

Internet Engineering Task Force (IETF)  
Request for Comments: 7552  
Updates: 5036, 6720  
Category: Standards Track  
ISSN: 2070-1721

R. Asati  
C. Pignataro  
K. Raza  
Cisco  
V. Manral  
Ionos Networks  
R. Papneja  
Huawei  
June 2015

## Updates to LDP for IPv6

### Abstract

The Label Distribution Protocol (LDP) specification defines procedures to exchange label bindings over either IPv4 or IPv6 networks, or both. This document corrects and clarifies the LDP behavior when an IPv6 network is used (with or without IPv4). This document updates RFCs 5036 and 6720.

### 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/rfc7552>.

## Copyright Notice

Copyright (c) 2015 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

## Table of Contents

1. Introduction .....	4
1.1. Topology Scenarios for Dual-Stack Environment .....	5
1.2. Single-Hop vs. Multi-Hop LDP Peering .....	6
2. Specification Language .....	6
3. LSP Mapping .....	7
4. LDP Identifiers .....	8
5. Neighbor Discovery .....	8
5.1. Basic Discovery Mechanism .....	8
5.1.1. Maintaining Hello Adjacencies .....	9
5.2. Extended Discovery Mechanism .....	10
6. LDP Session Establishment and Maintenance .....	10
6.1. Transport Connection Establishment .....	10
6.1.1. Dual-Stack: Transport Connection Preference and Role of an LSR .....	12
6.2. LDP Session Maintenance .....	14
7. Binding Distribution .....	15
7.1. Address Distribution .....	15
7.2. Label Distribution .....	16
8. LDP Identifiers and Duplicate Next-Hop Addresses .....	17
9. LDP TTL Security .....	18
10. IANA Considerations .....	18
11. Security Considerations .....	19
12. References .....	19
12.1. Normative References .....	19
12.2. Informative References .....	20
Appendix A. Additional Considerations .....	21
A.1. LDPv6 and LDPv4 Interoperability Safety Net .....	21
A.2. Accommodating Implementations Not Compliant with RFC 5036 ..	21
A.3. Why prohibit IPv4-mapped IPv6 addresses in LDP? .....	22
A.4. Why a 32-bit value even for the IPv6 LDP Router Id? .....	22
Acknowledgments .....	23
Contributors .....	23
Authors' Addresses.....	24

## 1. Introduction

The LDP specification [RFC5036] defines procedures and messages for exchanging FEC-label bindings over either IPv4 or IPv6 networks, or both (i.e., Dual-stack networks).

However, RFC 5036 has the following deficiencies (i.e., lacks details) in regard to IPv6 usage (with or without IPv4):

1. Label Switched Path (LSP) Mapping: No rule for mapping a particular packet to a particular LSP that has an Address Prefix Forwarding Equivalence Class (FEC) element containing the IPv6 address of the egress router
2. LDP Identifier: No details specific to IPv6 usage
3. LDP Discovery: No details for using a particular IPv6 destination (multicast) address or the source address
4. LDP Session Establishment: No rule for handling both IPv4 and IPv6 Transport Address optional objects in a Hello message, and subsequently two IPv4 and IPv6 transport connections
5. LDP Address Distribution: No rule for advertising IPv4 and/or IPv6 address bindings over an LDP session
6. LDP Label Distribution: No rule for advertising IPv4 and/or IPv6 FEC-label bindings over an LDP session, or for handling the coexistence of IPv4 and IPv6 FEC Elements in the same FEC TLV
7. Next-Hop Address Resolution: No rule for accommodating the usage of duplicate link-local IPv6 addresses
8. LDP Time to Live (TTL) Security: No rule for a built-in Generalized TTL Security Mechanism (GTSM) in LDP with IPv6 (this is a deficiency in [RFC6720])

This document addresses the above deficiencies by specifying the desired behavior/rules/details for using LDP in IPv6-enabled networks (IPv6-only or Dual-stack networks). This document closes the IPv6 MPLS gap discussed in Sections 3.2.1, 3.2.2, and 3.3.1.1 of [RFC7439].

Note that this document updates [RFC5036] and [RFC6720].

### 1.1. Topology Scenarios for Dual-Stack Environment

Two Label Switching Routers (LSRs) may involve Basic and/or Extended LDP Discovery in IPv6 and/or IPv4 address families in various topology scenarios.

