

Multiple LDP Instances
draft-pdutta-mpls-multi-ldp-instance-00

Abstract

This document defines an extension to Label Distribution Protocol (LDP) [[RFC5036](#)] for implementation of multiple LDP instances in a network node, where all such instances share the common data plane. Multiple LDP instances provide a method for operators for fate separation of various LDP FEC Types as well as for network segmentation. The methods defined in this extension are backward compatible with procedures defined in [[RFC5036](#)]

Requirements Language

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](#)].

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 October 7, 2012.

Copyright Notice

Copyright (c) 2012 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	3
2.	Multiple LDP Instances	4
2.1.	Procedures for multi-instance peering	4
2.1.1.	Case 1	5
2.1.2.	Case 2	6
2.1.3.	Case 3	7
2.1.4.	Case 4	8
3.	Detection of multi-instance peering	8
4.	LDP Address Distribution with multi-instance peering	9
5.	LDP State Sharing between instances	9
6.	Applicability	9
7.	IANA Considerations	9
8.	Security Considerations	10
9.	Acknowledgements	10
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	10
Appendix A.	An Appendix	11
	Authors' Addresses	11

1. Introduction

The Multi-Protocol Label Switching (MPLS) architecture is described in [RFC3031]. Label Distribution Protocol (LDP) is a signaling protocol for setup and maintenance of MPLS LSPs (Label Switched Paths) and the protocol specification is defined in [RFC5036].

Two Label Switched Routers (LSR) that use LDP to exchange label/FEC mapping information are known as "LDP Peers" with respect to that information, and we speak of there being an "LDP Session" between them. A single LDP session allows each peer to learn the other's label mappings. Each LSR is identified by an LDP identifier. An LDP Identifier is a six octet quantity used to identify an LSR label space. The 4 octets identify the LSR and is a globally unique value, such as a 32-bit router Id assigned to the LSR. The last two octets identify a specific label space within the LSR. The last two octets of LDP Identifiers for platform-wide label spaces are always both zero. This document uses the following representation for LDP Identifiers:

<LSR Id> : <label space id>

e.g, lsr171:0, lsr19:2 etc

As per [RFC5036] an LSR that manages and advertises multiple label spaces uses a different LDP Identifier for each such label space. This means for a single label space there can be only one router-id that is can be assigned to the node that exclusively owns that label space. For example, it is not possible to have two LSRs like lsr100:0 and lsr200:0 to be created in the a single node.

A LDP peering session between two LSRs may exchange labels for setting up LSPs that may belong to different FEC types. Operators may need the flexibility for fate separation of different FEC types in LDP protocol signaling when all such fec types share the same common label space. This is not possible with the current paradigm of single peering session between two LSRs and it requires one session per fate separated group of FEC types to exchange labels. Thus multiple LDP sessions are required between two peering nodes. One example could be fate separation between IP transport network and the overlay network of Pseudowires (PW). Procedures for PW set-up and maintenance using LDP are defined in [RFC4447]. It may be also desirable for fate separation IPv4 and IPv6 LSP set-up and maintenance in LDP in case of which two separate LDP sessions need to be formed between two peering nodes.

Although [RFC5036] does not specify that the 4 byte router-id of the LDP identifier be routable IP addresses, for various operational

simplicity implementations may map the 32 bit router-id to a IPv4 address configured in the node which is routable. In that way uniqueness of the 4 byte router-id can be achieved over a single routing domain. Interior Gateway Protocols (IGPs) like OSPF provide the option of creation of multiple instances for segmentation of a network into multiple routing domains. When LDP is deployed in such networks it is required to segment LDP network to align with multiple routing domains. When a node is connected to multiple such domains, LDP peering sessions over all such domains cannot use a common IPv4 router-id which is local to that node, since the IPv4 mapped router-id may not be routable across all such domains for security purposes. There are applications such as BGP Autodiscovery of L2VPNs or Dynamic MS-PW set-up that may auto-instantiate Targeted LDP sessions where BGP IPv4 next-hop addresses for respective NLRI are mapped to peer LDP identifiers. Such next-hop addresses may not be routable between two routing domains. Thus there is need to host multiple LSRs by a network node that shares the same label space but each with unique router-ids.

This document describes a method to implement multiple instances of LDP in a network node that shares same label space. The method is generic and is backward compatible with nodes that supports procedures defined in [[RFC5036](#)] but does not support the procedures defined in this document. The procedures defined in this document would be referred as "Multi-Instance LDP".

2. Multiple LDP Instances

The solution defines the concept of implementing multiple LDP instances on a single network node that shares the single label space and thus shares the common data plane. Each such LDP instance is identified by a unique 4 byte router-id but same label space. Since the multi-instance procedures use same LDP Identifier as defined in [[RFC5036](#)], it makes the node running multiple instances to be backward compatible with the node that support the multi-instance LDP procedures.

