

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 21, 2014

IJ. Wijnands, Ed.
Cisco Systems
P. Hitchen
BT
N. Leymann
Deutsche Telekom
W. Henderickx
Alcatel-Lucent
A. Gulko
Thomson Reuters
J. Tantsura
Ericsson
January 17, 2014

**Multipoint Label Distribution Protocol
In-Band Signaling in a VRF Context
draft-ietf-l3vpn-mldp-vrf-in-band-signaling-03**

Abstract

An IP Multicast Distribution Tree (MDT) may traverse both label switching (i.e. - Multi-Protocol Label Switching, or MPLS) and non-label switching regions of a network. Typically the MDT begins and ends in non-MPLS regions, but travels through an MPLS region. In such cases, it can be useful to begin building the MDT as a pure IP MDT, then convert it to an MPLS Multipoint Label Switched Path (MP-LSP) when it enters an MPLS-enabled region, and then convert it back to a pure IP MDT when it enters a non-MPLS-enabled region. Other documents specify the procedures for building such a hybrid MDT, using Protocol Independent Multicast (PIM) in the non-MPLS region of the network, and using Multipoint Extensions to Label Distribution Protocol (mLDP) in the MPLS region. This document extends those procedures to handle the case where the link connecting the two regions is a "Virtual Routing and Forwarding Table" (VRF) link, as defined in the "BGP IP/MPLS VPN" specifications. However, this document is primarily aimed at particular use cases where VRFs are used to support multicast applications other than Multicast VPN.

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 July 21, 2014.

Copyright Notice

Copyright (c) 2014 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
1.1.	Conventions used in this document	5
1.2.	Terminology	5
2.	VRF In-band signaling for MP LSPs	6
3.	Encoding the Opaque Value of an LDP MP FEC	7
3.1.	Transit VPNv4 Source TLV	7
3.2.	Transit VPNv6 Source TLV	8
3.3.	Transit VPNv4 bidir TLV	9
3.4.	Transit VPNv6 bidir TLV	10
4.	Security Considerations	11
5.	IANA considerations	11
6.	Acknowledgments	11
7.	References	11
7.1.	Normative References	11
7.2.	Informative References	12
	Authors' Addresses	12

1. Introduction

Sometimes an IP multicast distribution tree (MDT) traverses both MPLS-enabled and non-MPLS-enabled regions of a network. Typically the MDT begins and ends in non-MPLS regions, but travels through an MPLS region. In such cases, it can be useful to begin building the MDT as a pure IP MDT, then convert it to an MPLS Multipoint LSP (Label Switched Path) when it enters an MPLS-enabled region, and then convert it back to a pure IP MDT when it enters a non-MPLS-enabled region. Other documents specify the procedures for building such a hybrid MDT, using Protocol Independent Multicast (PIM) in the non-MPLS region of the network, and using Multipoint Extensions to Label Distribution Protocol (mLDP) in the MPLS region. This document extends the procedures from [\[RFC6826\]](#) to handle the case where the link connecting the two regions is a "Virtual Routing and Forwarding Table" (VRF) link, as defined in the "BGP IP/MPLS VPN" specifications [\[RFC6513\]](#). However, this document is primarily aimed at particular use cases where VRFs are used to support multicast applications other than Multicast VPN.

In PIM, a tree is identified by a source address (or in the case of bidirectional trees [\[RFC5015\]](#), a rendezvous point address or "RPA") and a group address. The tree is built from the leaves up, by sending PIM control messages in the direction of the source address or the RPA. In mLDP, a tree is identified by a root address and an "opaque value", and is built by sending mLDP control messages in the direction of the root. The procedures of [\[RFC6826\]](#) explain how to convert a PIM <source address or RPA, group address> pair into an mLDP <root node, opaque value> pair and how to convert the mLDP <root node, opaque value> pair back into the original PIM <source address or RPA, group address> pair.

However, the procedures in [\[RFC6826\]](#) assume that the routers doing the PIM/mLDP conversion have routes to the source address or RPA in their global routing tables. Thus the procedures cannot be applied exactly as specified when the interfaces connecting the non-MPLS-enabled region to the MPLS-enabled region are interfaces that belong to a VRF as described in [\[RFC4364\]](#). This specification extends the procedures of [\[RFC6826\]](#) so that they may be applied in the VRF context.

