

July 8, 2008

MPLS Upstream Label Assignment for LDP

[draft-ietf-mpls-ldp-upstream-03.txt](#)

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>.

Abstract

This document describes procedures for distributing upstream-assigned labels for Label Distribution Protocol (LDP). It also describes how these procedures can be used for avoiding branch LSR traffic replication on a LAN for LDP point-to-multipoint (P2MP)LSPs.

Table of Contents

1	Specification of requirements	2
2	Introduction	2
3	LDP Upstream Label Assignment Capability	3
4	Distributing Upstream-Assigned Labels in LDP	4
4.1	Procedures	4
5	LDP Tunnel Identifier Exchange	5
6	LDP Point-to-Multipoint LSPs on a LAN	6
7	IANA Considerations	8
8	Acknowledgements	8
9	References	9
9.1	Normative References	9
9.2	Informative References	9
10	Author's Address	10
11	Intellectual Property Statement	10
12	Full Copyright Statement	11

[1. Specification of requirements](#)

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. Introduction](#)

This document describes procedures for distributing upstream-assigned labels [[MPLS-UPSTREAM](#)] for Label Distribution Protocol (LDP). These procedures follow the architecture for MPLS Upstream Label Assignment described in [[MPLS-UPSTREAM](#)].

This document describes extensions to LDP that a LSR can use to advertise to its neighboring LSRs whether the LSR supports upstream label assignment.

This document also describes extensions to LDP to distribute upstream-assigned labels.

The usage of MPLS upstream label assignment using LDP for avoiding branch LSR traffic replication on a LAN for LDP P2MP LSPs [[MLDP](#)] is also described.

3. LDP Upstream Label Assignment Capability

According to [[MPLS-UPSTREAM](#)], upstream-assigned label bindings MUST NOT be used unless it is known that a downstream LSR supports them. This implies that there MUST be a mechanism to enable a LSR to advertise to its LDP neighbor LSR(s) its support of upstream-assigned labels.

A new Capability Parameter, the LDP Upstream Label Assignment Capability, is introduced to allow an LDP peer to exchange with its peers, its support of upstream label assignment. This parameter follows the format and procedures for exchanging Capability Parameters defined in [[LDP-CAP](#)].