2.1. Procedures for multi-instance peering

When multiple LDP instances are set-up between two peering nodes for fate separation reasons then there can be various ways Hello adjacencies can be formed over the interfaces between the nodes. Further multi-instance peering for fate separation results in multiple parallel sessions between two peering nodes.

While running parallel multi-instance LDP sessions between two peering nodes,

1. Each peering session MUST use separate transport address.
2. The FEC label mappings exchanged over each peering session MUST be a disjoint set from one another.

The above rules does not apply between multi-instance LDP sessions with different peering LDP nodes.

This document describes the following cases and defines the rules and procedures with each case.

2.1.1. Case 1




Figure 1.

In this case the operator wants to separate the fate of FECs exchanged between the nodes into two separate groups - Group 1 and Group 2. For example Group 1 can contain all transport specific FEC types such as IPV4 FEC Element Type and LDP Multi-point (MP) FEC types etc. LDP Multi-point FEC types are described in [\[RFC6388\]](#). Group 2 contains various Pseudowire (PW) FEC types. PW setup and maintenance using LDP is described in [\[RFC4447\]](#).

Two separate LSR-IDs are provisioned in each node - one LSR is dedicated for FEC Group 1 and another for FEC Group 2.

There are two parallel interfaces between Node-A and Node-B as IF1 and IF2 respectively.

The traffic for LSPs set-up for FEC Group 1 may use both IF1 and IF2. Thus both IF1 and IF2 would exchange Hello Packets using LSR-A1:0 and LSR-A2:0 for setting up Hello adjacency for the LDP instance assigned for FEC Group 1. The Hello messages exchanged over IF1 and IF2 MUST carry the LDP Adjacency Capabilities for each FEC Types in FEC Group 1. LDP Adjacency Capabilities are defined in [\[LDP-ADJ-CAP\]](#). This would result in formation of a LDP session between Node A and Node B for the instance identified by LSR-A1:0 and LSR-B1:0 respectively. The LDP session SHOULD be set-up with Capabilities of FEC Group 1.

LDP session specific capability negotiation is described in [[RFC5561](#)]

A Targeted LDP (t-LDP) hello adjacency would be formed between node A and node B using LSR-A2:0 and LSR-B2:0 respectively. The t-Ldp Hello Messages exchanged between the nodes MUST carry the LDP Adjacency Capabilities for each FEC Types in FEC Group 2. This would result in a LDP session between Node A and Node B for the instance identified by LSR-A2:0 and LSR-B2:0 respectively. The LDP session SHOULD be set-up with capabilities of FEC Group 2.

[2.1.2.](#) Case 2




Figure 2.

This is a variant of case 1 where the operator may choose to further separate the fate of IPV4 FEC Element Type and MP FEC Element types into "Unicast" and "Multicast" Groups. Thus there are three FEC Groups here and fate separation is required for all three FEC Groups.

FEC Group 1 : IPV4 FEC Element Type.

FEC Group 2: MP FEC Element Types.

FEC Group 3: PW FEC Element Types.

LDP Instance 1: The LDP instance with peering LSR-A1:0 and LSR-B1:0 are assigned for FEC Group 1.

LDP Instance 2: The LDP Instance with peering LSR-A2:0 and LSR-B2:0 are assigned for FEC Group 2.

LDP Instance 3: The LDP instance with peering LSR-A3:0 and LSR-B3:0 are assigned for FEC Group 3.

In this case, both IF1 and IF2 are associated with LDP instances 1 and 2. Each of IF1 and IF2 would originate two separate Hello Messages using the same source IP address, one Hello Message for each

instance . This would result in two hello adjacencies per interface - one for Instance 1 and Instance 2. Each Hello Adjacencies SHOULD advertise capabilities using rules described in case 1.

Such case may also arise when operator wants to do fate separation of IPV4 and IPV6 LDP based LSPs but IF1 and IF2 are single stack interfaces only - that is either IPV4 or IPV6. Thus an operator may provision single stack interfaces IF1 and IF2 and yet can provision fate separation of IPV4 and IPV6 LSPs.

The t-Ldp Hello Adjacency would be formed for LDP Instance 3 using the PW Capabilities.

2.1.3. Case 3




Figure 3.

This case is a variant of case 2 where, both interfaces IF1 and IF2 are dual-stack (IPV4 and IPV6) interfaces and operator wants fate separation of IPV4 and IPV6 LSPs. Without loss of generality, hereby IPV4 or IPV6 FECs may include all FEC types that are associated with IPV4 or IPV6. For example, [\[I-D.ietf-mpls-mldp-in-band-signaling\]](#) defines several in-band MP FEC Types that may be classified into IPV6.

LDP Instance 1: The LDP instance with peering LSR-A1:0 and LSR-B1:0 are assigned for IPV4 FEC Types.

LDP Instance 2: The LDP Instance with peering LSR-A2:0 and LSR-B2:0 are assigned for IPV6 FEC Types.

