

Internet Engineering Task Force (IETF)
Request for Comments: 7060
Category: Standards Track
ISSN: 2070-1721

M. Napierala
AT&T
E. Rosen
IJ. Wijnands
Cisco Systems, Inc.
November 2013

Using LDP Multipoint Extensions on Targeted LDP Sessions

Abstract

Label Distribution Protocol (LDP) can be used to set up Point-to-Multipoint (P2MP) and Multipoint-to-Multipoint (MP2MP) Label Switched Paths. However, the specification for the Multipoint Extensions to LDP presupposes that the two endpoints of an LDP session are directly connected. The LDP base specification allows for the case where the two endpoints of an LDP session are not directly connected; such a session is known as a "Targeted LDP" session. This document provides the specification for using the LDP Multipoint Extensions over a Targeted LDP session.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7060>.

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	2
2. Targeted mLDP and the Upstream LSR	3
2.1. Selecting the Upstream LSR	3
2.2. Sending Data from U to D	4
3. Applicability of Targeted mLDP	4
4. LDP Capabilities	5
5. Targeted mLDP with Unicast Replication	5
6. Targeted mLDP with Multicast Tunneling	6
7. Security Considerations	8
8. Acknowledgments	8
9. Normative References	8

1. Introduction

Label Distribution Protocol (LDP) extensions for setting up Point-to-Multipoint (P2MP) Label Switched Paths (LSPs) and Multipoint-to-Multipoint (MP2MP) LSPs are specified in [mLDP]. This set of extensions is generally known as "Multipoint LDP" (mLDP).

A pair of Label Switched Routers (LSRs) that are the endpoints of an LDP session are considered to be "LDP peers". When a pair of LDP peers are "directly connected" (e.g., they are connected by a layer 2 medium or are otherwise considered to be neighbors by the network's interior routing protocol), the LDP session is said to be a "directly connected" LDP session. When the pair of LDP peers are not directly connected, the session between them is said to be a "Targeted" LDP session.

The base specification for mLDP does not explicitly cover the case where the LDP multipoint extensions are used over a Targeted LDP session. This document provides that specification.

We will use the term "Multipoint" to mean "either P2MP or MP2MP".

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].

2. Targeted mLDP and the Upstream LSR

2.1. Selecting the Upstream LSR

In mLDP, a multipoint LSP (MP-LSP) has a unique identifier that is an ordered pair of the form <root, opaque value>. The first element of the ordered pair is the IP address of the MP-LSP's "root node". The second element of the ordered pair is an identifier that is unique in the context of the root node.

If LSR D is setting up the MP-LSP <R, X>, D must determine the "upstream LSR" for <R, X>. In [mLDP], the upstream LSR for <R, X>, U, is defined to be the "next hop" on D's path to R, and "next hop" is tacitly assumed to mean "IGP next hop". It is thus assumed that there is a direct LDP session between D and U. In this specification, we extend the notion of "upstream LSR" to cover the following cases:

- U is the "BGP next hop" on D's path to R, where U and D are not necessarily IGP neighbors, and where there is a Targeted LDP session between U and D. In this case, we allow D to select U as the "upstream LSR" for <R,X>.
- If the "next-hop interface" on D's path to R is an RSVP Traffic Engineering (RSVP-TE) P2P tunnel whose remote endpoint is U, and if there is known to be an RSVP-TE P2P tunnel from U to D, and if there is a Targeted LDP session between U and D, then we allow D to select U as the "upstream LSR" for <R,X>. This is useful when D and U are part of a network area that is fully meshed via RSVP-TE P2P tunnels.

The particular method used to select an "upstream LSR" is determined by the Service Provider (SP) and must be made known a priori (i.e., by provisioning) to all the LSRs involved.

Other methods than the two specified above MAY be used; however, the specification of other methods is outside the scope of this document.

2.2. Sending Data from U to D

By using Targeted mLDP, we can construct an MP-LSP $\langle R, X \rangle$ containing an LSR U, where U has one or more downstream LSR neighbors (D1, ..., Dn) to which it is not directly connected. In order for a data packet to travel along this MP-LSP, U must have some way of transmitting the packet to D1, ..., Dn. We will cover two methods of transmission:

- Unicast Replication

In this method, U creates n copies of the packet and unicasts each copy to exactly one of D1, ..., Dn.