This document addresses the following three topology scenarios in which the LSRs may be connected via one or more Dual-stack LDP-enabled interfaces (Figure 1), or one or more Single-stack LDP-enabled interfaces (Figures 2 and 3):

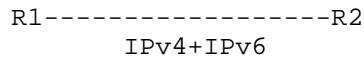


Figure 1: LSRs Connected via a Dual-Stack Interface

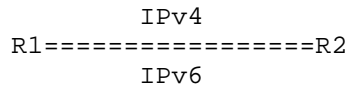


Figure 2: LSRs Connected via Two Single-Stack Interfaces

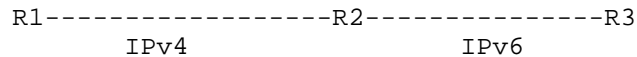


Figure 3: LSRs Connected via a Single-Stack Interface

Note that the topology scenario illustrated in Figure 1 also covers the case of a Single-stack LDP-enabled interface (say, IPv4) being converted to a Dual-stack LDP-enabled interface (by enabling IPv6 routing as well as IPv6 LDP), even though the LDP-over-IPv4 (LDPoIPv4) session may already be established between the LSRs.

Note that the topology scenario illustrated in Figure 2 also covers the case of two routers getting connected via an additional Single-stack LDP-enabled interface (IPv6 routing and IPv6 LDP), even though the LDPoIPv4 session may already be established between the LSRs over the existing interface(s).

This document also addresses the scenario in which the LSRs do the Extended Discovery in IPv6 and/or IPv4 address families:

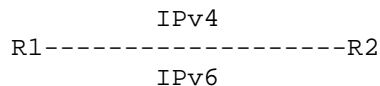


Figure 4: LSRs Involving IPv4 and IPv6 Address Families

## 1.2. Single-Hop vs. Multi-Hop LDP Peering

The LDP TTL Security mechanism specified by this document applies only to single-hop LDP peering sessions, not to multi-hop LDP peering sessions, in line with Section 5.5 of [RFC5082]. [RFC5082] describes the Generalized TTL Security Mechanism (GTSM).

As a consequence, any LDP feature that relies on a multi-hop LDP peering session would not work with GTSM and will warrant (statically or dynamically) disabling GTSM. Please see Section 9.

## 2. Specification 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 [RFC2119].

### Abbreviations:

LDP	Label Distribution Protocol
LDPoIPv4	LDP-over-IPv4 transport connection
LDPoIPv6	LDP-over-IPv6 transport connection
FEC	Forwarding Equivalence Class
TLV	Type Length Value
LSR	Label Switching Router
LSP	Label Switched Path
LSPv4	IPv4-signaled Label Switched Path
LSPv6	IPv6-signaled Label Switched Path
AFI	Address Family Identifier

LDP Id	LDP Identifier
Single-stack LDP	LDP supporting just one address family (for discovery, session setup, address/label binding exchange, etc.)
Dual-stack LDP	LDP supporting two address families (for discovery, session setup, address/label binding exchange, etc.)
Dual-stack LSR	LSR supporting Dual-stack LDP for a peer
Single-stack LSR	LSR supporting Single-stack LDP for a peer

Note that an LSR can be a Dual-stack and Single-stack LSR at the same time for different peers. This document loosely uses the term "address family" to mean "IP address family".

### 3. LSP Mapping

Section 2.1 of [RFC5036] specifies the procedure for mapping a particular packet to a particular LSP using three rules. Quoting the third rule from [RFC5036]:

If it is known that a packet must traverse a particular egress router, and there is an LSP that has an Address Prefix FEC element that is a /32 address of that router, then the packet is mapped to that LSP.

This rule is correct for IPv4 (to set up LSPv4), but not for IPv6 (to set up LSPv6), since an IPv6 router may even have a /64 or /96 or /128 (or whatever prefix length) address. Hence, that rule is updated here to use IPv4 or IPv6 addresses instead of /32 or /128 addresses, as shown below:

If it is known that a packet must traverse a particular egress router, and there is an LSP that has an Address Prefix FEC element that is an IPv4 or IPv6 address of that router, then the packet is mapped to that LSP.

#### 4. LDP Identifiers

In line with Section 2.2.2 of [RFC5036], this document specifies the usage of a 32-bit (unsigned non-zero integer) LSR Id on an IPv6-enabled LSR (with or without Dual-stacking).

This document also qualifies the first sentence of the last paragraph of Section 2.5.2 of [RFC5036] to be per address family.

From Section 2.5.2 of [RFC5036]:

An LSR MUST advertise the same transport address in all Hellos that advertise the same label space.

Updated by this document, as follows:

For a given address family, an LSR MUST advertise the same transport address in all Hellos that advertise the same label space.

This rightly enables the per-platform label space to be shared between IPv4 and IPv6.

