MPLS Working Group   Rajiv Asati
Internet Draft   Carlos Pignataro
Updates: 5036, 6720 (if approved)   Kamran Raza
Intended status: Standards Track   Cisco
Expires: August 2015

                                                      Vishwas Manral
                                              Hewlett-Packard, Inc

                                                      Rajiv Papneja
                                                             Huawei

                                                  February 26, 2015

                           Updates to LDP for IPv6
                        draft-ietf-mpls-ldp-ipv6-17

Abstract

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

Table of Contents

1. Introduction

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

However, RFC5036 specification has the following deficiency (or lacks details) in regards to IPv6 usage (with or without IPv4):

1) LSP Mapping: No rule for mapping a particular packet to a particular LSP that has an Address Prefix FEC element containing 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 or/and IPv6 Address bindings over an LDP session

6) LDP Label Distribution: No rule for advertising IPv4 or/and IPv6
   FEC-label bindings over an LDP session, and for handling the co-
   existence 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 TTL Security: No rule for 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 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 3 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 (figure 2 and figure 3):


```
              R1-----------------R2
                    IPv4+IPv6
```

Figure 1 LSRs connected via a Dual-stack Interface


```
                  IPv4
              R1================R2
                  IPv6
```

Figure 2 LSRs connected via two Single-stack Interfaces

```
         R1-----------------R2---------------R3
              IPv4                   IPv6
```
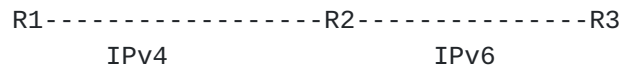
Figure 3 LSRs connected via a Single-stack Interface

Note that the topology scenario illustrated in figure 1 also covers
the case of a Single-stack LDP enabled interface (IPv4, say) being
converted to a Dual-stacked LDP enabled interface (by enabling IPv6
routing as well as IPv6 LDP), even though the 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:

```
                    IPv4
          R1------------------R2
                    IPv6
```
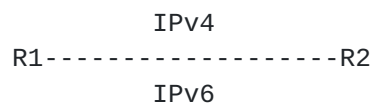
Figure 4 LSRs involving IPv4 and IPv6 address-families

1.2. Single-hop vs. Multi-hop LDP Peering

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

As a consequence, any LDP feature that relies on multi-hop LDP
peering session would not work with GTSM and will warrant
(statically or dynamically) disabling GTSM. Please see section 10.

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 [RFC4798]

LSPv6   - IPv6-signaled Label Switched Path [RFC4798]

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
3rd 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, but not for IPv6, since an IPv6
router may even have a /64 or /96 or /128 (or whatever prefix
length) address. Hence, that rule is updated to use IPv4 or IPv6
address 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 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 last paragraph of
Section 2.5.2 of [RFC5036] to be per address family and therefore
updates that sentence to the following:

> "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 (same LSR Id aka LDP Router Id as well as a common Label
space id) for both IPv4 and IPv6 address families.

5. Neighbor Discovery

   If Dual-stack LDP is enabled (e.g. 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 (e.g. LDP enabled in either 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 "all routers on this subnet" group
   multicast IP address.

   Interesting 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 e.g.
   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 Generalized TTL
   Security Mechanism (GTSM) 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
(e.g. 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. Maintaining Hello Adjacencies

In case of 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 Rx-side Link Hello
adjacency only for the address family that has been used for the
establishment of the LDP session (whether 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

The forthcoming sub-section 6.1 discusses the LDP consideration for
IPv6 and/or Dual-stacking in the context of session establishment,
whereas sub-section 6.2 discusses the LDP consideration 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 an optional
   transport address object (TLV) in LDP Hello message to convey the
   transport (IP) address, however, it does not specify the behavior of
   LDP if both IPv4 and IPv6 transport address objects (TLV) 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 the session establishment.

   This document specifies that:

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

   3. An LSR MUST send separate Hello messages (each containing
      either 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 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 IPv6
      transport address optional object of outgoing Link Hellos, if
      it had to choose between global unicast IPv6 address and
      unique-local or link-local IPv6 address.

   6. A Single-stack LSR MUST establish either 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
   they already have an LDPoIPv4 or LDPoIPv6 session (for the same
   LDP Identifier) established.

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

8. A Dual-stack LSR SHOULD prefer establishing an LDPoIPv6 session
   (instead of 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 case of Dual-
   stack LDP being 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. Determining Transport connection Roles

