

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 12, 2009

I. Wijnands (Editor)
T. Eckert
Cisco Systems, Inc.
N. Leymann
Deutsche Telekom
M. Napierala
AT&T Labs
September 8, 2008

**In-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint
Label Switched Paths
draft-wijnands-mpls-mldp-in-band-signaling-00**

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on March 12, 2009.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

When an IP multicast tree needs to pass through an MPLS domain, it is advantageous to map the tree to a Point-to-Multipoint or Multipoint-to-Multipoint Label Switched Path. This document specifies a way to

provide a one-one mapping between IP multicast trees and Label Switched Paths. The IP multicast control messages are translated into MPLS control messages when they enter the MPLS domain, and are translated back into IP multicast control messages at the far end of the MPLS domain. The IP multicast control information is coded into the MPLS control information in such a way as to ensure that a single Multipoint Label Switched Path gets set up for each IP multicast tree.

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 IP multicast source trees	5
2.2.	Transiting IP multicast bidirectional trees	5
2.3.	Transiting IP multicast shared Trees	6
3.	LSP opaque encodings	6
3.1.	Transit IPv4 Source TLV	6
3.2.	Transit IPv6 Source TLV	7
3.3.	Transit IPv4 bidir TLV	7
3.4.	Transit IPv6 bidir TLV	8
4.	Security Considerations	8
5.	IANA considerations	8
6.	Acknowledgments	9
7.	Contributing authors	9
8.	References	10
8.1.	Normative References	10
8.2.	Informative References	10
	Authors' Addresses	10
	Intellectual Property and Copyright Statements	12

1. Introduction

The mLDP specification [3] describes mechanisms for creating point-to-multipoint (P2MP) and multipoint-to-multipoint MP2MP LSPs. These LSPs are typically used for transporting enduser multicast packets. However, the mLDP specification [3] does not provide any rules for associating particular enduser multicast packets with any particular LSP. Other drafts, like [7], 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 draft 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 element. When an IP multicast tree (either a source-specific tree or a bidirectional tree) enters the MPLS network, the IP multicast control messages used to set up the tree are translated into mLDP messages. The (S,G) or (*,G) information from the IP multicast control messages 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 messages that are sent outside the MPLS domain. Note that although the IP multicast 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. This type of service works well if the number of LSPs that are created is under control of the MPLS network operator, or if the number of LSPs for a particular service are known to be limited in number.

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](#) [2].

1.2. Terminology

IP multicast tree : An IP multicast distribution tree identified by an source IP address and/or IP multicast destination address, also referred to as (S,G) and (*,G).

mLDP : Multicast LDP.

Transit LSP : An 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 signal multicast route information.

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, acting indifferently 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

Suppose an LSR, call it D, is attached to a network that is capable of MPLS multicast and IP multicast, and D receives a PIM Join from the IP multicast interface. The PIM Join identifies a particular IP multicast tree. Suppose that D can determine that the IP multicast tree needs to travel through the MPLS network until it reaches some other LSR, U. For instance, when D looks up the route to the Source or Rendezvous Point (RP) [4] 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 identified by the PIM Join. 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 out of scope of this document.

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

In order to send the multicast stream via a P2MP or MP2MP LSP using in-band signaling the source and the group will be encoded into an mLDP opaque TLV encoding [3]. The type of encoding depends on the IP version. The tree type (P2MP or MP2MP) depends on whether this is a source specific or a bidirectional multicast stream. The root of the tree is Ingress LSR that was found during the route lookup on the source or RP. Using this information a mLDP FEC is created and the LSP is build towards the root of the LSP.

When an LSR receives a label mapping or withdraw and discovers it is the root of the identified P2MP or MP2MP LSP, then the following procedure will be executed. If the opaque encoding of the FEC indicates this is an Transit LSP (indicated by the opaque type), the opaque TLV will be decoded and the multicast source and group is passed to the multicast code. If the multicast tree information was received via a label mapping, the multicast code will effectively treat this as having received a PIM join from the MPLS network. If it was due to a label withdraw, the multicast code will effectively treat this as having received a PIM prune from the MPLS network. From this point on normal PIM process will occur and multicast packets are forwarded to the LSP or pruned from the LSP.

