

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 2, 2013

IJ. Wijnands, Ed.
T. Eckert
Cisco Systems, Inc.
N. Leymann
Deutsche Telekom
M. Napierala
AT&T Labs
November 29, 2012

Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-
to-Multipoint Label Switched Paths
[draft-ietf-mpls-mldp-in-band-signaling-08](#)

Abstract

Consider an IP multicast tree, constructed by Protocol Independent Multicast (PIM), needs to pass through an MPLS domain in which Multipoint LDP (mLDP) Point-to-Multipoint and/or Multipoint-to-Multipoint Labels Switched Paths (LSPs) can be created. The part of the IP multicast tree that traverses the MPLS domain can be instantiated as a multipoint LSP. When a PIM Join message is received at the border of the MPLS domain, information from that message is encoded into mLDP messages. When the mLDP messages reach the border of the next IP domain, the encoded information is used to generate PIM messages that can be sent through the IP domain. The result is an IP multicast tree consisting of a set of IP multicast sub-trees that are spliced together with a multipoint LSP. This document describes procedures how IP multicast trees are spliced together with multipoint LSPs.

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 June 2, 2013.

Copyright Notice

Copyright (c) 2012 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	3
1.1.	Conventions used in this document	3
1.2.	Terminology	3
2.	In-band signaling for MP LSPs	4
2.1.	Transiting Unidirectional IP multicast Shared Trees	6
2.2.	Transiting IP multicast source trees	6
2.3.	Transiting IP multicast bidirectional trees	7
3.	LSP opaque encodings	7
3.1.	Transit IPv4 Source TLV	7
3.2.	Transit IPv6 Source TLV	8
3.3.	Transit IPv4 bidir TLV	9
3.4.	Transit IPv6 bidir TLV	9
4.	Security Considerations	10
5.	IANA considerations	10
6.	Acknowledgments	10
7.	Contributing authors	11
8.	References	11
8.1.	Normative References	11
8.2.	Informative References	12
	Authors' Addresses	12

1. Introduction

The mLDP (Multipoint LDP) [[RFC6388](#)] specification describes mechanisms for creating point-to-multipoint (P2MP) and multipoint-to-multipoint (MP2MP) LSPs (Label Switched Paths). These LSPs are typically used for transporting end-user multicast packets. However, the mLDP specification does not provide any rules for associating particular end-user multicast packets with any particular LSP. Other documents, like [[RFC6513](#)], describe applications in which out-of-band signaling protocols, such as PIM and BGP, are used to establish the mapping between an LSP and the multicast packets that need to be forwarded over the LSP.

This document describes an application in which the information needed to establish the mapping between an LSP and the set of multicast packets to be forwarded over it is carried in the "opaque value" field of an mLDP FEC (Forwarding Equivalence Class) element. When an IP multicast tree (either a source-specific tree or a bidirectional tree) enters the MPLS network the (S,G) or (*,G) information from the IP multicast control plane state is carried in the opaque value field of the mLDP FEC message. As the tree leaves the MPLS network, this information is extracted from the FEC element and used to build the IP multicast control plane. PIM messages can be sent outside the MPLS domain. Note that although the PIM control messages are sent periodically, the mLDP messages are not.

Each IP multicast tree is mapped one-to-one to a P2MP or MP2MP LSP in the MPLS network. A network operator should expect to see as many LSPs in the MPLS network as there are IP multicast trees. A network operator should be aware how IP multicast state is created in the network to ensure it does not exceed the scalability numbers of the protocol, either PIM or mLDP.

1.1. 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](#)].

1.2. Terminology

IP multicast tree : An IP multicast distribution tree identified by a IP multicast Group address and optionally a Source IP address, also referred to as (S,G) and (*,G).

RP: The PIM Rendezvous Point.

SSM: PIM Source Specific Multicast.

ASM: PIM Any Source Multicast.

mLDP : Multipoint LDP.

Transit LSP : A P2MP or MP2MP LSP whose FEC element contains the (S,G) or (*,G) identifying a particular IP multicast distribution tree.

In-band signaling : Using the opaque value of a mLDP FEC element to carry the (S,G) or (*,G) identifying a particular IP multicast tree.

P2MP LSP: An LSP that has one Ingress LSR and one or more Egress LSRs.

MP2MP LSP: An LSP that connects a set of leaf nodes that may each independently act as ingress or egress.

MP LSP: A multipoint LSP, either a P2MP or an MP2MP LSP.

Ingress LSR: Source of the P2MP LSP, also referred to as root node.

Egress LSR: One of potentially many destinations of an LSP, also referred to as leaf node in the case of P2MP and MP2MP LSPs.

Transit LSR: An LSR that has one or more directly connected downstream LSRs.

2. In-band signaling for MP LSPs

Consider the following topology;


```

|--- IPM ---|--- MPLS --|--- IPM ---|
S/RP -- (A) - (U) - (C) - (D) -- (B) -- R

```

Figure 1.