Both the interfaces IF1 and IF2 are associated with each of the LDP instances 1 and 2 respectively. Here the operator may choose to use IPV4 addresses on the interfaces for sending Hello Messages for Instance 1 and IPV6 addresses on the interfaces for sending Hello Messages for Instance 2.

2.1.4. Case 4

This case is variant of case 3 where interface IF1 is dedicated for IPV4 LSP Types and IF2 is dedicated for IPV6 LSP Types. This provides fate-separation of both control plane and data plane for LSP types.

3. Detection of multi-instance peering

While running parallel multi-instance LDP sessions between two peering nodes, it is important to detect that such sessions with the same peer node. If a node receives the same FEC label mapping from parallel multi-LSR peering sessions it may result in a loop for some applications. An example of such application can be LDP based Virtual Private LAN Service (VPLS) described [RFC4762]. So it is important to detect and prevent such loops.

This document defines a new LDP Node-ID TLV that uniquely identifies the node that hosts multiple LDP instances. The LDP Node ID TLV is OPTIONAL and is carried in LDP Hello Messages sent out by the node in its Optional Parameters. The encoding of the LDP Node ID TLV is as follows:




Figure 4

The Value field is a 48 bit identifier and MUST be unique identifier across the network.

1. All the multi-instance LDP LSRs MUST advertise the same LDP Node-ID TLV in all Hello Messages originated by that node. One example of the value can be a IEEE Vendor specific MAC Address that can uniquely identify a node in the network.
2. When a LSR receives a FEC label mapping from a peering session but same FEC mapping has been already received over another peering session associated with same Node-ID then the receiving LSR MUST send a Label Release to the peering session with status code

LOOP_DETECTED.

4. LDP Address Distribution with multi-instance peering

An LSR maintains learned labels in a Label Information Base (LIB). When operating in Downstream Unsolicited mode, the LIB entry for an address prefix associates a collection of (LDP Identifier, label) pairs with the prefix, one such pair for each peer advertising a label for the prefix. When the next hop for a prefix changes, the LSR retrieves the label advertised by the new next hop from the LIB for use in forwarding. To retrieve the label, the LSR should be able to map the next hop address for the prefix to an LDP Identifier. Similarly, when the LSR learns a label for a prefix from an LDP peer, it should be able to determine whether that peer is currently a next hop for the prefix to determine whether it needs to start using the newly learned label when forwarding packets that match the prefix. To make that decision, the LSR should be able to map an LDP Identifier to the peer's addresses to check whether any are a next hop for the prefix. To enable LSRs to map between a peer LDP Identifier and the peer's addresses, LSRs advertise their addresses using LDP Address and Withdraw Address messages as per procedures defined in [[RFC5036](#)]

However while running multi-instance LDP peering between two nodes, it is possible that all such sessions would distribute same set of local addresses in each node. An implementation MAY segregate the local address space in each node among the multiple ldp instances to avoid duplication of address distribution.

5. LDP State Sharing between instances

TBD.

6. Applicability

This solution described in this document is applicable for multi-instance LDP sessions for fate separation as well as for segmentation of LDP network domains. More details would be covered in next revisions of the document.

7. IANA Considerations

This document requests the following code points:

- LDP Node-ID TLV type.

Note to RFC Editor: this section may be removed on publication as an RFC.

8. Security Considerations

[I-D.ietf-mpls-mpls-and-gmpls-security-framework] describes the security framework for MPLS networks. whereas [RFC5036] describes the security considerations that apply to the base LDP specification. The same security framework and considerations apply to the capability mechanism described in this document.

9. Acknowledgements

The authors would like to thank Wim Henderickx for insightful comments and probing questions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), January 2001.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", [RFC 5561](#), July 2009.
- [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.

10.2. Informative References

- [I-D.ietf-mpls-mldp-in-band-signaling]
Wijnands, I., Eckert, T., Leymann, N., and M. Napierala,
"Multipoint LDP in-band signaling for Point-to-Multipoint
and Multipoint- to-Multipoint Label Switched Paths",

[draft-ietf-mpls-mlbp-in-band-signaling-05](#) (work in progress), December 2011.

- [I-D.ietf-mpls-mpls-and-gmpls-security-framework]
Fang, L. and M. Behringer, "Security Framework for MPLS and GMPLS Networks",
[draft-ietf-mpls-mpls-and-gmpls-security-framework-09](#) (work in progress), March 2010.
- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", [RFC 4447](#), April 2006.
- [RFC4762] Lasserre, M. and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), January 2007.

[Appendix A](#). An Appendix

Authors' Addresses

Pranjal Kumar Dutta
Alcatel-Lucent
701 E Middlefield Road
Mountain View, California 94043
USA

Phone:
Fax:
Email: pranjal.dutta@alcatel-lucent.com

Mustapha Aissaoui
Alcatel-Lucent
600 May Road
Kanata, ON
Canada

Phone:
Fax:
Email: mustapha.aissaoui@alcatel-lucent.com
URI:

