

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2014

IJ. Wijnands, Ed.
E. Rosen
Cisco
A. Gulko
Thomson Reuters
U. Joorde
Deutsche Telekom
J. Tantsura
Ericsson
October 9, 2013

mLDP In-Band Signaling with Wildcards
draft-wijnands-mpls-mlbp-in-band-wildcard-encoding-01

Abstract

There are scenarios in which an IP multicast tree traverses an MPLS domain. In these scenarios, it can be desirable to convert the IP multicast tree "seamlessly" to an MPLS multipoint label switched path (MP-LSP) when it enters the MPLS domain, and then to convert it back to an IP multicast tree when it exists the MPLS domain. Previous documents specify procedures that allow certain kinds of IP multicast trees (either "Source-Specific Multicast" trees or "Bidirectional multicast" trees) to be attached to an MPLS multipoint label switched path (MP-LSP), either in the context of a particular VPN or in a "global" (non-VPN) context. However, the previous documents do not specify procedures for attaching IP "Any Source Multicast" trees to MP-LSPs, nor do they specify procedures "aggregating" multiple IP multicast trees onto a single MP-LSP. This document specifies the procedures to support these functions. It does so by defining "wildcard" encodings making it possible to specify, when setting up an MP-LSP, that a set of IP multicast trees or a shared IP multicast tree should be attached to that MP-LSP.

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 April 12, 2014.

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.	Introduction	4
2.	Terminology and Definitions	5
3.	Wildcards in mLDP Opaque Value TLVs	7
4.	Some Wildcard Use Cases	8
4.1.	PIM shared tree forwarding	8
4.2.	IGMP/MLD proxying	9
4.3.	Selective Source mapping	10
5.	Procedures for Wildcard Source Usage	10
6.	Procedures for Wildcard Group Usage	11
7.	Determining the MP-LSP Root (Ingress LSR)	11
8.	Anycast RP	12
9.	Acknowledgements	12
10.	IANA Considerations	12
11.	Security Considerations	12
12.	References	12
12.1.	Normative References	12
12.2.	Informative References	13
	Authors' Addresses	13

1. Introduction

[RFC6826] and [[I-D.ietf-l3vpn-mldp-vrf-in-band-signaling](#)] specify procedures for mLDP (Multiple Extensions to the Label Distribution Protocol) that allow an IP multicast tree (either a "Source-Specific Multicast" tree or a "Bidirectional multicast" tree) to be attached "seamlessly" to an MPLS multipoint label switched path (MP-LSP). This can be useful, for example, when there is multicast data that originates in a domain that supports IP multicast, then has to be forwarded across a domain that supports MPLS multicast, then has to be forwarded across another domain that supports IP multicast. By attaching an IP multicast tree to an MP-LSP, data that is traveling along the IP multicast tree is moved to the MP-LSP. The data then travels along the MP-LSP through the MPLS domain. When it reaches the boundary of the MPLS domain, it can be seamlessly passed along an IP multicast tree. This can be useful in either VPN context or global context.

When following the procedures of those documents, a given MP-LSP can carry data from only a single IP source-specific multicast tree (i.e., a single "(S,G) tree"). However, there are scenarios in which it would be desirable to "aggregate" a number of (S,G) trees on a single MP-LSP. Aggregation allows a given number of IP multicast trees to use a smaller number of MP-LSPs, thus saving state in the network.

In addition, the previous documents do not support the attachment of an "Any Source Multicast" (ASM) shared tree to an MP-LSP (except in the case where the ASM shared tree is a "bidirectional" tree (i.e., a tree set up by BIDIR-PIM [[RFC5015](#)]). However, there are scenarios in which it would be desirable to attach a non-bidirectional ASM shared tree to an MP-LSP.

In mLDP, every MP-LSP is identified by the combination of a "root node" (or "Ingress LSR") and an "Opaque Value" that uniquely identifies the MP-LSP in the context of the root node. When mLDP in-band signaling is used (for non-bidirectional trees), the Opaque Value has an IP Source Address (S) and an IP Group Address (G) encoded into it, thus enabling it to identify a particular IP multicast (S,G) tree.