2.1. Transiting IP multicast source trees

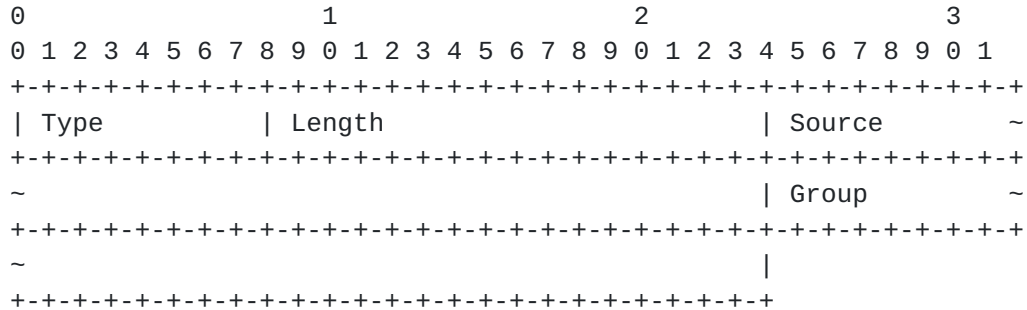
IP multicast source trees can either be created via PIM operating in SSM mode [5] or ASM mode [4] and MUST be transporting across the MPLS network using a P2MP LSP. A Transit LSP may be setup to forward the IP multicast traffic across an MPLS core. If the multicast source is reachable in a global table the source and group addresses are encoded into the a transit TLV. Depending on the IP version it is either [Section 3.1](#) or [Section 3.2](#).

2.2. Transiting IP multicast bidirectional trees

Bidirectional IP multicast trees [6] MUST be transported across a MPLS network using MP2MP LSPs. A bidirectional tree does not have a specific source address; only the group address and subnet mask are relevant for multicast forwarding. The RP for the Multicast group is used to select the ingress PE and root of the LSP. How the RP is discovered for the multicast group is out the scope of this document. The group address is encoded in either [Section 3.3](#) or [Section 3.4](#), depending on the IP version. The subnet mask associated with the

Group: IPv4 multicast group address, 4 octets.

3.2. Transit IPv6 Source TLV




Type: 3 (to be assigned by IANA).

Length: 32

Source: IPv6 multicast source address, 16 octets.

Group: IPv6 multicast group address, 16 octets.

3.3. Transit IPv4 bidir TLV



Type: 4 (to be assigned by IANA).

Length: 5

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

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          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|
|                               Group                               ~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
~
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Type: 4 (to be assigned by IANA).

Length: 17

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

Group: IPv6 multicast group address, 16 octets.

4. Security Considerations

The same security considerations apply as for the base LDP specification, as described in [1].

5. IANA considerations

This document requires allocation from the LDP MP Opaque Value Element type name space managed by IANA. The values requested are:

Transit IPv4 Source TLV type - requested 2

Transit IPv6 Source TLV type - requested 3

Transit IPv4 Bidir TLV type - requested 4

Transit IPv6 Bidir TLV type - requested 5

6. Acknowledgments

Thanks to Eric Rosen for his valuable comments on this draft.

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
Goslarer Ufer 35
Berlin, 10589
Germany
E-mail: nicolai.leymann@t-systems.com

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

- [1] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [3] Minei, I., "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", [draft-ietf-mppls-ldp-p2mp-05](#) (work in progress), June 2008.

8.2. Informative References

- [4] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", [RFC 4601](#), August 2006.
- [5] Holbrook, H. and B. Cain, "Source-Specific Multicast for IP", [RFC 4607](#), August 2006.
- [6] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", [RFC 5015](#), October 2007.
- [7] Aggarwal, R., Bandi, S., Cai, Y., Morin, T., Rekhter, Y., Rosen, E., Wijnands, I., and S. Yasukawa, "Multicast in MPLS/BGP IP VPNs", [draft-ietf-l3vpn-2547bis-mcast-07](#) (work in progress), July 2008.

Authors' Addresses

IJsbrand Wijnands
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
Goslarer Ufer 35
Berlin 10589
Germany

Email: nicolai.leymann@t-systems.com

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

Email: mnapierala@att.com

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