In summary, this document mandates the usage of a common LDP Identifier (the same LSR Id and label space id) for both IPv4 and IPv6 address families.

#### 5. Neighbor Discovery

If Dual-stack LDP is enabled (i.e., LDP enabled in both IPv6 and IPv4 address families) on an interface or for a targeted neighbor, then the LSR MUST transmit both IPv6 and IPv4 LDP (Link or targeted) Hellos and include the same LDP Identifier (assuming per-platform label space usage) in them.

If Single-stack LDP is enabled (i.e., LDP enabled in either an IPv6 or IPv4 address family), then the LSR MUST transmit either IPv6 or IPv4 LDP (Link or targeted) Hellos, respectively.

##### 5.1. Basic Discovery Mechanism

Section 2.4.1 of [RFC5036] defines the Basic Discovery mechanism for directly connected LSRs. Following this mechanism, LSRs periodically send LDP Link Hellos destined to the "all routers on this subnet" group multicast IP address.



Interestingly enough, per the IPv6 addressing architecture [RFC4291], IPv6 has three "all routers on this subnet" multicast addresses:

ff01:0:0:0:0:0:0:2 = Interface-local scope

ff02:0:0:0:0:0:0:2 = Link-local scope

ff05:0:0:0:0:0:0:2 = Site-local scope

[RFC5036] does not specify which particular IPv6 "all routers on this subnet" group multicast IP address should be used by LDP Link Hellos.

This document specifies the usage of link-local scope (i.e., ff02:0:0:0:0:0:0:2) as the destination multicast IP address in IPv6 LDP Link Hellos. An LDP Link Hello packet received on any of the other destination addresses MUST be dropped. Additionally, the link-local IPv6 address MUST be used as the source IP address in IPv6 LDP Link Hellos.

Also, the LDP Link Hello packets MUST have their IPv6 Hop Limit set to 255, be checked for the same upon receipt (before any LDP-specific processing), and be handled as specified in Section 3 of [RFC5082]. The built-in inclusion of GTSM automatically protects IPv6 LDP from off-link attacks.

More importantly, if an interface is a Dual-stack LDP interface (i.e., LDP enabled in both IPv6 and IPv4 address families), then the LSR MUST periodically transmit both IPv6 and IPv4 LDP Link Hellos (using the same LDP Identifier per Section 4) on that interface and be able to receive them. This facilitates discovery of IPv6-only, IPv4-only, and Dual-stack peers on the interface's subnet and ensures successful subsequent peering using the appropriate (address family) transport on a multi-access or broadcast interface.

#### 5.1.1.1. Maintaining Hello Adjacencies

In the case of a Dual-stack LDP-enabled interface, the LSR SHOULD maintain Link Hello adjacencies for both IPv4 and IPv6 address families. This document, however, allows an LSR to maintain Receive-side Link Hello adjacencies only for the address family that has been used for the establishment of the LDP session (whether an LDPoIPv4 or LDPoIPv6 session).

## 5.2. Extended Discovery Mechanism

The Extended Discovery mechanism (defined in Section 2.4.2 of [RFC5036]), in which the targeted LDP Hellos are sent to a unicast IPv6 address destination, requires only one IPv6-specific consideration: the link-local IPv6 addresses MUST NOT be used as the targeted LDP Hello packet's source or destination addresses.

## 6. LDP Session Establishment and Maintenance

Section 2.5.1 of [RFC5036] defines a two-step process for LDP session establishment, once the neighbor discovery has completed (i.e., LDP Hellos have been exchanged):

1. Transport connection establishment
2. Session initialization

Section 6.1 discusses the LDP considerations for IPv6 and/or Dual-stacking in the context of session establishment, whereas Section 6.2 discusses the LDP considerations for IPv6 and/or Dual-stacking in the context of session maintenance.

### 6.1. Transport Connection Establishment

Section 2.5.2 of [RFC5036] specifies the use of a Transport Address optional object (TLV) in LDP Hello messages to convey the transport (IP) address; however, it does not specify the behavior of LDP if both IPv4 and IPv6 Transport Address objects (TLVs) are sent in a Hello message or separate Hello messages. More importantly, it does not specify whether both IPv4 and IPv6 transport connections should be allowed if both IPv4 and IPv6 Hello adjacencies were present prior to session establishment.