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

To ensure deterministic transport connection (active/passive) role
in 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 Dual-
stack capability TLV, as defined below:


```
 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 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|1|0|  Dual-stack capability  |        Length                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|TR     |     Reserved        |    MBZ                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

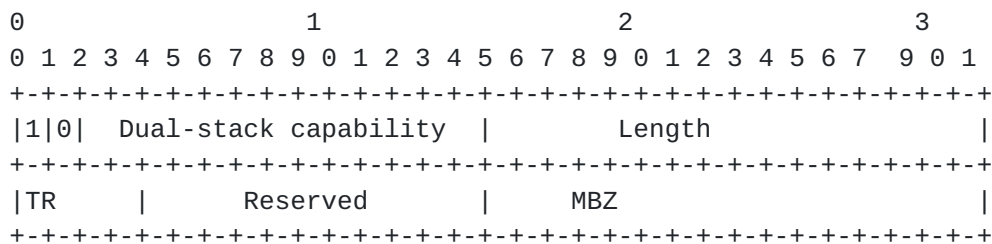   Figure 5 Dual-stack capability TLV

Where:

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

Dual-stack capability: TLV code point (to be assigned by IANA).

TR,   Transport Connection Preference.

This document defines the following 2 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. LSR supporting Dual-stack LDP for a peer)
MUST include "Dual-stack capability" TLV in all of its LDP Hellos,
and MUST set the "TR" field to announce its preference for either
LDPoIPv4 or LDPoIPv6 transport connection for that peer. The default
preference is LDPoIPv6.

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

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

   If LDP session was already in place, then LSR MUST send a fatal
   Notification message with status code [Transport Connection
   mismatch, IANA allocation TBD] and reset the session.

2. If "Dual-stack capability" TLV is present, and remote
   preference matches with the local preference, then:

   a) If TR=0100 (LDPoIPv4), then determine the active/passive
      roles for TCP connection using 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 TCP connection by using IPv6 transport address
          as defined in section 2.5.2 of RFC 5036.

   3. If "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 session from that neighbor, then the
          neighbor is deemed as a non-compliant 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 Non-Compliance', IANA allocation TBD).

       b) Only IPv6 hellos are received, then the neighbor is deemed
          as an IPv6-only LSR (supporting Single-stack LDP) and
          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 session from that neighbor, then the
          neighbor is deemed as a non-compliant 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 Non-Compliance', IANA allocation TBD).

       c) Both IPv4 and IPv6 hellos are received, then the neighbor
          is deemed as a non-compliant Dual-stack neighbor, and is
          not allowed to have any LDP session. A Notification
          message should be sent (with status code of 'Dual-Stack
          Non-Compliance', IANA allocation TBD).

  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 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 same TR value) or none do).

A Single-stack LSR do 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 (IPv4, say)
over another AFI (IPv6, say) 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 LDP Hello message, however, it will get used even in a
Single-stack IPv6 LDP networks and linger on forever, even though
Dual-stack will not. Hence, this alternative is discarded.

6.2. LDP Sessions Maintenance

This document specifies that two LSRs maintain a single LDP session
regardless of 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 (e.g. figure 1) on either LSR;
- an additional Single-stack or Dual-stack LDP enabled interface is
  added or removed between two LSRs (e.g. 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
a Single-stack LDP enabled interfaces on one LSR etc.), 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 establishing 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 LDPoIPv4 or LDPoIPV6
   session between them.

   However, there might be some legacy LSRs that are fully RFC 5036
   compliant for IPv4, but non-compliant for IPv6 (say, 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 Appendix 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 peer has been determined as a legacy LSR.

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

7.1. **Address Distribution**

   An LSR MUST NOT advertise (via ADDRESS message) any IPv4-mapped IPv6
   addresses (defined in section 2.5.5.2 of [RFC4291]), and 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 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. Is NOT able to 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. Is able to find the Dual-stack capability in the incoming IPv4
        (or IPv6) LDP Hello messages from that peer, then it MUST
        advertise (via ADDRESS message) its local IPv4 and IPv6
        addresses to that peer.

   3. Is NOT able to find the Dual-stack capability in the incoming
      IPv6 LDP Hello messages, then it MUST advertise (via 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 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 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 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. Is NOT able to 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 IPv6 address family was done).

   2. Is able to find the Dual-stack capability 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. Is NOT able to find the Dual-stack capability 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 [IPPWCap] document. 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

RFC5036 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 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 'IPv6 Addressing Architecture [RFC4291]' allows a link-local IPv6 address to be used on more than one links.

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 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 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 recommends enabling Generalized TTL Security Mechanism
(GTSM) for LDP, as specified in [RFC6720], for the LDP/TCP transport
connection over IPv6 (i.e. LDPoIPv6). The GTSM inclusion is intended
to automatically protect IPv6 LDP peering session from off-link
attacks.