Following is the format of the LDP Upstream Label Assignment Capability Parameter:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1|0| Upstream Lbl Ass Cap(IANA)|          Length (= 1)          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|1| Reserved          |
+--+--+--+--+--+--+--+

```

If a LSR includes the Upstream Label Assignment Capability in LDP Initialization Messages it implies that the LSR is capable of both distributing upstream-assigned label bindings and receiving upstream-assigned label bindings. The reserved bits MUST be set to zero on transmission and ignored on receipt. The Upstream Label Assignment Capability Parameter can be exchanged only in LDP initialization messages.

4. Distributing Upstream-Assigned Labels in LDP

An optional LDP TLV, Upstream-Assigned Label Request TLV, is introduced. This TLV MUST be carried in a Label Request message if an upstream-assigned label is being requested.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0| Upstream Ass Lbl Req (TBD)|          Length          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reserved                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

An optional LDP TLV, Upstream-Assigned Label TLV is introduced to signal an upstream-assigned label. Upstream-Assigned Label TLVs are carried by the messages used to advertise, release and withdraw upstream assigned label mappings.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0| Upstream Ass Label (TBD) |          Length          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reserved                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Label                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Label

This is a 20-bit label value as specified in [[RFC3032](#)] represented as a 20-bit number in a 4 octet field.

4.1. Procedures

Procedures for Label Mapping, Label Request, Label Abort, Label Withdraw and Label Release follow [[RFC3036](#)] other than the modifications pointed out in this section.

A LDP LSR MUST NOT distribute the Upstream Assigned Label TLV to a neighboring LSR if the neighboring LSR had not previously advertised the Upstream Label Assignment Capability in its LDP Initialization messages. A LDP LSR MUST NOT send the Upstream Assigned Label Request TLV to a neighboring LSR if the neighboring LSR had not previously advertised the Upstream Label Assignment Capability in its

LDP Initialization messages.

As described in [[MPLS-UPSTREAM](#)] the distribution of upstream-assigned labels is similar to either ordered LSP control or independent LSP control of the downstream assigned labels.

When the label distributed in a Label Mapping message is an upstream-assigned label, the Upstream Assigned Label TLV MUST be included in the Label Mapping message. When a LSR receives a Label Mapping message with an Upstream Assigned Label TLV and it does not recognize the TLV, it MUST generate a Notification message with a status code of "Unknown TLV" [[RFC3036](#)]. If it does recognize the TLV but is unable to process the upstream label, it MUST generate a Notification message with a status code of "No Label Resources". If the Label Mapping message was generated in response to a Label Request message, the Label Request message MUST contain an Upstream Assigned Label Request TLV. A LSR that generates an upstream assigned label request to a neighbor LSR, for a given FEC, MUST NOT send a downstream label mapping to the neighbor LSR for that FEC unless it withdraws the upstream-assigned label binding. Similarly if a LSR generates a downstream assigned label request to a neighbor LSR, for a given FEC, it MUST NOT send an upstream label mapping to that LSR for that FEC, unless it aborts the downstream assigned label request.

The Upstream Assigned Label TLV may be optionally included in Label Withdraw and Label Release messages that withdraw/release a particular upstream assigned label binding.

5. LDP Tunnel Identifier Exchange

As described in [[MPLS-UPSTREAM](#)] an upstream LSR Ru MAY transmit a MPLS packet, the top label of which (L) is upstream-assigned, to a downstream LSR Rd, by encapsulating it in an IP or MPLS tunnel. In this case the fact that L is upstream-assigned is determined by Rd by the tunnel on which the packet is received. There must be a mechanism for Ru to inform Rd that a particular tunnel from Ru to Rd will be used by Ru for transmitting MPLS packets with upstream-assigned MPLS labels.

When LDP is used for upstream label assignment, the Interface ID TLV [[RFC3472](#)] is used for signaling the Tunnel Identifier. If Ru uses an IP or MPLS tunnel to transmit MPLS packets with upstream assigned labels to Rd, Ru MUST include the Interface ID TLV in the Label Mapping messages along with the Upstream Assigned Label TLV. Three new Interface ID TLVs are introduced to support RSVP-TE P2MP LSPs, IP Multicast Tunnels and context labels. The TLV value acts as the tunnel identifier.

1. RSVP-TE P2MP LSP TLV. Type = TBD. Value of the TLV is the RSVP-TE P2MP Session Object and optionally the P2MP Sender Template Object [RFC4875]. The TLV value identifies the RSVP-TE P2MP LSP. It allows Ru to tunnel an "inner" LDP P2MP LSP, the label for which is upstream assigned, over an "outer" RSVP-TE P2MP LSP that has leaves <Rd1...Rdn>. The P2MP LSP IF_ID TLV allows Ru to signal to <Rd1...Rdn> the binding of the inner LDP P2MP LSP to the outer RSVP-TE P2MP LSP. The control plane signaling between Ru and <Rd1...Rdn> for the inner P2MP LSP uses targeted LDP signaling messages

2. IP Multicast Tunnel TLV. Type = TBD. In this case the TLV value is a <Source Address, Multicast Group Address> tuple. Source Address is the IP address of the root of the tunnel i.e. Ru, and Multicast Group Address is the Multicast Group Address used by the tunnel.

3. MPLS Context Label TLV. Type = TBD. In this case the TLV value is a <Source Address, MPLS Context Label> tuple. The Source Address belongs to Ru and the MPLS Context Label is an upstream assigned label, assigned by Ru. This allows Ru to tunnel an "inner" LDP P2MP LSP, the label of which is upstream assigned, over an "outer" one-hop MPLS LSP, where the outer one-hop LSP has the following property:

- + The label pushed by Ru for the outer MPLS LSP is an upstream assigned context label, assigned by Ru. When <Rd1...Rdn> perform a MPLS label lookup on this label a combination of this label and the incoming interface MUST be sufficient for <Rd1...Rdn> to uniquely determine Ru's context specific label space to lookup the next label on the stack in. <Rd1...Rdn> MUST receive the data sent by Ru with the context specific label assigned by Ru being the top label on the label stack.