This document specifies the following:

1. An LSR MUST NOT send a Hello message containing both IPv4 and IPv6 Transport Address optional objects. In other words, there MUST be at most one Transport Address optional object in a Hello message. An LSR MUST include only the transport address whose address family is the same as that of the IP packet carrying the Hello message.
2. An LSR SHOULD accept the Hello message that contains both IPv4 and IPv6 Transport Address optional objects but MUST use only the transport address whose address family is the same as that of the IP packet carrying the Hello message. An LSR SHOULD accept only the first Transport Address optional object for a given address

family in the received Hello message and ignore the rest if the LSR receives more than one Transport Address optional object for a given address family.

3. An LSR MUST send separate Hello messages (each containing either an IPv4 or IPv6 Transport Address optional object) for each IP address family if Dual-stack LDP is enabled (for an interface or neighbor).
4. An LSR MUST use a global unicast IPv6 address in an IPv6 Transport Address optional object of outgoing targeted Hellos and check for the same in incoming targeted Hellos (i.e., MUST discard the targeted Hello if it failed the check).
5. An LSR MUST prefer using a global unicast IPv6 address in an IPv6 Transport Address optional object of outgoing Link Hellos if it had to choose between a global unicast IPv6 address and a unique-local or link-local IPv6 address.
6. A Single-stack LSR MUST establish either an LDPoIPv4 or LDPoIPv6 session with a remote LSR as per the enabled address family.
7. A Dual-stack LSR MUST NOT initiate or accept the request for a TCP connection for a new LDP session with a remote LSR if it already has an LDPoIPv4 or LDPoIPv6 session for the same LDP Identifier established with that remote LSR.

This means that only one transport connection is established, regardless of IPv6 and/or IPv4 Hello adjacencies present between two LSRs.

8. A Dual-stack LSR SHOULD prefer establishing an LDPoIPv6 session (instead of an LDPoIPv4 session) with a remote Dual-stack LSR by following the 'transport connection role' determination logic in Section 6.1.1.

Additionally, to ensure the above preference in the case where Dual-stack LDP is enabled on an interface, it would be desirable that IPv6 LDP Link Hellos are transmitted before IPv4 LDP Link Hellos, particularly when an interface is coming into service or being reconfigured.

6.1.1.1. Dual-Stack: Transport Connection Preference and Role of an LSR

Section 2.5.2 of [RFC5036] specifies the rules for determining active/passive roles in setting up a TCP connection. These rules are clear for Single-stack LDP but not for Dual-stack LDP, in which an LSR may assume different roles for different address families, causing the LDP session to not get established.

To ensure a deterministic transport connection (active/passive) role in the case of Dual-stack LDP, this document specifies that the Dual-stack LSR conveys its transport connection preference in every LDP Hello message. This preference is encoded in a new TLV, named the "Dual-Stack capability" TLV, as defined below:




Figure 5: Dual-Stack Capability TLV

Where:

U and F bits: 1 and 0 (as specified by [RFC5036])

Dual-Stack capability: TLV code point (0x0701)

TR: Transport Connection Preference

This document defines the following two values:

0100: LDPoIPv4 connection

0110: LDPoIPv6 connection (default)

Reserved

This field is reserved. It MUST be set to zero on transmission and ignored on receipt.

A Dual-stack LSR (i.e., an LSR supporting Dual-stack LDP for a peer) MUST include the Dual-Stack capability TLV in all of its LDP Hellos and MUST set the "TR" field to announce its preference for either an LDPoIPv4 or LDPoIPv6 transport connection for that peer. The default preference is LDPoIPv6.

A Dual-stack LSR MUST always check for the presence of the Dual-Stack capability TLV in the received Hello messages and take appropriate action, as follows:

1. If the Dual-Stack capability TLV is present and the remote preference does not match the local preference (or does not get recognized), then the LSR MUST discard the Hello message and log an error.

If an LDP session was already in place, then the LSR MUST send a fatal Notification message with status code of 'Transport Connection Mismatch' (0x00000032) and reset the session.

2. If the Dual-Stack capability TLV is present and the remote preference matches the local preference, then:
  - a) If TR=0100 (LDPoIPv4), then determine the active/passive roles for the TCP connection using an IPv4 transport address as defined in Section 2.5.2 of RFC 5036.
  - b) If TR=0110 (LDPoIPv6), then determine the active/passive roles for the TCP connection by using an IPv6 transport address as defined in Section 2.5.2 of RFC 5036.
3. If the Dual-Stack capability TLV is NOT present and
  - a) only IPv4 Hellos are received, then the neighbor is deemed as a legacy IPv4-only LSR (supporting Single-stack LDP); hence, an LDPoIPv4 session SHOULD be established (similar to that of 2a above).

