2012

A. Farrel
Huawei Technologies
H. Endo
Hitachi, Ltd.
R. Winter
NEC
Y. Koike
NTT
M. Paul
Deutsche Telekom
October 31, 2011

Handling MPLS-TP OAM Packets Targeted at Internal MIPS
draft-farrel-mpls-tp-mip-mep-map-05

Abstract

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

This document describes a way of forming OAM messages so that they can be targeted at MIPS on incoming or MIPS on outgoing interfaces, forwarded correctly through the "switch fabric", and handled efficiently in node implementations where there is no distinction between the incoming and outgoing MIP.

The material in this document is provided for discussion within the MPLS-TP community in the expectation that this idea or some similar mechanism will be subsumed into a more general MPLS-TP OAM document.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on May 3, 2012.

Copyright Notice

Copyright (c) 2011 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.	Introduction	4
2.	Requirements notation	5
3.	Terminology	5
4.	Summary of the Problem Statement	5
5.	Overview	7
5.1.	Rejected Partial Solution	10
6.	Per-Interface MIP Message Handling	11
7.	Security Considerations	12
8.	IANA Considerations	12
9.	Acknowledgments	12
10.	References	12
10.1.	Normative References	12
10.2.	Informative References	13
Appendix A.	Previously considered solutions	13
A.1.	GAL TTL	13
A.2.	A separate channel type for the out-MIP	13
A.3.	Decrement TTL once per MIP	13
A.4.	Using an ACH reserved bit	14
	Authors' Addresses	14

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 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 both sides 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) [\[RFC4385\]](#).

In case 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.

This document describes a way of forming OAM messages so that they can be targeted at incoming MIPs and outgoing MIPs, forwarded correctly through the "switch fabric", and handled efficiently in node implementations where there is no distinction between the incoming and outgoing MIP.

The material in this document is provided for discussion within the MPLS-TP community in the expectation that this idea or some similar mechanisms will be subsumed into a more general MPLS-TP OAM document.

This document is a product of a joint Internet Engineering Task Force (IETF)/International Telecommunication Union Telecommunication Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architecture to support the capabilities and functionalities of a packet transport network.

Note that the acronym "OAM" is used in conformance with [\[RFC6291\]](#).

2. Requirements notation

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

In this document we use the term in-MIP (incoming MIP) to refer to the MIP which 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.

4. Summary of the Problem Statement

Figure 1 shows an abstract functional representation of an MPLS-TP node. It is decomposed as an incoming interface, a cross-connect (XC), and an outgoing interface. 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 such as:

- o CV between a MEP and a MIP
- o traceroute over an MPLS-TP LSP and/or an MPLS-TP PW containing MIPs
- o OAM control at a MIP
- o data-plane loopback at a MIP
- o diagnostic tests

The MIPs in these OAM functions may equally 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 targeted performance monitoring or diagnostic test. 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 the cause of a fault down quickly to initiate a timely, and well-directed maintenance action to resume normal network operation.

The following two figures illustrate the fundamental difference of 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. Figure 3 on the other hand shows the same basic network but for OAM operations per-interface maintenance points are configured.