Currently the usage of the context label TLV is limited only to LDP P2MP LSPs on a LAN as specified in the next section. The context label TLV MUST NOT be used for any other purposes.

6. LDP Point-to-Multipoint LSPs on a LAN

This section describes one application of upstream label assignment using LDP. Further applications are to be described in separate documents.

[MLDP] describe how to setup P2MP LSPs using LDP. On a LAN the solution relies on "ingress replication". A LSR on a LAN, that is a branch LSR for a P2MP LSP, (say Ru) sends a separate copy of a packet that it receives on the P2MP LSP to each of the downstream LSRs on the LAN (say <Rd1...Rdn> that are adjacent to it in the P2MP LSP.

It is desirable for Ru to send a single copy of the packet for the LDP P2MP LSP on the LAN, when there are multiple downstream routers on the LAN that are adjacent to Ru in that LDP P2MP LSP. This requires that each of <Rd1...Rdn> must be able to associate the label L, used by Ru to transmit packets for the P2MP LSP on the LAN, with that P2MP LSP. It is possible to achieve this using LDP upstream-assigned labels with the following procedures.

Consider a LSR Rd that receives the LDP P2MP FEC [[MLDP](#)] from its downstream LDP peer. Further the upstream interface to reach LSR Ru which is the next-hop to the P2MP LSP root address, Pr, in the LDP P2MP FEC, is a LAN interface, Li. Further Rd and Ru support upstream-assigned labels. In this case Rd instead of sending a Label Mapping message as described in [[MLDP](#)] sends a Label Request message to Ru. This Label Request message MUST contain an Upstream Assigned Label Request TLV.

Ru on receiving this message sends back a Label Mapping message to Rd with an upstream-assigned label. This message also contains a MPLS Context Label TLV, as described in the previous section, with the value of the MPLS label set to a value assigned by Ru on interface Li as specified in [[MPLS-UPSTREAM](#)]. Processing of the Label Request and Label Mapping messages for LDP upstream-assigned labels is as described in [section 4.2](#). If Ru receives a Label Request for an upstream assigned label for the same P2MP FEC from multiple downstream LSRs on the LAN, <Rd1...Rdn>, it MUST send the same upstream-assigned label to each of <Rd1...Rdn>.

Ru transmits the MPLS packet using the procedures defined in [[MPLS-UPSTREAM](#)] and [[MPLS-MCAST-ENCAPS](#)]. The MPLS packet transmitted by Ru contains as the top label the context label assigned by Ru on the LAN interface, Li. The bottom label is the upstream label assigned by Ru to the LDP P2MP LSP. The top label is looked up in the context of the LAN interface, Li, [[MPLS-UPSTREAM](#)] by a downstream LSR on the LAN. This lookup enables the downstream LSR to determine the context specific label space to lookup the inner label in.

Note that <Rd1...Rdn> may have more than one equal cost next-hop on the LAN to reach Pr. It MAY be desirable for all of them to send the label request to the same upstream LSR and they MAY select one upstream LSR using the following procedure:

1. The candidate upstream LSRs are numbered from lower to higher IP address
2. The following hash is performed: $H = (\text{Sum Opaque value}) \text{ modulo } N$, where N is the number of candidate upstream LSRs. Opaque value is defined in [[MLDP](#)] and comprises the P2MP LSP identifier.

3. The selected upstream LSR U is the LSR that has the number H.

This allows for load balancing of a set of LSPs among a set of candidate upstream LSRs, while ensuring that on a LAN interface a single upstream LSR is selected. It is also to be noted that the procedures in this section can still be used by Rd and Ru if other LSRs on the LAN do not support upstream label assignment. Ingress replication and downstream label assignment will continue to be used for LSRs that do not support upstream label assignment.

7. IANA Considerations

This document defines a new LDP Upstream Label Assignment Capability Parameter. IANA is requested to assign the value 0x0507 to this Parameter.

This document defines a new LDP Upstream-Assigned Label Request TLV, IANA is requested to assign the type value of this TLV.

This document defines a new LDP Upstream-Assigned Label TLV, IANA is requested to assign the type value of this TLV.

This document defines three new Interface ID TLVs:

- RSVP-TE P2MP LSP TLV
- IP Multicast Tunnel TLV
- MPLS Context Label TLV

IANA is requested to assign the type values of these TLVs.

8. Acknowledgements

Thanks to Yakov Rekhter for his contribution. Thanks to Ina Minei and Thomas Morin for their comments. The hashing algorithm used on LAN interfaces is taken from [[MLDP](#)].

9. References

9.1. Normative References

[RFC3031] "MPLS Architecture", E. Rosen, A. Viswanathan, R. Callon, [RFC 3031](#).

[MPLS-UPSTREAM] R. Aggarwal, Y. Rekhter, E. Rosen, "MPLS Upstream Label Assignment and Context Specific Label Space", [draft-ietf-mpls-upstream-label-06.txt](#)

[MPLS-MCAST-ENCAPS] T. Eckert, E. Rosen, R. Aggarwal, Y. Rekhter, [draft-ietf-mpls-codepoint-10.txt](#)

[RFC2119] "Key words for use in RFCs to Indicate Requirement Levels.", Bradner, March 1997

[RFC3472] Ashwood-Smith, P. and L. Berger, Editors, " Generalized Multi-Protocol Label Switching (GMPLS) Signaling - Constraint-based Routed Label Distribution Protocol (CR-LDP) Extensions", [RFC 3472](#), January 2003.

[RFC3471] Berger, L. Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#) January 2003.

[RFC3036] L. Andersson, et. al., "LDP Specification", January 2001.

9.2. Informative References

[MVPN] E. Rosen, R. Aggarwal [Editors], "Multicast in BGP/MPLS VPNs", [draft-ietf-l3vpn-2547bis-mcast-06.txt](#)

[RFC4875] R. Aggarwal, D. Papadimitriou, S. Yasukawa [Editors], "Extensions to RSVP-TE for Point to Multipoint TE LSPs", [RFC 4875](#)

[MLDP] I. Minei et. al, "Label Distribution Protocol Extensions for Point-to-Multipoint Label Switched Paths", [draft-ietf-mpls-ldp-p2mp-05.txt](#)

[LDP-CAP] B. Thomas, et. al., "LDP Capabilities", [draft-thomas-mpls-ldp-capabilities-02.txt](#)

10. Author's Address

Rahul Aggarwal
Juniper Networks
1194 North Mathilda Ave.
Sunnyvale, CA 94089
Phone: +1-408-936-2720
Email: rahul@juniper.net

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
France
E-mail: jeanlouis.leroux@orange-ftgroup.com

11. Intellectual Property Statement

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.

12. Full Copyright Statement

Copyright (C) The IETF Trust (2007). 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.