- Multicast Tunneling

In this method, U becomes the root node of a multicast tunnel, with D1, ..., Dn as leaf nodes. When a packet traveling along the MP-LSP $\langle R, X \rangle$ arrives at U, U transmits it through the multicast tunnel, and as a result it arrives at D1, ..., Dn.

When this method is used, it may be desirable to carry traffic of multiple MP-LSPs through a single multicast tunnel. We specify procedures that allow for the proper demultiplexing of the MP-LSPs at the leaf nodes of the multicast tunnel. We do not assume that all the leaf nodes of the tunnel are on all the MP-LSPs traveling through the tunnel; thus, some of the tunnel leaf nodes may need to discard some of the packets received through the tunnel. For example, suppose MP-LSP $\langle R1, X1 \rangle$ contains node U with downstream LSRs D1 and D2, while MP-LSP $\langle R2, X2 \rangle$ contains node U with downstream LSRs D2 and D3. Suppose also that there is a multicast tunnel with U as root and with D1, D2, and D3 as leaf nodes. U can aggregate both MP-LSPs in this one tunnel. However, D1 will have to discard packets that are traveling on $\langle R2, X1 \rangle$, while D3 will have to discard packets that are traveling on $\langle R1, X2 \rangle$.

3. Applicability of Targeted mLDP

When LSR D is setting up MP-LSP $\langle R, X \rangle$, it MUST NOT use Targeted mLDP unless D implements a procedure that can select an LSR U that is a Targeted mLDP peer of D as the "upstream LSR" for $\langle R, X \rangle$. See Section 2.1.

Whether D uses Targeted mLDP when this condition holds is determined by provisioning or by other methods that are outside the scope of this specification.

When Targeted mLDP is used, the choice between unicast replication and multicast tunneling is determined by provisioning or by other methods that are outside the scope of this specification. It is presupposed that all nodes will have a priori knowledge of whether to use unicast replication or to use multicast tunneling. If the latter, it is presupposed that all nodes will have a priori knowledge of the type of multicast tunneling to use.

4. LDP Capabilities

Per [mLDP], any LSR that needs to set up an MP-LSP must support the procedures of [LDP-CAP], and in particular must send and receive the P2MP Capability and/or the MP2MP Capability. This specification does not define any new capabilities; the advertisement of the P2MP and/or MP2MP Capabilities on a Targeted LDP session means that the advertising LSR is capable of following the procedures set forth in this document.

Some of the procedures described in this document require the use of upstream-assigned labels [LDP-UP]. In order to use upstream-assigned labels as part of Targeted mLDP, an LSR must advertise the LDP Upstream-Assigned Label Capability [LDP-UP] on the Targeted LDP session.

5. Targeted mLDP with Unicast Replication

When unicast replication is used, the mLDP procedures are exactly the same as described in [mLDP], with the following exception. If LSR D is setting up MP-LSP <R,X>, its "upstream LSR" is selected according to the procedures of Section 2.1, and is not necessarily the "IGP next hop" on D's path to R.

Suppose that LSRs D1 and D2 are both setting up the P2MP MP-LSP <R,X>, and that LSR U is the upstream LSR on each of their paths to R. D1 and D2 each binds a label to <R,X> and each uses a Label Mapping message to inform U of the label binding. Suppose D1 has assigned label L1 to <R,X> and D2 has assigned label L2 to <R,X>. (Note that L1 and L2 could have the same value or different values; D1 and D2 do not coordinate their label assignments.) When U has a packet to transmit on the MP-LSP <R,X>, it makes a copy of the packet, pushes on label L1, and unicasts the resulting packet to D1. It also makes a second copy of the packet, pushes on label L2, and then unicasts the resulting packet to D2.

This procedure also works when the MP-LSP <R,X> is an MP2MP LSP. Suppose that in addition to labels L1 and L2 described above, U has assigned label L3 for <R,X> traffic received from D1 and label L4 for <R,X> traffic received from D2. When U processes a packet with label

L3 at the top of its label stack, it knows the packet is from D1, so U sends a unicast copy of the packet to D2, after swapping L3 for L2. U does not send a copy back to D1.

Note that all labels used in this procedure are downstream-assigned labels.

The method of unicast is a local matter, outside the scope of this specification. The only requirement is that D1 will receive the copy of the packet carrying label L1 and that D1 will process the packet by looking up label L1. (And similarly, D2 must receive the copy of the packet carrying label L2 and must process the packet by looking up label L2.)