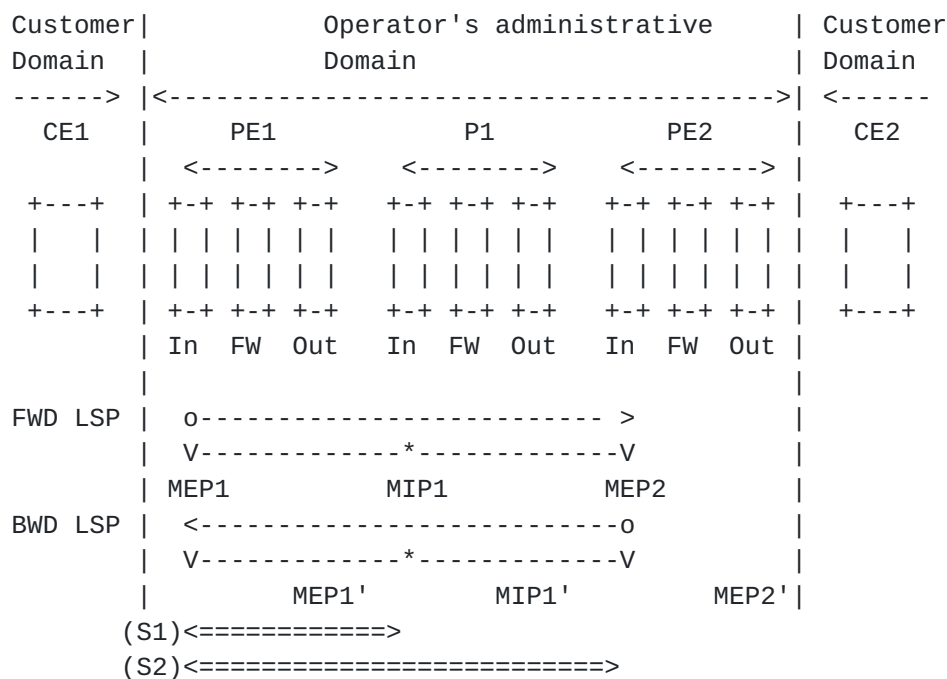


Figure 2: Example of OAM relying on per-node MIPs and MEPs

To illustrate the difference between these two modes of operation, we use fault detection as an example. Consider the case where the client traffic between CE1 and CE2 experiences a fault. Also assume that an on-demand CV test between PE1 and PE2 was successful. The scenario in Figure 2 therefore leaves the forwarding engine(FW) of PE2, the out-going interface of PE2, the transmission line between PE2 and CE2 or CE2 itself as a potential location of the fault as on-

demand CV can only be performed on segment S2.

The per-interface model in Figure 3 allows more fine-grained OAM operations to be performed. At first, CV on segment S'4 and in addition CV on segment S'5 can help to rule out e.g. the forwarding engine of PE2. This is of course only a single example, and other OAM functions and scenarios are trivially conceivable. The basic message is that with the per-interface OAM model, an operator can configure smaller segments on a transport path to which OAM operations apply. This enables a more fine-grained scoping of OAM operations such as fault localization and performance monitoring which gives operators better information to deal with adverse networking conditions.

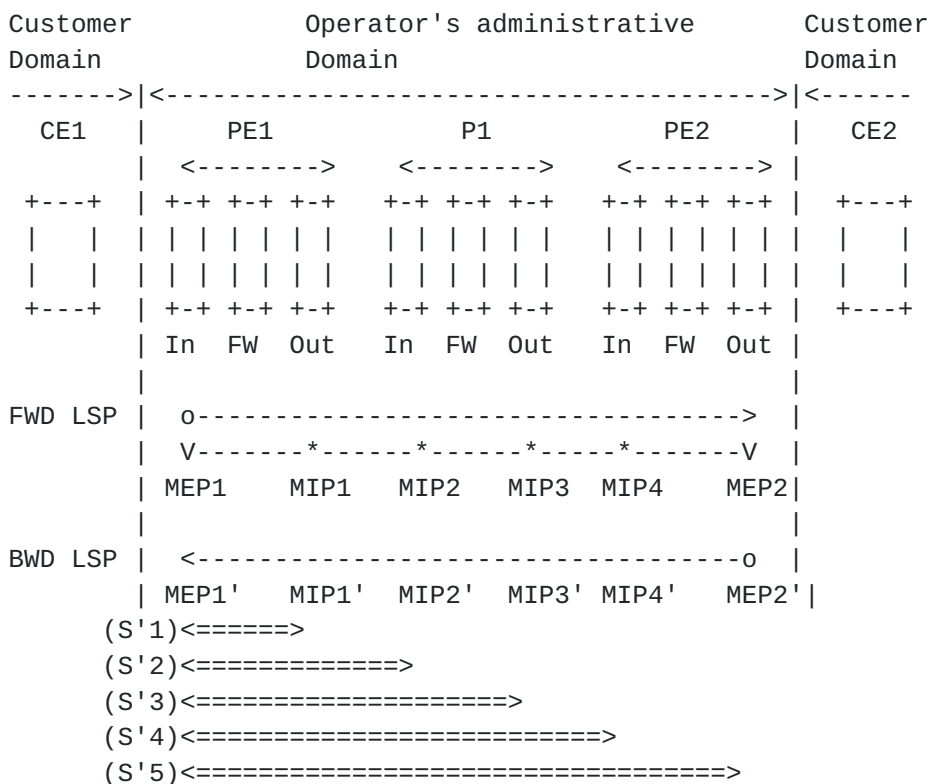


Figure 3: Example of OAM relying on per-interface MIPs and MEPs

5. Overview

In-band OAM messages are sent using the G-ACh [[RFC5586](#)] for MPLS-TP LSPs and MPLS-TP PWs, respectively. OAM messages for the transit points of PWs or LSPs are delivered using the expiration of the time-