This document specifies a way to encode an mLDP "Opaque Value" that replaces either the "S" or the "G" or both by a "wildcard". Procedures are described for using the wildcard encoding to map non-bidirectional ASM shared trees to MP-LSPs, and for mapping multiple (S,G) trees (with a common value of S or a common value of G) to a single MP-LSP.

Some example scenarios where wildcard encoding is useful are:

- o PIM Shared tree forwarding with threshold infinity.
- o IGMP/MLD proxying.
- o Selective Source mapping.

These scenarios are discussed in [Section 4](#). Note that this list of scenarios is not meant to be exhaustive.

This draft specifies only the mLDP procedures that are specific to the use of wildcards. mLDP in-band signaling procedures that are not specific to the use of wildcards can be found in [[RFC6826](#)] and [[I-D.ietf-l3vpn-mldp-vrf-in-band-signaling](#)]. Unless otherwise specified in this document, those procedures still apply when wildcards are used.

2. Terminology and Definitions

Readers of this document are assumed to be familiar with the terminology and concepts of the documents listed as Normative References. For convenience, some of the more frequently used terms appear below.

IGMP:

Internet Group Management Protocol.

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.

Ingress LSR:

Root node of a MP-LSP. When in-band signaling is used, the Ingress LSR receives mLDP messages about a particular MP-LSP from "downstream", and emits IP multicast control messages "upstream". The set of IP multicast control messages that are emitted upstream depends upon the contents of the LDP Opaque Value TLVs. The Ingress LSR also receives IP multicast data messages from "upstream" and sends them "downstream" as MPLS packets on a MP-LSP.

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).

MLD:

Multicast Listener Discovery.

mLDP:

Multipoint LDP.

MP-LSP:

A P2MP or MP2MP LSP.

PIM:

Protocol Independent Multicast.

PIM-ASM:

PIM Any Source Multicast.

PIM-SM:

PIM Sparse Mode

PIM-SSM:

PIM Source Specific Multicast.

RP:

The PIM Rendezvous Point.

Egress LSR:

The Egress LSRs of an MP-LSP are LSPs that receive MPLS multicast data packets from "upstream" on that MP-LSP, and that forward that data "downstream" as IP multicast data packets. The Egress LSRs also receive IP multicast control messages from "downstream", and send mLDP control messages "upstream". When in-band signaling is used, the Egress LSRs construct Opaque Value TLVs that contain IP source and/or group addresses, based on the contents of the IP multicast control messages received from downstream.

Threshold Infinity:

A PIM-SM procedure where no source specific multicast (S,G) trees are created for multicast packets that are forwarded down the shared tree (*,G).

TLV:

A protocol element consisting of a type field, followed by a length field, followed by a value field. Note that the value field of a TLV may be sub-divided into a number of sub-fields.

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](#)].

3. Wildcards in mLDP Opaque Value TLVs

[RFC6826] and [[I-D.ietf-l3vpn-mlbp-vrf-in-band-signaling](#)] define the following TLVs: Transit IPv4 Source TLV, Transit IPv6 Source TLV, Transit VPNv4 Source TLV, and Transit VPNv6 Source TLV. The value fields of each such TLV is divided into a number of sub-fields, one of which contains an IP source address, and one of which contains an IP group address. Per those documents, these fields must contain valid IP addresses.

This document extends the definition of those TLVs by allowing either the IP Source Address field or the IP Group Address field (or both) to specify a "wildcard" rather than a valid IP address.

A value of all zeroes in the IP Source Address sub-field is used to represent a wildcard source address. A value of all zeroes in the IP Group Address sub-field is used to represent the wildcard group address. Note that the lengths of these sub-fields is as specified in the previous documents.

If the IP Source Address sub-field contains the wildcard, and the IP Group Address sub-field contains an IP multicast group address, say G, that is NOT in the SSM address range (see [Section 4.8 of \[RFC4601\]](#)), the TLV identifies a PIM-SM shared tree.

If the IP Source Address sub-field contains the wildcard, and the IP Group Address sub-field contains an IP multicast group address, say G, that is in the SSM address range, the TLV identifies the collection of PIM-SSM trees with the given group address.

