MPLS Working Group
Internet Draft
Updates: 5036 (if approved)
Intended status: Standards Track
Expires: January 2015

Rajiv Asati
Cisco

Vishwas Manral
Hewlett-Packard, Inc.

Rajiv Papneja
Huawei

Carlos Pignataro
Cisco

July 3, 2014

Updates to LDP for IPv6
draft-ietf-mpls-ldp-ipv6-13

Status of this Memo

Copyright Notice

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.

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 (with or without IPv4 co-existence)

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 FEC-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).

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 interfaces (figure 1), or one or more single-stack 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 interface (IPv4, say) being converted to a dual-stacked interface by enabling IPv6 routing as well as IPv6 LDP, even though the IPv4 LDP 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 interface (IPv6 routing and IPv6 LDP), even though the IPv4 LDP session may already be established between the LSRs over the existing interface(s).

This document also addresses the scenario in which the LSRs do 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 session

LDPoIPv6 - LDP over IPv6 transport session

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

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, it is reasonable to say IPv4 or IPv6 address instead of /32 or /128 addresses as shown below in the updated rule:

> "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 on a dual-stack LSR.

5. Neighbor Discovery

If an LSR is enabled with dual-stack LDP (e.g. LDP enabled in both IPv6 and IPv4 address families), then the LSR MUST advertise both IPv6 and IPv4 LDP Link or targeted Hellos and include the same LDP Identifier (assuming per-platform label space usage) in them.

If an LSR is enabled with single-stack LDP (e.g. LDP enabled in either IPv6 or IPv4 address family), then the LSR MUST advertise 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 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 send 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.

An implementation MUST send IPv6 LDP link Hellos before sending IPv4 LDP Link Hellos on a dual-stack interface.

5.1.1. Maintaining Hello Adjacencies

In case of dual-stack LDP interface (e.g. LDP enabled in both IPv6 and IPv4 address families), 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 for the address family that has been used for the establishment of the LDP session (either IPv4 or IPv6).

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 pre-configured (unicast) destination IPv6 address, 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 peer discovery has completed (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 there were both IPv4 and IPv6 Hello adjacencies.

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 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.

3. An LSR MUST send separate Hello messages (each containing either IPv4 or IPv6 transport address optional object) for each IP address family, if LDP was enabled for both IP address families.

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 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. An LSR SHOULD NOT create (or honor the request for creating) a TCP connection for a new LDP session with a remote LSR, if they already have an LDP session (for the same LDP Identifier) established over whatever IP version transport.

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

7. An LSR SHOULD prefer the LDP/TCP connection over IPv6 for a new LDP session with a remote LSR, if it is able to determine the

IPv6 presence (e.g. IPv6 Hello adjacency), by following the 'transport connection role' determination logic in section 6.1.1.

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 (IPv4 or IPv6) LDP, but not for a dual-stack (IPv4 and IPv6) 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 for dual-stack LDP peering, this document specifies that the LSR convey its transport connection preference in every LDP Hello message. A new optional parameter, encoded as a TLV, (section 3.5.2 of RFC5036) is defined as follows (for Hello Message):

```
 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|   IPv4orIPv6 Preference |         Length                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|TR |  Reserved               |      MBZ                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5 IPv4 or IPv6 Transport Preference TLV

Where:

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

IPv4orIPv6 Preference: TLV code point for IPv4 or IPv6 Preference (to be assigned by IANA).

TR, Transport Preference

00: IPv4

01: IPv6 (default value)

Reserved

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

A dual-stack LDP enabled LSR (capable of supporting both IPv4 and IPv6 transports for LDP) MUST include "IPv4orIPv6 Transport Preference" optional parameter in all of its LDP Hellos, and MUST set the "TR" field to announce its preference for either IPv4 or IPv6 transport connection. The default preference is IPv6.

Upon receiving the hello message with this TLV, a dual-stack capable receiving LSR MUST do the following:

1. If it understands the TLV, and if neighbor's preference does not match with the local preference, then it discards the hello (and no adjacency is formed) and logs an error.

2. If it understands the TLV, and if neighbor's preference matches with the local preference, then:

   a) If TR=0 (IPv4), 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=1 (IPv6), then determine the active/passive roles for TCP connection by comparing the LSR Id part of the LDP Identifiers of LSRs.

      The LSR with higher LSR Id MUST assume the active role and other LSR MUST assume the passive role for the IPv6 TCP connection.

3. If it does not understand the TLV, then it MUST silently discard this TLV and process the rest of the Hello message.

If an LSR receives the hello message without the "IPv4orIPv6 Transport Preference" TLV, then it MUST proceed with session establishment using single-stack rules, as per section 2.5.2 of RFC 5036.

An LSR MUST convey the same transport connection preference ("TR" field) 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.

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 Capability TLV could be a new Flag value in LDP Hello message, however, it will get used even in a single-stack IPv6 scenarios 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).

The procedures defined in section 6.1 SHOULD result in preferring LDPoIPv6 session only after the loss of an existing LDP session (because of link failure, node failure, reboot etc.).

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) SHOULD be reset. Otherwise, the LDP session SHOULD stay intact.

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

7. Address Distribution

If an LSR is enabled with dual-stack LDP (i.e. LDP in both IPv4 and IPv6 address families) for any (discovered or targeted) peer, then it MUST advertise (via ADDRESS message) its local IPv4 and IPv6 addresses to that peer by default, independent of the transport connection (address family) used for that peering.

If an LSR, compliant with this specification, is enabled with single-stack LDP (i.e. LDP in either IPv6 or IPv4 address family) for any (discovered or targeted) peer, then it MUST advertise (via ADDRESS message) its local IP addresses as per the enabled address family by default, and SHOULD accept a received Address message containing both IPv4 and IPv6 addresses.

8. 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.

Additionally, to ensure backward compatibility (and interoperability with IPv4-only LDP implementations) in light of section 3.4.1.1 of RFC5036, as rationalized in the Appendix section A.1 later, this document specifies that -

1. An LSR MUST NOT send a label mapping message with a FEC TLV containing two or more Prefix FEC Elements of different address families. This means that a FEC TLV in the label mapping message must contain all the Prefix FEC Elements belonging to IPv6 address family or IPv4 address family, but not both.

If an LSR is enabled with dual-stack LDP (i.e. LDP in both IPv4 and IPv6 address families) for any peer, then it MUST advertise the FEC-Label bindings for both IPv4 and IPv6 address families to that peer. However, an LSR MAY 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.

If an LSR is configured to move an interface or peer from single-stack (IPv6 or IPv4 address family) to dual-stack LDP (IPv6 and IPv4

address families), then an LSR SHOULD use Typed Wildcard FEC procedures [RFC5918] to request the FEC-label bindings for the enabled address family. This helps to relearn the FEC-label bindings that may have been discarded before without resetting the peering.

9. 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 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.

10. 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. Suffice to say that such 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].

11. IANA Considerations

This document defines a new optional parameter for the LDP Hello Message. The type code needs to be assigned by IANA.

12. 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 [RFC4835] as well as [RFC5920].

13. 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, 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.

14. 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

15. References

15.1. 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.

15.2. Informative References

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

[RFC4835] Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4835, 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, June 2011.

[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.

Appendix A.

A.1. LDPv6 and LDPv4 Interoperability Safety Net

It is naive to assume that RFC5036 compliant implementations have supported IPv6 address family (IPv6 FEC processing, in particular) in label advertisement all along. And if that assumption turned out to be not true, then section 3.4.1.1 of RFC5036 would cause LSRs to abort processing the entire label mapping message and generate an error.

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. 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].

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.

Author's Addresses

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

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

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

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