As in [\[RFC6826\]](#), the scope of this document is limited to the case where the multicast content is distributed in the non-MPLS-enabled regions via PIM-created Source-Specific or Bidirectional trees. Bidirectional trees are always mapped onto Multipoint-to-Multipoint LSPs, and source-specific trees are always mapped onto Point-to-Multipoint LSPs.

Note that the procedures described herein do not support non-bidirectional PIM ASM groups, do not support the use of multicast trees other than mLDP multipoint LSPs in the core, and do not provide the capability to aggregate multiple PIM trees onto a single multipoint LSP. If any of those features are needed, they can be provided by the procedures of [\[RFC6513\]](#) and [\[RFC6514\]](#). However, there are cases where multicast services are offered through interfaces associated with a VRF, and where mLDP is used in the core, but where aggregation is not desired. For example, some service providers offer multicast content to their customers, but have chosen to use VRFs to isolate the various customers and services. This is a simpler scenario than one in which the customers provide their own multicast content, out of the control of the service provider, and can be handled with a simpler solution. Also, when PIM trees are mapped one-to-one to multipoint LSPs, it is helpful for troubleshooting purposes to have the PIM source/group addresses encoded into the mLDP FEC element.

In order to use the mLDP in-band signaling procedures for a particular group address in the context of a particular set of VRFs, those VRFs MUST be configured with a range of multicast group addresses for which mLDP in-band signaling is to be enabled. This configuration is per VRF ("Virtual Routing and Forwarding table", defined in [\[RFC4364\]](#)). For those groups, and those groups only, the procedures of this document are used. For other groups the general purpose Multicast VPN procedures MAY be used, although it is more likely this VRF is dedicated to mLDP in-band signaling procedures and all other groups are discarded. The configuration MUST be present in all PE routers that attach to sites containing senders or receivers for the given set of group addresses. Note, since the provider most likely owns the multicast content and how it is transported across the network is transparent to the end-user, no co-ordination needs to happen between the end-user and the provider.

[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 source IP address and/or IP multicast destination address, also referred to as (S,G) and (*,G).

mLDP : Multicast LDP.

In-band signaling: Using the opaque value of a mLDP FEC element to encode 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 (see [[RFC6388](#)]).

MP2MP LSP: An LSP that connects a set of leaf nodes, acting indifferently as ingress or egress (see [[RFC6388](#)]).

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

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

VRF: Virtual Routing and Forwarding table.

2. VRF In-band signaling for MP LSPs

Suppose that a PE router, PE1, receives a PIM Join(S,G) message over one of its interfaces that is associated with a VRF. Following the procedure of [section 5.1 of \[RFC6513\]](#), PE1 determines the "upstream RD", the "upstream PE", and the "upstream multicast hop" (UMH) for the source address S.

In order to transport the multicast tree via an MP LSP using VRF in-band signaling, an mLDP Label Mapping Message is sent by PE1. This message will contain either a P2MP FEC or an MP2MP FEC (see [[RFC6388](#)]), depending upon whether the PIM tree being transported is a source-specific tree, or a bidirectional tree, respectively. The FEC contains a "root" and an "opaque value".

If the UMH and the upstream PE have the same IP address (i.e., the Upstream PE is the UMH), then the root of the Multipoint FEC is set to the IP address of the Upstream PE. If, in the context of this VPN, (S,G) refers to a source-specific MDT, then the values of S, G, and the upstream RD are encoded into the opaque value. If, in the context of this VPN, G is a bidirectional group address, then S is replaced with the value of the RPA associated with G.

The encoding details are specified in [Section 3](#). Conceptually, the Multipoint FEC can be thought of as an ordered pair: {root=<Upstream-PE>; opaque_value=<S or RPA , G, Upstream-RD>}. The mLDP Label Mapping Message is then sent by PE1 on its LDP session to the "next hop" on the message's path to the upstream PE. The "next hop" is usually the directly connected next hop, but see [\[RFC7060\]](#) for cases in which the next hop is not directly connected.