If the IP Source Address sub-field contains a non-zero IP address, and the IP Group Address sub-field contains the wildcard, the TLV identifies the collection of PIM-SSM trees that have the source address as their root.

Procedures for the use of the wildcards are discussed in Sections 4, 5 and 6. Please note that, as always, the structure of the Opaque Value TLVs does not actually affect the operation of mLDP, but only affects the interface between mLDP and IP multicast.

At the present time, there are no procedures defined for the use of a wildcard group in the following TLVs that are defined in [[RFC6826](#)] or [[I-D.ietf-l3vpn-mlbp-vrf-in-band-signaling](#)]: Transit IPv4 Bidir TLV, Transit IPv6 Bidir TLV, Transit VPNv4 Bidir TLV, Transit VPNv6 Bidir TLV. Such procedures may be added in a later revision of this document. Note that the Bidir TLVs do not have a "Source Address" sub-field, and hence the notion of a wildcard source is not applicable to them.

At the present time, there are no procedures defined for the use of both a wildcard source and a wildcard group in the same TLV. Such procedures may be added in a later revision of this document.

4. Some Wildcard Use Cases

This section discusses a number of wildcard use cases. The set of use cases here is not meant to be exhaustive. In each of these use cases, the Egress LSRs construct mLDP Opaque Value TLVs that contain wildcards in the IP Source Address or IP Group Address sub-fields.

4.1. PIM shared tree forwarding

PIM [[RFC4601](#)] has the concept of a "shared tree", identified as (*,G). This concept is only applicable when G is an IP Multicast Group address that is not in the SSM address range (i.e., is an ASM group address). Every ASM group is associated with a Rendezvous Point (RP), and the (*,G) tree is built towards the RP (i.e., its root is the RP). The RP for group G is responsible for forwarding packets down the (*,G) tree. The packets forwarded down the (*,G) tree may be from any multicast source, as long as they have an IP destination address of G.

The RP learns about all the multicast sources for a given group, and then joins a source-specific tree for each such source. I.e., when the RP for G learns that S has multicast data to send to G, the RP joins the (S,G) tree. When the RP receives multicast data from S that is destined to G, the RP forwards the data down the (*,G) tree. There are several different ways that the RP may learn about the sources for a given group. The RP may learn of sources via PIM Register messages [[RFC4601](#)], via MSDP [[RFC3618](#)] or by observing packets from a source that is directly connected to the RP.

In PIM, a PIM router that has receivers for a particular ASM multicast group G (known as a "last hop" router for G) will first join the (*,G) tree. As it receives multicast traffic on the (*,G) tree, it learns (by examining the IP headers of the multicast data packets) the sources that are transmitting to G. Typically, when a last hop router for group G learns that source S is transmitting to G, the last hop router joins the (S,G) tree, and "prunes" S off the (*,G) tree. This allows each last hop router to receive the multicast data along the shortest path from the source to the last hop router. (Full details of this behavior can be found in [[RFC4601](#)].)

In some cases, however, a last hop router for group G may decide not to join the source trees, but rather to keep receiving all the

traffic for G from the (*,G) tree. In this case, we say that the last hop router has "threshold infinity" for group G. This is optional behaviour documented in [\[RFC4601\]](#). "Threshold infinity" is often used in deployments where the RP is between the multicast sources and the multicast receivers for group G, i.e., in deployments where it is known that the shortest path from any source to any receiver of the group goes through the RP. In these deployments, there is no advantage for a last hop router to join a source tree, since the data is already traveling along the shortest path. The only effect of executing the complicated procedures for joining a source tree and pruning the source off the shared tree would be to increase the amount of multicast routing state that has to be maintained in the network.

To efficiently use mLDP in-band signaling in this scenario, it is necessary for the Egress LSRs to construct an Opaque Value TLV that identifies a (*,G) tree. This is done by using the wildcard in the IP Source Address sub-field, and setting the IP Group Address sub-field to G.

Note that these mLDP in-band signaling procedures do not support PIM-ASM in scenarios where "threshold infinity" is not used.

4.2. IGMP/MLD proxying

There are scenarios where the multicast senders and receivers are directly connected to an MPLS routing domain, and where it is desired to use mLDP rather than PIM to set up "trees" through that domain.