Nodes A and B are IP Multicast capable routers and respectively connect to a Source/RP and a Receiver. Nodes U, C and D are MPLS Label Switched Routers (LSRs).

Label Switched Router D is attached to a network that is capable of MPLS multicast and IP multicast (see figure 1), and D is required to create a IP multicast tree due to a certain IP multicast event, like a PIM Join, MSDP Source Announcement (SA) [[RFC3618](#)], BGP Source Active auto-discovery route [[I-D.rekhter-pim-sm-over-mldp](#)] or Rendezvous Point (RP) discovery. Suppose that D can determine that the IP multicast tree needs to travel through the MPLS network until it reaches LSR U. For instance, when D looks up the route to the Source or RP [[RFC4601](#)] of the IP multicast tree, it may discover that the route is a BGP route with U as the BGP next hop. Then D may chose to set up a P2MP or MP2MP LSP, with U as root, and to make that LSP become part of the IP multicast distribution tree. Note that other methods are possible to determine that an IP multicast tree is to be transported across an MPLS network using P2MP or MP2MP LSPs, these methods are outside the scope of this document.

In order to establish a multicast tree via a P2MP or MP2MP LSP using "in-band signaling", LSR D encodes a P2MP or MP2MP FEC Element, with the IP address of LSR U as the "Root Node Address", and with the source and the group encoded into the "opaque value" ([RFC6388](#), [section 2.2](#) and 3.2). Several different opaque value types are defined in this document; LSR D MUST NOT use a particular opaque value type unless it knows (through provisioning, or through some other means outside the scope of this document) that LSR U supports the root node procedures for that opaque value type.

The particular type of FEC Element and opaque value used depends on the IP address family being used, and on whether the multicast tree being established is a source specific or a bidirectional multicast tree.

When an LSR receives a label mapping or withdraw whose FEC Element contains one of the opaque value types defined in this document, and that LSR is not the one identified by the "Root Node Address" field of that FEC element, the LSR follows the procedures of [RFC 6388](#).

When an LSR receives a label mapping or withdraw whose FEC Element contains one of the opaque value types defined in this document, and

that LSR is the one identified by the "Root Node Address" field of that FEC element, then the following procedure is executed. The multicast source and group are extracted and passed to the multicast code. If a label mapping is being processed, the multicast code will add the downstream LDP neighbor to the olist of the corresponding (S,G) or (*,G) state, creating such state if it does not already exist. If a label withdraw is being processed, the multicast code will remove the downstream LDP neighbor from the olist of the corresponding (S,G) or (*,G) state. From this point on normal PIM processing will occur.

Note that if the LSR identified by the "Root Node Address" field does not recognize the opaque value type, the MP LSP will be established, but the root node will not send any multicast data packets on it.

Source or RP addresses that are reachable in a VPN context are outside the scope of this document.

Multicast groups that operate in PIM Dense-Mode are outside the scope of this document.

2.1. Transiting Unidirectional IP multicast Shared Trees

Nothing prevents PIM shared trees, used by PIM-SM in the ASM service model, from being transported across a MPLS core. However, it is not possible to prune individual sources from the shared tree without the use of an additional out-of-band signaling protocol, like PIM or BGP [[I-D.rekhter-pim-sm-over-mldp](#)]. For that reason transiting Shared Trees across a Transit LSP is outside the scope of this document.

2.2. Transiting IP multicast source trees

IP multicast source trees can either be created via PIM operating in SSM mode [[RFC4607](#)] or ASM mode [[RFC4601](#)]. When PIM-SM is used in ASM mode, the usual means of discovering active sources is to join a sparse mode shared tree. However, this document does not provide any method of establishing a sparse mode shared tree across an MPLS network. To apply the technique of this document to PIM-SM in ASM mode, there must be some other means of discovering the active sources. One possible means is the use of MSDP [[RFC3618](#)]. Another possible means is to use BGP Source Active auto-discovery routes, as documented in [[I-D.rekhter-pim-sm-over-mldp](#)]. However, the method of discovering the active sources is outside the scope of this document, and as a result this document does not specify everything that is needed to support the ASM service model using in-band signaling.

The source and group addresses are encoded into the a transit TLV as specified in [Section 3.1](#) and [Section 3.2](#).

2.3. Transit IP multicast bidirectional trees

If a Bidirectional IP multicast trees [[RFC5015](#)] has to be transported over a MPLS network using in-band signaling, as described in this document, it MUST be transported using a MP2MP LSPs. A bidirectional tree does not have a specific source address; the group address, subnet mask and RP are relevant for multicast forwarding. This document does not provide procedures to discover RP to group mappings dynamically across an MPLS network and assumes the RP is statically defined. Support of dynamic RP mappings in combination with in-band signaling is outside the scope of this document.