[RFC6720] allows for the implementation to statically
(configuration) and/or dynamically override the default behavior
(enable/disable GTSM) on a per-peer basis. Such a configuration an
option could be set on either LSR (since GTSM negotiation would
ultimately disable GTSM between 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 requires a code point from the
TLV Type Name Space. IANA is requested to allocated a code point
from the IETF Consensus range 0x0700-0x07ff for the 'Dual-Stack
capability' TLV.

The 'Transport Connection Mismatch' status code requires a code
point from the Status Code Name Space. IANA is requested to allocate
a code point from the IETF Consensus range and mark the E bit column
with a '1'.

The 'Dual-Stack Non-Compliance' status code requires a code point
from the Status Code Name Space. IANA is requested to allocate a
code point from the IETF Consensus range and mark 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 the [RFC5036] are relevant
for this document as well, this document reduces the chances of off-
link attacks when using 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. 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 providing
insightful comments and multiple improvements.

This document was prepared using 2-Word-v2.0.template.dot.

13. Additional Contributors

The following individuals contributed to this document:

Kamran Raza
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, ON K2K-3E8, Canada
Email: skraza@cisco.com

Nagendra Kumar
Cisco Systems, Inc.
SEZ Unit, Cessna Business Park,
Bangalore, KT, India
Email: naikumar@cisco.com

Andre Pelletier
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, ON K2K-3E8, Canada
Email: apelleti@cisco.com

14. References

14.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
          Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC4291] Hinden, R. and S. Deering, "Internet Protocol Version 6
          (IPv6) Addressing Architecture", RFC 4291, February 2006.

[RFC5036] Andersson, L., Minei, I., and Thomas, B., "LDP
          Specification", RFC 5036, October 2007.

[RFC5082] Pignataro, C., Gill, V., Heasley, J., Meyer, D., and
          Savola, P., "The Generalized TTL Security Mechanism
          (GTSM)", RFC 5082, October 2007.

[RFC5918] Asati, R., Minei, I., and Thomas, B., "Label Distribution
          Protocol (LDP) 'Typed Wildcard Forward Equivalence Class
          (FEC)", RFC 5918, October 2010.

14.2. Informative References

[RFC4301] Kent, S. and K. Seo, "Security Architecture and Internet
          Protocol", RFC 4301, December 2005.

[RFC7321] Manral, V., "Cryptographic Algorithm Implementation
          Requirements for Encapsulating Security Payload (ESP) and
          Authentication Header (AH)", RFC 7321, April 2007.

[RFC5920] Fang, L., "Security Framework for MPLS and GMPLS
          Networks", RFC 5920, July 2010.

[RFC4798] De Clercq, et al., "Connecting IPv6 Islands over IPv4 MPLS
          Using IPv6 Provider Edge Routers (6PE)", RFC 4798,
          February 2007.

[IPPWCap] Raza, K., "LDP IP and PW Capability", draft-ietf-mpls-ldp-
          ip-pw-capability, October 2014.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF
          for IPv6", RFC 5340, July 2008.

   [RFC6286] E. Chen, and J. Yuan, "Autonomous-System-Wide Unique BGP
             Identifier for BGP-4", RFC 6286, June 2011.

   [RFC6720] R. Asati, and C. Pignataro, "The Generalized TTL Security
             Mechanism (GTSM) for the Label Distribution Protocol
             (LDP)", RFC 6720, August 2012.

   [RFC4038] M-K. Shin, Y-G. Hong, J. Hagino, P. Savola, and E. M.
             Castro, "Application Aspects of IPv6 Transition", RFC
             4038, March 2005.

   [RFC7439] W. George, and C. Pignataro, "Gap Analysis for Operating
             IPv6-Only MPLS Networks", RFC 7439, January 2015.

Appendix A.

A.1. LDPv6 and LDPv4 Interoperability Safety Net

It is not safe to assume that RFC5036 compliant implementations have supported handling IPv6 address family (IPv6 FEC label) in Label Mapping message all along.

If a router upgraded with 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 label FECs), as per the section 3.4.1.1 of RFC5036.

This would result in LDPv6 to be somewhat undeployable in existing production networks.

The change proposed in 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 Non-RFC5036-compliant implementations

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

If a router upgraded with this specification advertised IPv6 addresses (with or without IPv4 addresses) in 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 tear down the LDP session.

This would result in LDPv6 to be somewhat undeployable in existing production networks.

The changes proposed in section 6 and 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.3. Why prohibit IPv4-mapped IPv6 addresses in LDP

Per discussion with 6MAN and V6OPS working groups, the overwhelming
consensus was to not promote IPv4-mapped IPv6 addresses appear 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 32-bit value even for IPv6 LDP Router ID

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

Please note that 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, nor 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 BGP Identifier [RFC6286] and OSPF router ID [RFC5340].

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

Author's Addresses

   Rajiv Asati
   Cisco Systems, Inc.
   7025 Kit Creek Road
   Research Triangle Park, NC 27709-4987
   Email: rajiva@cisco.com

   Vishwas Manral
   Hewlet-Packard, Inc.
   19111 Pruneridge Ave., Cupertino, CA, 95014
   Phone: 408-447-1497
   Email: vishwas@ionosnetworks.com

   Kamran Raza
   Cisco Systems, Inc.,
   2000 Innovation Drive,
   Ottawa, ON K2K-3E8, Canada.
   E-mail: skraza@cisco.com

   Rajiv Papneja
   Huawei Technologies
   2330 Central Expressway
   Santa Clara, CA 95050
   Phone: +1 571 926 8593
   EMail: rajiv.papneja@huawei.com

   Carlos Pignataro
   Cisco Systems, Inc.
   7200 Kit Creek Road
   Research Triangle Park, NC 27709-4987
   Email: cpignata@cisco.com