In these scenarios we can apply "IGMP/MLD proxying" and eliminate the use of PIM. The senders and receivers consider the MPLS domain to be single hop between each other. [\[RFC4605\]](#) documents procedures where a multicast routing protocol is not necessary to build a 'simple tree'. Within the MPLS domain, mLDP will be used to build a MP-LSP, but this is hidden from the senders and receivers. The procedures defined in [\[RFC4605\]](#) are applicable, since the senders and receivers are considered to be one hop away from each other.

For mLDP to build the necessary MP-LSP, it needs to know the root of the tree. Following the procedures as defined in [\[RFC4605\]](#) we depend on manual configuration of the mLDP root for the ASM multicast group. Since the MP-LSP for a given ASM multicast group will carry traffic from all the sources for that group, the Opaque Value TLV used to construct the MP-LSP will contain a wildcard in the IP Source Address sub-field.

4.3. Selective Source mapping

In many IPTV deployments, the content servers are gathered into a small number of sites. Popular channels are often statically configured, and forwarded over a core MPLS network to the egress routers. Since these channels are statically defined, they MAY also be forwarded over a multipoint LSP with wildcard encoding. The sort of wildcard encoding that needs to be used (Source and/or Group) depends on the Source/Group allocation policy of the IPTV provider. Other options are to use MSDP [[RFC3618](#)] or BGP "Auto-Discovery" procedures [[RFC6513](#)] for source discovery by the Ingress LSR. Based on the received wildcard, the Ingress LSR can select from the set of IP multicast streams for which it has state.

5. Procedures for Wildcard Source Usage

The IP multicast component on an Egress LSR determines when a wildcard is to be used in the IP Source Address sub-field of an mLDP Opaque Value TLV. How the IP multicast component determines this is a local matter, and may need to be explicitly configured. It MAY however use the following rules (with or without explicit configuration);

1. Suppose that PIM is enabled, and an Egress LSR needs to join a non-bidirectional ASM group G, and the RP for G is reachable via a BGP route. The Egress LSR MAY choose the BGP Next Hop of the route to the RP to be the Ingress LSR (root node) of the MP-LSP corresponding to the (*,G) tree. (See also [Section 7](#).) The Egress LSR MAY identify the (*,G) tree by using an mLDP Opaque Value TLV whose IP Source Address sub-field contains a wildcard, and whose IP Group Address sub-field contains G.
2. If PIM is not enabled for group G, and an IGMP/MLD group membership report for G has been received, the Egress LSR may determine the "proxy device" for G (following the procedures defined in [[RFC4605](#)]). It can then set up an MP-LSP using the proxy device as the Ingress LSR. The Egress LSR then needs to signal the Ingress LSR that the MP-LSP is to carry traffic belonging to group G. It does this by using an Opaque Value TLV whose IP Source Address sub-field contains a wildcard, and whose IP Group Address sub-field contains G.

When the Ingress LSR receives an mLDP Opaque Value TLV that has been defined for in-band signaling, the information from the sub-fields of that TLV is passed to the IP multicast component of the Ingress LSR. If the IP Source Address sub-field contains a wildcard, the IP multicast component must determine how to process it. How the

wildcard is processed is a local matter, subject to the rules below:

1. If PIM is enabled and the group identified in the Opaque Value TLV is a non-bidirectional ASM group, the Ingress LSR acts as if it had received a (*,G) IGMP/MLD report from a downstream node, and the procedures defined in [[RFC4601](#)] are followed.
2. If PIM is enabled and the identified group is a PIM-SSM group, all multicast sources known for the group on the Ingress LSR are to be forwarded down the MP-LSP.
3. If PIM is not enabled for identified group, the Ingress LSR acts as if it had received a (*,G) IGMP/MLD report from a downstream node, and the procedures as defined in [[RFC4605](#)] are followed.

6. Procedures for Wildcard Group Usage

The IP multicast component on an Egress LSR determines when a wildcard is to be used in the IP Group Address sub-field of an mLDP Opaque Value TLV. How the IP multicast component determines this is a local matter, and may need to be explicitly configured.