However, if IPv6 Hellos are also received at any time during the life of the session from that neighbor, then the neighbor is deemed as a noncompliant Dual-stack LSR (similar to that of 3c below), resulting in any established LDPoIPv4 session being reset and a fatal Notification message being sent (with status code of 'Dual-Stack Noncompliance', 0x00000033).

- b) only IPv6 Hellos are received, then the neighbor is deemed as an IPv6-only LSR (supporting Single-stack LDP) and an LDPoIPv6 session SHOULD be established (similar to that of 2b above).

However, if IPv4 Hellos are also received at any time during the life of the session from that neighbor, then the neighbor is deemed as a noncompliant Dual-stack LSR (similar to that of 3c below), resulting in any established LDPoIPv6 session being reset and a fatal Notification message being sent (with status code of 'Dual-Stack Noncompliance', 0x00000033).

- c) both IPv4 and IPv6 Hellos are received, then the neighbor is deemed as a noncompliant Dual-stack neighbor and is not allowed to have any LDP session. A Notification message should be sent (with status code of 'Dual-Stack Noncompliance', 0x00000033).

A Dual-stack LSR MUST convey the same transport connection preference ("TR" field value) in all (link and targeted) Hellos that advertise the same label space to the same peer and/or on the same interface. This ensures that two LSRs linked by multiple Hello adjacencies using the same label spaces play the same connection establishment role for each adjacency.

A Dual-stack LSR MUST follow Section 2.5.5 of [RFC5036] and check for matching Hello messages from the peer (either all Hellos also include the Dual-Stack capability (with the same TR value) or none do).

A Single-stack LSR does not need to use the Dual-Stack capability in Hello messages and SHOULD ignore this capability if received.

An implementation may provide an option to favor one AFI (say, IPv4) over another AFI (say, IPv6) for the TCP transport connection, so as to use the favored IP version for the LDP session and force deterministic active/passive roles.

Note: An alternative to this new capability TLV could be a new Flag value in an LDP Hello message; however, it would be used even in Single-stack IPv6 LDP networks and linger on forever, even though Dual-stack will not. Hence, the idea of this alternative has been discarded.

## 6.2. LDP Session Maintenance

This document specifies that two LSRs maintain a single LDP session, regardless of the number of Link or targeted Hello adjacencies between them, as described in Section 6.1. This is independent of whether:

- they are connected via a Dual-stack LDP-enabled interface(s) or via two (or more) Single-stack LDP-enabled interfaces;
- a Single-stack LDP-enabled interface is converted to a Dual-stack LDP-enabled interface (see Figure 1) on either LSR;
- an additional Single-stack or Dual-stack LDP-enabled interface is added or removed between two LSRs (see Figure 2).

If the last Hello adjacency for a given address family goes down (e.g., due to Dual-stack LDP-enabled interfaces being converted into Single-stack LDP-enabled interfaces on one LSR) and that address family is the same as the one used in the transport connection, then the transport connection (LDP session) MUST be reset. Otherwise, the LDP session MUST stay intact.

If the LDP session is torn down for whatever reason (LDP disabled for the corresponding transport, Hello adjacency expiry, preference mismatch, etc.), then the LSRs SHOULD initiate the establishment of a new LDP session as per the procedures described in Section 6.1 of this document.

## 7. Binding Distribution

LSRs by definition can be enabled for Dual-stack LDP globally and/or per peer so as to exchange the address and label bindings for both IPv4 and IPv6 address families, independent of any LDPoIPv4 or LDPoIPv6 session between them.

However, there might be some legacy LSRs that are fully compliant with RFC 5036 for IPv4 but are noncompliant for IPv6 (for example, see Section 3.5.5.1 of RFC 5036), causing them to reset the session upon receiving IPv6 address bindings or IPv6 FEC (Prefix) label bindings from a peer compliant with this document. This is somewhat undesirable, as clarified further in Appendices A.1 and A.2.

To help maintain backward compatibility (i.e., accommodate IPv4-only LDP implementations that may not be compliant with RFC 5036, Section 3.5.5.1), this specification requires that an LSR MUST NOT send any IPv6 bindings to a peer if the peer has been determined to be a legacy LSR.

The Dual-Stack capability TLV, which is defined in Section 6.1.1, is also used to determine whether or not a peer is a legacy (IPv4-only Single-stack) LSR.

### 7.1. Address Distribution

An LSR MUST NOT advertise (via an Address message) any IPv4-mapped IPv6 addresses (as defined in Section 2.5.5.2 of [RFC4291]) and MUST ignore such addresses if ever received. Please see Appendix A.3.