Note that if the method of unicast is MPLS, U will need to push another label on each copy of the packet before transmitting it. This label needs to ensure that delivery of the packet to the appropriate LSR, D1 or D2. Use of penultimate-hop popping for that label is perfectly legitimate.

6. Targeted mLDP with Multicast Tunneling

Suppose that LSRs D1 and D2 are both setting up MP-LSP <R,X> and that LSR U is the upstream LSR on each of their paths to R. Since multicast tunneling is being used, when U has a packet to send on this MP-LSP, it does not necessarily send two copies, one to D1 and one to D2. It may send only one copy of the packet, which will get replicated somewhere downstream in the multicast tunnel. Therefore, the label that gets bound to the MP-LSP must be an upstream-assigned label assigned by U. This requires a change from the procedures of [mLDP]. D1 and D2 do not send Label Mapping messages to U; instead, they send Label Request messages to U, following the procedures of Section 4 of [LDP-UP], asking U to assign a label to the MP-LSP <R,X>. U responds with a Label Mapping message containing an upstream-assigned label L (using the procedures specified in [LDP-UP]). As part of the same Label Mapping message, U also sends an Interface TLV (as specified in [LDP-UP]) identifying the multicast tunnel in which data on the MP-LSP will be carried. When U transmits a packet on this tunnel, it first pushes on the upstream-assigned label L and then pushes on the label that corresponds to the multicast tunnel.

If the numerical value L of the upstream-assigned label is the value 3, defined in [LDP] and [RFC3032] as "Implicit NULL", then the specified multicast tunnel will carry only the specified MP-LSP. That is, aggregation of multiple MP-LSPs into a single multicast

tunnel is not being done. In this case, no upstream-assigned label is pushed onto a packet that is transmitted through the multicast tunnel.

Various types of multicast tunnel may be used. The choice of tunnel type is determined by provisioning, or by some other method that is outside the scope of this document. [LDP-UP] specifies encodings allowing U to identify an mLDP MP-LSP, and RSVP-TE P2MP LSP, as well as other types of multicast tunnel.

Procedures for tunneling MP2MP LSPs through P2MP or MP2MP LSPs are outside the scope of this document.

If the multicast tunnel is an mLDP MP-LSP or an RSVP-TE P2MP LSP, when U transmits a packet on the MP-LSP <R,X>, the upstream-assigned label L will be the second label in the label stack. Penultimate-hop popping MUST NOT be done, because the top label provides the context in which the second label is to be interpreted. See [RFC5331].

When LSR U uses these procedures to inform LSR D that a particular MP-LSP is being carried in a particular multicast tunnel, U and D MUST take appropriate steps to ensure that the packets U sends into this tunnel will be received by D. The exact steps to take depend on the tunnel type. As long as U is D's upstream LSR for any MP-LSP that has been assigned to this tunnel, D must remain joined to the tunnel.

Note that U MAY assign the same multicast tunnel for multiple different MP-LSPs. However, U MUST assign a distinct upstream-assigned label to each MP-LSP. This allows the packets traveling through the tunnel to be demultiplexed into the proper MP-LSPs.

If U has an MP-LSP <R1,X1> with downstream LSRs D1 and D2, and an MP-LSP <R2,X2> with downstream LSRs D2 and D3, U may assign both MP-LSPs to the same multicast tunnel. In this case, D3 will receive packets traveling on <R1,X1>. However, the upstream-assigned label carried by those packets will not be recognized by D3, hence D3 will discard those packets. Similarly, D1 will discard the <R2,X2> packets.

This document does not specify any rules for deciding whether to aggregate two or more MP-LSPs into a single multicast tunnel. Such rules are outside the scope of this document.

Except for the procedures explicitly detailed in this document, the procedures of [mLDP] and [LDP-UP] apply unchanged.

7. Security Considerations

This document raises no new security considerations beyond those discussed in [LDP], [LDP-UP], and [RFC5331].

8. Acknowledgments

The authors wish to thank Lizhong Jin and Lizhen Bin for their comments.

9. Normative References

- [LDP] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [LDP-CAP] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, July 2009.
- [mLDP] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, November 2011.
- [LDP-UP] Aggarwal, R. and JL. Le Roux, "MPLS Upstream Label Assignment for LDP", RFC 6389, November 2011.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, August 2008.

Authors' Addresses

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

EMail: mnapierala@att.com

Eric C. Rosen
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA, 01719
USA

EMail: erosen@cisco.com

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

EMail: ice@cisco.com