to-live (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.

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 LSE under the top LSE [[RFC5586](#)]. In PWs, the mechanism used is the presence of the PW Associated Channel Header [[RFC4385](#)].

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

Additionally, many MPLS-TP nodes contain an optimization such that all queuing and the forwarding function is 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 in particular the case for existing deployed implementations.

Any solution that attempts to send OAM to the outgoing interface of an MPLS-TP node must not cause any problems when such implementations are present.

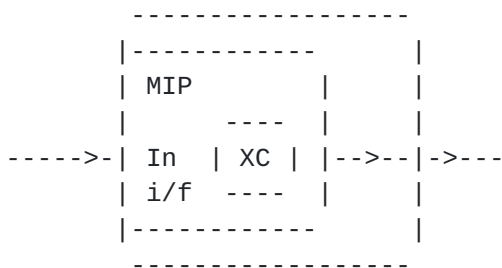


Figure 4: Abstract Functional Representation of an Optimized MPLS-TP Node

Lastly, OAM must operate on MPLS-TP nodes that are branch points on point-to-multipoint (P2MP) trees. That 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, as well as to handle the optimized P2MP node as shown in Figure 6.




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




Figure 6: Abstract Functional Representation of an Optimized MPLS-TP Node Supporting P2MP

In summary, the solution for OAM message delivery must support the following features:

- o Forwarding of OAM packets exactly as data packets.
- 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 Packet inspection at the incoming and outgoing interfaces must be minimized.

Note that although this issue appears superficially to be an implementation matter local to an individual node, the format of the message needs to be standardised 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 optimized implementation.

Note that the last bullet point describes a safety net and an implementation should avoid that this situation ever arises.

5.1. Rejected Partial Solution

A rejected solution is depicted in Figure 7. 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.

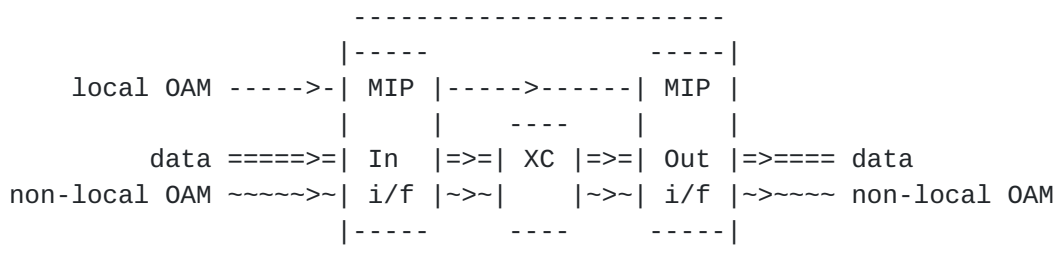


Figure 7: OAM Control Message Delivery Bypassing the Switching Fabric

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 switch fabric, this solution cannot correctly perform OAM loopback, connectivity verification, LSP tracing, or performance measurement.

6. Per-Interface MIP Message Handling

The preferred solution to per-interface MIP message handling is presented in this section. The appendix of this document contains a few solutions that the authors have discarded which have been left in the document for informational purposes.

Per-interface MIP addressing should not require changes to existing OAM solutions. Therefore, ID information that specific OAM mechanisms already contain should be (re-)used to address the MIPs on a node. Upcoming OAM solutions therefore need to individually make sure that enough of that information is present to support the per-interface model. In particular, the MIP identifiers as described in [\[RFC6370\]](#) need to be present in OAM messages. [\[RFC6370\]](#) defines a format that supports the per-interface model which is sufficient for this purpose. In addition, some constraints must be agreed on.