If an LSR is enabled with Single-stack LDP for any peer, then it MUST advertise (via an Address message) its local IP addresses as per the enabled address family to that peer and process received Address messages containing IP addresses as per the enabled address family from that peer.

If an LSR is enabled with Dual-stack LDP for a peer and

1. does not find the Dual-Stack capability TLV in the incoming IPv4 LDP Hello messages from that peer, then the LSR MUST NOT advertise its local IPv6 addresses to the peer.
2. finds the Dual-Stack capability TLV in the incoming IPv4 (or IPv6) LDP Hello messages from that peer, then it MUST advertise (via an Address message) its local IPv4 and IPv6 addresses to that peer.
3. does not find the Dual-Stack capability TLV in the incoming IPv6 LDP Hello messages, then it MUST advertise (via an Address message) only its local IPv6 addresses to that peer.

This last point helps to maintain forward compatibility (no need to require this TLV in the case of IPv6 Single-stack LDP).

## 7.2. Label Distribution

An LSR MUST NOT allocate and MUST NOT advertise FEC-label bindings for link-local or IPv4-mapped IPv6 addresses (defined in Section 2.5.5.2 of [RFC4291]), and it MUST ignore such bindings if ever received. Please see Appendix A.3.

If an LSR is enabled with Single-stack LDP for any peer, then it MUST advertise (via a Label Mapping message) FEC-label bindings for the enabled address family to that peer and process received FEC-label bindings for the enabled address family from that peer.

If an LSR is enabled with Dual-stack LDP for a peer and

1. does not find the Dual-Stack capability TLV in the incoming IPv4 LDP Hello messages from that peer, then the LSR MUST NOT advertise IPv6 FEC-label bindings to the peer (even if IP capability negotiation for the IPv6 address family was done).
2. finds the Dual-Stack capability TLV in the incoming IPv4 (or IPv6) LDP Hello messages from that peer, then it MUST advertise FEC-label bindings for both IPv4 and IPv6 address families to that peer.
3. does not find the Dual-Stack capability TLV in the incoming IPv6 LDP Hello messages, then it MUST advertise FEC-label bindings for IPv6 address families to that peer.

This last point helps to maintain forward compatibility (no need to require this TLV for IPv6 Single-stack LDP).



An LSR MAY further constrain the advertisement of FEC-label bindings for a particular address family by negotiating the IP capability for a given address family, as specified in [RFC7473]. This allows an LSR pair to neither advertise nor receive the undesired FEC-label bindings on a per-address-family basis to a peer.

If an LSR is configured to change an interface or peer from Single-stack LDP to Dual-stack LDP, then an LSR SHOULD use Typed Wildcard FEC procedures [RFC5918] to request the label bindings for the enabled address family. This helps to relearn the label bindings that may have been discarded before, without resetting the session.

#### 8. LDP Identifiers and Duplicate Next-Hop Addresses

RFC 5036, Section 2.7 specifies the logic for mapping the IP routing next hop (of a given FEC) to an LDP peer so as to find the correct label entry for that FEC. The logic involves using the IP routing next-hop address as an index into the (peer address) database (which is populated by the Address message containing a mapping between each peer's local addresses and its LDP Identifier) to determine the LDP peer.

However, this logic is insufficient to deal with duplicate IPv6 (link-local) next-hop addresses used by two or more peers. The reason is that all interior IPv6 routing protocols (can) use link-local IPv6 addresses as the IP routing next hops, and "IP Version 6 Addressing Architecture" [RFC4291] allows a link-local IPv6 address to be used on more than one link.

Hence, this logic is extended by this specification to use not only the IP routing next-hop address but also the IP routing next-hop interface to uniquely determine the LDP peer(s). The next-hop address-based LDP peer mapping is to be done through the LDP peer address database (populated by Address messages received from the LDP peers), whereas next-hop interface-based LDP peer mapping is to be done through the LDP Hello adjacency/interface database (populated by Hello messages received from the LDP peers).

This extension solves the problem of two or more peers using the same link-local IPv6 address (in other words, duplicate peer addresses) as the IP routing next hops.

Lastly, for better scale and optimization, an LSR may advertise only the link-local IPv6 addresses in the Address message, assuming that the peer uses only the link-local IPv6 addresses as static and/or dynamic IP routing next hops.

## 9. LDP TTL Security

This document mandates the use of the Generalized TTL Security Mechanism (GTSM) [RFC6720] for LDP Link Hello packets over IPv6 (see Section 5.1).

This document further recommends enabling GTSM for the LDP/TCP transport connection over IPv6 (i.e., LDPoIPv6). This GTSM inclusion is intended to automatically protect IPv6 LDP peering sessions from off-link attacks.