The RP for the group is used to select the ingress LSR and root of the LSP. The group address is encoded according to the rules of [Section 3.3](#) or [Section 3.4](#), depending on the IP version. The subnet mask associated with the bidirectional group is encoded in the Transit TLV. There are two types of bidirectional states in IP multicast, the group specific state and the RP state. The first type is typically created due to receiving a PIM join and has a subnet mask of 32 for IPv4 and 128 for IPv6. The latter is typically created via the static RP mapping and has a variable subnet mask. The RP state is used to build a tree to the RP and used for sender only branches. Each state (group specific and RP state) will result in a separate MP2MP LSP. The merging of the two MP2MP LSPs will be done by PIM on the root LSR. No special procedures are necessary for PIM to merge the two LSPs, each LSP is effectively treated as a PIM enabled interface. Please see [[RFC5015](#)] for more details.

For transporting the packets of a sender only branch we create a MP2MP LSP. Other sender only branches will receive these packets and will not forward them because there are no receivers. These packets will be dropped. If that affect is undesirable some other means of transport has to be established to forward packets to the root of the tree, like a Multi-Point to Point LSP for example. A technique to unicast packets to the root of a P2MP or MP2MP LSP is documented in [[I-D.rosen-l3vpn-mvpn-mspmi](#)] [section 3.2.2.1](#).

3. LSP opaque encodings

This section documents the different transit opaque encodings.

3.1. Transit IPv4 Source TLV


```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type           | Length           | Source           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               | Group           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 3 (to be assigned by IANA).

Length: 8 (octet size of Source and Group fields)

Source: IPv4 multicast source address, 4 octets.

Group: IPv4 multicast group address, 4 octets.

[3.2.](#) Transit IPv6 Source TLV

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type           | Length           | Type           | Source           | ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~               | Group           | ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 4 (to be assigned by IANA).

Length: 32 (octet size of Source and Group fields)

Source: IPv6 multicast source address, 16 octets.

Group: IPv6 multicast group address, 16 octets.

3.3. Transit IPv4 bidir TLV

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type           | Length           | Mask Len       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     RP                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Group                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: 5 (to be assigned by IANA).

Length: 9 (octet size of Mask Len, RP and Group fields)

Mask Len: The number of contiguous one bits that are left justified and used as a mask, 1 octet. Maximum value allowed is 32.

RP: Rendezvous Point (RP) IPv4 address used for encoded Group, 4 octets.

Group: IPv4 multicast group address, 4 octets.

3.4. Transit IPv6 bidir TLV

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Type           | Length           | Mask Len       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     RP                                     ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                                                                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Group                                ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


7. Contributing authors

Below is a list of the contributing authors in alphabetical order:

Toerless Eckert
Cisco Systems, Inc.
170 Tasman Drive
San Jose, CA, 95134
USA
E-mail: eckert@cisco.com

Nicolai Leymann
Deutsche Telekom
Winterfeldtstrasse 21
Berlin, 10781
Germany
E-mail: n.leymann@telekom.de

Maria Napierala
AT&T Labs
200 Laurel Avenue
Middletown, NJ 07748
USA
E-mail: mnapierala@att.com

IJsbrand Wijnands
Cisco Systems, Inc.
De kleetlaan 6a
1831 Diegem
Belgium
E-mail: ice@cisco.com

8. References

8.1. Normative References

- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-

Multipoint and Multipoint-to-Multipoint Label Switched Paths", [RFC 6388](#), November 2011.

8.2. Informative References

- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", [RFC 4601](#), August 2006.
- [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", [RFC 4607](#), August 2006.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", [RFC 5015](#), October 2007.
- [RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", [RFC 3618](#), October 2003.
- [RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), February 2012.
- [I-D.rekhter-pim-sm-over-mldp]
Rekhter, Y. and R. Aggarwal, "Carrying PIM-SM in ASM mode Trees over P2MP mLDP LSPs",
[draft-rekhter-pim-sm-over-mldp-04](#) (work in progress),
August 2011.
- [I-D.rosen-l3vpn-mvpn-mspmsi]
Cai, Y., Rosen, E., Wijnands, I., Napierala, M., and A. Boers, "MVPN: Optimized use of PIM via MS-PMSIs",
[draft-rosen-l3vpn-mvpn-mspmsi-10](#) (work in progress),
February 2012.

Authors' Addresses

IJsbrand Wijnands (editor)
Cisco Systems, Inc.
De kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Toerless Eckert
Cisco Systems, Inc.
170 Tasman Drive
San Jose CA, 95134
USA

Email: eckert@cisco.com

Nicolai Leymann
Deutsche Telekom
Winterfeldtstrasse 21
Berlin 10781
Germany

Email: n.leymann@telekom.de

Maria Napierala
AT&T Labs
200 Laurel Avenue
Middletown NJ 07748
USA

Email: mnapierala@att.com