Regarding the features we required earlier from a solution this translates to:

- o Feature 1: "Forwarding of OAM packets exactly as data packets"
 - * This way of internal-MIP addressing has no implications on the way data packets and non-local OAM packets are handled as the label stack is not altered for the purpose of MIP addressing. Also the TTL processing remains untouched. This means that the TTL will expire on the ingress.
- o Feature 2: "Delivery of OAM messages to the correct MPLS-TP node"
 - * The node itself is addresses using TTL expiry, comparable to the per-node MIP addressing case.
- o Feature 3: "Delivery of OAM instructions to the correct MIP within an MPLS-TP node"
 - * The ID information contained in the OAM packet is used to tell whether the packet is for the in or out-MIP.
- o Feature 4: "Packet inspection at the incoming and outgoing interfaces must be minimized"
 - * Additional packet inspection compared to the per-node case is inevitably needed. The ID information inside the OAM message needs to be considered in order to deliver the packet to the right MIP.

The above illustrates how in principle per-MIP addressing operates. Another issue of concern was the correct filtering of OAM messages at a downstream node, that were intended for an upstream node's outgoing

MIP. Since OAM messages expire on the ingress, the legacy upstream neighbor will actually process the packet. Since the ID information is not correct, the node should discard that packet. Leakage should therefore not occur.

7. Security Considerations

OAM security is discussed in [[RFC6371](#)] and security aspects specific to MPLS-TP in general are outlined in [[I-D.ietf-mpls-tp-security-framework](#)].

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 strongly advised to use such mechanisms.

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

8. IANA Considerations

This revision of this document does not make any requests of IANA.

9. Acknowledgments

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

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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., Vigoureux, M., and S. Bryant, "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. and D. Allan, "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", [RFC 6371](#), September 2011.

[10.2. Informative References](#)

- [I-D.ietf-mpls-tp-security-framework]
Fang, L., Niven-Jenkins, B., and S. Mansfield, "MPLS-TP Security Framework",
[draft-ietf-mpls-tp-security-framework-01](#) (work in progress), May 2011.
- [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.

[Appendix A. Previously considered solutions](#)

[A.1. GAL TTL](#)

The use of the GAL TTL has been considered before. This transforms the GAL TTL into some kind of node-internal TTL, i.e. a GAL TTL of 0 would address the in-MIP and a GAL TTL of 1 the out-MIP. The main drawback of this approach is that it (as of now at least) would only be applicable to LSPs and not to PWs.

[A.2. A separate channel type for the out-MIP](#)

This approach would require two channel types for the exact same OAM type, one to address the in-MIP and another one to address the out-MIP. This seems like a waste of channel types, however it appears that there is no expected shortage of them. Legacy nodes will discard the packets as the new channel types are unknown. Having two channel types for the same OAM however feels a bit hacky.

[A.3. Decrement TTL once per MIP](#)

Decrementing the TTL more than once per node seems a "natural" way of per-interface MIP addressing since TTL expiry is all that is needed for the per-node MIP case. In other words, by decrementing the TTL once per MIP (twice per node) no extra mechanism is needed to solve the internal MIP addressing problem. The solution has been discarded

since it does not represent the typical mode of network operation today (since also for normal data packets the TTL needs to be decremented more than once).

A.4. Using an ACH reserved bit

The ACH contains eight reserved bits which currently all need to be set to zero and ignored on reception. One bit could be reserved as an out-MIP address flag. In other words, in case the bit is set, the out-MIP is addressed. An advantage of this approach is that there is no semantic overlap with anything that exists today, as the bits are not in use. Existing implementations need to ignore it. That means that existing implementations will process the OAM packets at the in-MIP/per-node MIP. ID information is still needed however for the P2MP case as a single bit is not enough.

Authors' Addresses

Adrian Farrel
Huawei Technologies

Email: adrian.farrel@huawei.com

Hideki Endo
Hitachi, Ltd.

Email: hideki.endo.es@hitachi.com

Rolf Winter
NEC

Email: rolf.winter@neclab.eu

Yoshinori Koike
NTT

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

Manuel Paul
Deutsche Telekom

Email: Manuel.Paul@telekom.de