[RFC6720] allows for the implementation to statically (via configuration) and/or dynamically override the default behavior (enable/disable GTSM) on a per-peer basis. Such an option could be set on either LSR in a peering session (since GTSM negotiation would ultimately disable GTSM between the LSR and its peer(s)).

LDP Link Hello packets MUST have their IPv6 Hop Limit set to 255 and be checked for the same upon receipt before any further processing, as per Section 3 of [RFC5082].

## 10. IANA Considerations

This document defines a new optional parameter for the LDP Hello message and two new status codes for the LDP Notification message.

The "Dual-Stack capability" parameter has been assigned a code point (0x0701) from the "TLV Type Name Space" registry. IANA has allocated this code point from the IETF Consensus range 0x0700-0x07ff for the Dual-Stack capability TLV.

The 'Transport Connection Mismatch' status code has been assigned a code point (0x00000032) from the "Status Code Name Space" registry. IANA has allocated this code point from the IETF Consensus range and marked the E bit column with a '1'.

The 'Dual-Stack Noncompliance' status code has been assigned a code point (0x00000033) from the "Status Code Name Space" registry. IANA has allocated this code point from the IETF Consensus range and marked the E bit column with a '1'.

## 11. Security Considerations

The extensions defined in this document only clarify the behavior of LDP; they do not define any new protocol procedures. Hence, this document does not add any new security issues to LDP.

While the security issues relevant for [RFC5036] are relevant for this document as well, this document reduces the chances of off-link attacks when using an IPv6 transport connection by including the use of GTSM procedures [RFC5082]. Please see Section 9 for LDP TTL Security details.

Moreover, this document allows the use of IPsec [RFC4301] for IPv6 protection; hence, LDP can benefit from the additional security as specified in [RFC7321] as well as [RFC5920].

## 12. References

### 12.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<http://www.rfc-editor.org/info/rfc4291>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<http://www.rfc-editor.org/info/rfc5036>>.
- [RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <<http://www.rfc-editor.org/info/rfc5082>>.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", RFC 5918, DOI 10.17487/RFC5918, August 2010, <<http://www.rfc-editor.org/info/rfc5918>>.

## 12.2. Informative References

- [RFC4038] Shin, M-K., Ed., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, DOI 10.17487/RFC4038, March 2005, <<http://www.rfc-editor.org/info/rfc4038>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <<http://www.rfc-editor.org/info/rfc4301>>.
- [RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <<http://www.rfc-editor.org/info/rfc5340>>.
- [RFC5920] Fang, L., Ed., "Security Framework for MPLS and GMPLS Networks", RFC 5920, DOI 10.17487/RFC5920, July 2010, <<http://www.rfc-editor.org/info/rfc5920>>.
- [RFC6286] Chen, E. and J. Yuan, "Autonomous-System-Wide Unique BGP Identifier for BGP-4", RFC 6286, DOI 10.17487/RFC6286, June 2011, <<http://www.rfc-editor.org/info/rfc6286>>.
- [RFC6720] Pignataro, C. and R. Asati, "The Generalized TTL Security Mechanism (GTSM) for the Label Distribution Protocol (LDP)", RFC 6720, DOI 10.17487/RFC6720, August 2012, <<http://www.rfc-editor.org/info/rfc6720>>.
- [RFC7321] McGrew, D. and P. Hoffman, "Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 7321, DOI 10.17487/RFC7321, August 2014, <<http://www.rfc-editor.org/info/rfc7321>>.
- [RFC7439] George, W., Ed., and C. Pignataro, Ed., "Gap Analysis for Operating IPv6-Only MPLS Networks", RFC 7439, DOI 10.17487/RFC7439, January 2015, <<http://www.rfc-editor.org/info/rfc7439>>.
- [RFC7473] Raza, K. and S. Boutros, "Controlling State Advertisements of Non-negotiated LDP Applications", RFC 7473, DOI 10.17487/RFC7473, March 2015, <<http://www.rfc-editor.org/info/rfc7473>>.

## Appendix A. Additional Considerations

### A.1. LDPv6 and LDPv4 Interoperability Safety Net

It is not safe to assume that implementations compliant with RFC 5036 have supported the handling of an IPv6 address family (IPv6 FEC-label) in a Label Mapping message all along.

If a router upgraded per this specification advertised both IPv4 and IPv6 FECs in the same Label Mapping message, then an IPv4-only peer (not knowing how to process such a message) may abort processing the entire Label Mapping message (thereby discarding even the IPv4 FEC-labels), as per Section 3.4.1.1 of [RFC5036].