If the UMH and the upstream PE do not have the same IP address, the procedures of [section 2 of \[RFC6512\]](#) should be applied. The root node of the multipoint FEC is set to the UMH. The recursive opaque value is then set as follows: the root node is set to the upstream PE, and the opaque value is set to the multipoint FEC described in the previous paragraph. That is, the multipoint FEC can be thought of as the following recursive ordered pair: {root=<UMH>; opaque_value=<root=Upstream-PE, opaque_value=<S or RPA, G, Upstream-RD>>}}.

The encoding of the multipoint FEC also specifies the "type" of PIM MDT being spliced onto the multipoint LSP. Four types of MDT are defined in [\[RFC6826\]](#): IPv4 source-specific tree, IPv6 source-specific tree, IPv4 bidirectional tree, and IPv6 bidirectional tree.

When a PE router receives an mLDP message with a P2MP or MP2MP FEC, where the PE router itself is the root node, and the opaque value is one of the types defined in [Section 3](#), then it uses the RD encoded in the opaque value field to determine the VRF context. (This RD will be associated with one of the PE's VRFs.) Then, in the context of that VRF, the PE follows the procedure specified in [section 2 of \[RFC6826\]](#).

3. Encoding the Opaque Value of an LDP MP FEC

This section documents the different transit opaque encodings.

3.1. Transit VPNv4 Source TLV

This opaque value type is used when transporting a source-specific mode multicast tree whose source and group addresses are IPv4 addresses.


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

Group: IPv4 multicast group address, 4 octets.

RD: Route Distinguisher, 8 octets.

3.4. Transit VPNv6 bildir TLV

This opaque value type is used when transporting a bidirectional multicast tree whose group address is an IPv6 address. The RP address is also an IPv6 address in this case.



Type: (to be assigned by IANA).

Length: 41

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

RP: Rendezvous Point (RP) IPv6 address used for encoded group, 16 octets.

Group: IPv6 multicast group address, 16 octets.

RD: Route Distinguisher, 8 octets.

4. Security Considerations

The same security considerations apply as for the base LDP specification, described in [[RFC5036](#)], and the base mLDP specification, described in [[RFC6388](#)]

5. IANA considerations

[RFC6388] defines a registry for the "LDP MP Opaque Value Element Basic Type". This document requires the assignment of four new code points in this registry:

Transit VPNv4 Source TLV type - requested 250

Transit VPNv6 Source TLV type - requested 251

Transit VPNv4 Bidir TLV type - requested TBD-1

Transit VPNv6 Bidir TLV type - requested TBD-2

6. Acknowledgments

Thanks to Eric Rosen, Andy Green, Yakov Rekhter and Eric Gray for their comments on the draft.

7. References

7.1. Normative References

- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", [RFC 4364](#), February 2006.
- [RFC5015] Handley, M., Kouvelas, I., Speakman, T., and L. Vicisano, "Bidirectional Protocol Independent Multicast (BIDIR-PIM)", [RFC 5015](#), October 2007.
- [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.
- [RFC6512] Wijnands, IJ., Rosen, E., Napierala, M., and N. Leymann, "Using Multipoint LDP When the Backbone Has No Route to the Root", [RFC 6512](#), February 2012.
- [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.

7.2. Informative References

- [RFC7060] Napierala, M., Rosen, E., and IJ. Wijnands, "Using LDP Multipoint Extensions on Targeted LDP Sessions", [RFC 7060](#), November 2013.
- [RFC6513] Rosen, E. and R. Aggarwal, "Multicast in MPLS/BGP IP VPNs", [RFC 6513](#), February 2012.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), February 2012.

Authors' Addresses

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

Email: ice@cisco.com

Paul Hitchen
BT
BT Adastral Park
Ipswich IP53RE
UK

Email: paul.hitchen@bt.com

Nicolai Leymann
Deutsche Telekom
Winterfeldtstrasse 21
Berlin 10781
Germany

Email: n.leymann@telekom.de

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
Antwerp 2018
Belgium

Email: wim.henderickx@alcatel-lucent.com

Arkadiy Gulko
Thomson Reuters
195 Broadway
New York NY 10007
USA

Email: arkadiy.gulko@thomsonreuters.com

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

Email: jeff.tantsura@ericsson.com