When the Ingress LSR (i.e., the root node of the MP-LSP) receives an mLDP Opaque Value TLV that has been defined for in-band signaling, the information from the sub-fields of that TLV is passed to the IP multicast component of the Ingress LSR. If the IP Group Address sub-field contains a wildcard, the Ingress LSR examines its IP multicast routing table, to find all the IP multicast streams whose IP source address is the address specified in the IP Source Address sub-field of the TLV. All these streams SHOULD be forwarded down the MP-LSP identified by the Opaque Value TLV. Note that some of these streams may have SSM group addresses, while some may have ASM group addresses.

7. Determining the MP-LSP Root (Ingress LSR)

Documents [[RFC6826](#)] and [[I-D.ietf-l3vpn-mldp-vrf-in-band-signaling](#)] describe procedures by which an Egress LSR may determine the MP-LSP root node address corresponding to a given IP multicast stream, based upon the IP address of the source of the IP multicast stream. When a wildcard source encoding is used, PIM is enabled, and the group is a non-bidirectional ASM group, a similar procedure is applied. The only difference from the above mentioned procedures is that the Proxy device or RP address is used instead of the Source to discover the mLDP root node address.

In all other cases some sort of manual configuration is applied in order to find the root node. Note, finding the root node is a local implementation matter and not limited to the solutions mentioned in this document.

8. Anycast RP

In the scenarios where in-band signaling is used, it is unlikely that the RP-to-Group mappings are being dynamically distributed over the MPLS core. It is more likely that the RP address is statically configured at each multicast site. In these scenarios, it is advisable to configure an Anycast RP Address at each site, in order to provide redundancy. See [[RFC3446](#)] for more details.

9. Acknowledgements

We would like to thank Loa Andersson for his review and comments.

10. IANA Considerations

There are no new allocations required from IANA.

11. Security Considerations

There are no security considerations other than ones already mentioned in [[RFC6826](#)] and [[I-D.ietf-l3vpn-mldp-vrf-in-band-signaling](#)].

12. References

12.1. Normative References

- [I-D.ietf-l3vpn-mldp-vrf-in-band-signaling]
Wijnands, I., Hitchen, P., Leymann, N., Henderickx, W.,
and a. arkadiy.gulko@thomsonreuters.com, "Multipoint Label
Distribution Protocol In-Band Signaling in a VRF Context",
[draft-ietf-l3vpn-mldp-vrf-in-band-signaling-01](#) (work in
progress), June 2013.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas,

"Protocol Independent Multicast - Sparse Mode (PIM-SM):
Protocol Specification (Revised)", [RFC 4601](#), August 2006.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick,
"Internet Group Management Protocol (IGMP) / Multicast
Listener Discovery (MLD)-Based Multicast Forwarding
("IGMP/MLD Proxying")", [RFC 4605](#), August 2006.

[RFC6826] Wijnands, IJ., Eckert, T., Leymann, N., and M. Napierala,
"Multipoint LDP In-Band Signaling for Point-to-Multipoint
and Multipoint-to-Multipoint Label Switched Paths",
[RFC 6826](#), January 2013.

12.2. Informative References

[RFC3446] Kim, D., Meyer, D., Kilmer, H., and D. Farinacci, "Anycast
Rendezvous Point (RP) mechanism using Protocol Independent
Multicast (PIM) and Multicast Source Discovery Protocol
(MSDP)", [RFC 3446](#), January 2003.

[RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery
Protocol (MSDP)", [RFC 3618](#), October 2003.

[RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano,
"Bidirectional Protocol Independent Multicast (BIDIR-
PIM)", [RFC 5015](#), October 2007.

[RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP
VPNs", [RFC 6513](#), February 2012.

Authors' Addresses

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

Phone:
Email: ice@cisco.com

Eric Rosen
Cisco
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone:
Email: erosen@cisco.com

Arkadiy Gulko
Thomson Reuters
195 Broadway
New York NY 10007
USA

Email: arkadiy.gulko@thomsonreuters.com

Uwe Joorde
Deutsche Telekom
Hammer Str. 216-226
Muenster D-48153
DE

Email: Uwe.Joorde@telekom.de

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, california 95134
usa

Email: jeff.tantsura@ericsson.com