This would result in LDPv6 being somewhat undeployable in existing production networks.

Section 7 of this document provides a good safety net and makes LDPv6 incrementally deployable without making any such assumption on the routers' support for IPv6 FEC processing in current production networks.

### A.2. Accommodating Implementations Not Compliant with RFC 5036

It is not safe to assume that implementations have been [RFC5036] compliant in gracefully handling an IPv6 address family (IPv6 Address List TLV) in an Address message all along.

If a router upgraded per this specification advertised IPv6 addresses (with or without IPv4 addresses) in an Address message, then an IPv4-only peer (not knowing how to process such a message) may not follow Section 3.5.5.1 of [RFC5036] and may tear down the LDP session.

This would result in LDPv6 being somewhat undeployable in existing production networks.

Sections 6 and 7 of this document provide a good safety net and make LDPv6 incrementally deployable without making any such assumption on the routers' support for IPv6 FEC processing in current production networks.

#### A.3. Why prohibit IPv4-mapped IPv6 addresses in LDP?

Per discussion with the 6MAN and V6OPS working groups, the overwhelming consensus was to not promote IPv4-mapped IPv6 addresses appearing in the routing table, as well as in LDP (address and label) databases.

Also, [RFC4038], Section 4.2 suggests that IPv4-mapped IPv6-addressed packets should never appear on the wire.

#### A.4. Why a 32-bit value even for the IPv6 LDP Router Id?

The first four octets of the LDP Identifier, the 32-bit LSR Id (i.e., LDP router Id), identify the LSR and provide a globally unique value within the MPLS network, regardless of the address family used for the LDP session.

Please note that the 32-bit LSR Id value would not map to any IPv4 address in an IPv6-only LSR (i.e., Single-stack), nor would there be an expectation of it being IP routable or DNS resolvable. In IPv4 deployments, the LSR Id is typically derived from an IPv4 address, generally assigned to a loopback interface. In IPv6-only deployments, this 32-bit LSR Id must be derived by some other means that guarantees global uniqueness within the MPLS network, similar to that of the BGP Identifier [RFC6286] and the OSPF router Id [RFC5340].

This document reserves 0.0.0.0 as the LSR Id and prohibits its usage with IPv6, in line with the OSPF router Id in OSPF version 3 [RFC5340].

## Acknowledgments

We acknowledge the authors of [RFC5036], since some text in this document is borrowed from [RFC5036].

Thanks to Bob Thomas for providing critical feedback to improve this document early on.

Many thanks to Eric Rosen, Lizhong Jin, Bin Mo, Mach Chen, Shane Amante, Pranjal Dutta, Mustapha Aissaoui, Matthew Bocci, Mark Tinka, Tom Petch, Kishore Tiruveedhula, Manoj Dutta, Vividh Siddha, Qin Wu, Simon Perreault, Brian E. Carpenter, Santosh Esale, Danial Johari, and Loa Andersson for thoroughly reviewing this document and for providing insightful comments and multiple improvements.

## Contributors

The following individuals contributed to this document:

Nagendra Kumar  
Cisco Systems, Inc.  
7200 Kit Creek Road  
Research Triangle Park, NC 27709, United States  
EMail: [naikumar@cisco.com](mailto:naikumar@cisco.com)

Andre Pelletier  
Cisco Systems, Inc.  
2000 Innovation Drive  
Kanata, ON K2K-3E8, Canada  
EMail: [apelletti@cisco.com](mailto:apelletti@cisco.com)

## Authors' Addresses

Rajiv Asati  
Cisco Systems, Inc.  
7025 Kit Creek Road  
Research Triangle Park, NC 27709-4987  
United States

EMail: rajiva@cisco.com

Carlos Pignataro  
Cisco Systems, Inc.  
7200 Kit Creek Road  
Research Triangle Park, NC 27709-4987  
United States

EMail: cpignata@cisco.com

Kamran Raza  
Cisco Systems, Inc.  
2000 Innovation Drive  
Ottawa, ON K2K-3E8  
Canada

EMail: skraza@cisco.com

Vishwas Manral  
Ionos Networks  
4100 Moorpark Ave., Ste. #122  
San Jose, CA 95117  
United States  
Phone: +1 408 447 1497

EMail: vishwas@ionosnetworks.com

Rajiv Papneja  
Huawei Technologies  
2330 Central Expressway  
Santa Clara, CA 95050  
United States  
Phone: +1 571 926 8593

EMail: rajiv.papneja@huawei.com