MPLS Working Group
Internet Draft
Intended status: Standards Track
Expires: April 2011

Vishwas Manral
IPInfusion Inc.

Rajiv Papneja
Isocore

Rajiv Asati
Cisco Systems

October 13, 2010

Updates to LDP for IPv6
draft-manral-mpls-ldp-ipv6-04

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.

Table of Contents

1. Introduction

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

However, RFC5036 specification has the following deficiencies in regards to IPv6 usage:

1) LSP mapping: No rule defined 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 multicast address (with or without IPv4 co-existence)

4) LDP Session establishment: No prescription for handling both IPv4 and IPv6 transport address optional objects in a Hello message, and subsequently two IPv4 and IPv6 transport connections.

This document addresses the above deficiencies by specifying the desired behavior.

Note that this document updates RFC5036.

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].

LDP - Label Distribution Protocol

FEC - Forwarding Equivalence Class

TLV - Type Length Value

LSR - Label Switch Router

LSP - Label Switched Path

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."

Suffice to say, this rule is correct for IPv4, but not for IPv6, since an IPv6 router may not have any /32 address.

This document proposes to modify this rule by also including a /128 address (for IPv6). In fact, it should be reasonable to just 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

Section 2.2.2 of [RFC5036] specifies formulating at least one LDP Identifier, however, it doesn't provide any consideration in case of IPv6 (with or without dual-stacking).

This document preserves the usage of 32-bit LSR Id on an IPv6 only LSR and allows the usage of a common LDP identifier i.e. same LSR-Id and same Label space id for IPv4 and IPv6 on a dual-stack LSR. This rightly enables the per-platform label space to be shared between IPv4 and IPv6.

> Editor's note: The possible conflict with last paragraph of section 2.5.2 of RFC5036 needs to be addressed or clarified.

Additionally, this document reserves 0.0.0.0 as the LSR-Id, and prohibits its usage.

5. Peer Discovery

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 sends LDP Link Hellos destined to "all routers on this subnet" group multicast IP address.

Interesting enough, per [IANA-IPv6] [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 for IPv6 LDP Link Hellos. An LDP Hello packet received on any of the other addresses should be dropped. Also, the LDP Link Hello packets must have their IPv6 Hop Limit set to 1.

More importantly, if an interface is a dual-stack interface (e.g. enabled with both IPv4 and IPv6 LDP), then the LSR must periodically send both IPv4 and IPv6 LDP Link Hellos and must separately maintain the Hello adjacency for IPv4 and IPv6. This ensures LDP peerings on a multi-access interface (even if there are IPv4-only, IPv6-only and dual-stack routers). Needless to say, the IPv4 and IPv6 LDP Link Hellos must carry the same LDP identifier (assuming per-platform label space usage).

5.2. Extended Discovery Mechanism

Suffice to say, the extended discovery mechanism (defined in section 2.4.2 of [RFC5036]) doesn't require any additional IPv6 specific consideration, since the targeted LDP Hellos are sent to a pre-configured destination IP address.

6. LDP Session Establishment

Section 2.5.1 of [RFC5036] defines a two-step process for LDP session establishment:

1. Transport connection establishment
2. Session initialization

Next two sections discuss the LDP consideration for IPv6 and/or dual-stacking.

6.1. Transport connection establishment

Section 2.5.2 of [RFC5036] specifies the use of an optional transport address object (TLV) in LDP Link Hello message, however, it does not specify the behavior of LDP in case of both IPv4 and IPv6 transport address objects (TLV) are sent in a Hello message. Additionally, it does not specify whether both IPv4 and IPv6 transport connections should be allowed, if there were Hello adjacencies for both IPv4 and IPv6.

This document specifies that:

- An LSR should not send the Hello containing both IPv4 and IPv6 transport address optional objects. In other words, there would be at most one optional Transport Address object in a Hello message. An LSR should include only the transport address whose address family is the same as that of the IP packet carrying Hello.

- An LSR should accept the Hello message that contains both IPv4 and IPv6 transport address optional objects, but use only the transport address whose address family is the same as that of the IP packet carrying Hello.

- 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 label spaces) established using whatever IP version. This means that only one transport connection is established, even if there are two Hello adjacencies (one for IPv4 and another for IPv6), as highlighted in the last paragraph of section 2.5.2.

- An LSR should close the lagging TCP connection for a new LDP session with a remote LSR, if they attempted two TCP connections using IPv4 and IPv6 transports simultaneously.

6.2. Session initialization

No additional consideration needed.

7. IANA Considerations

None.

8. Security Considerations

The extensions defined in this document only clarify the behavior of LDP, they do not define any new protocol procedures. All the security issues relevant for the [RFC5036] are relevant for this document as well.

Moreover, this document allows the use of IPsec [RFC4301] for IPv6 protection, hence, LDP can benefit from the additional security as specified in [RFC4835].

9. Acknowledgments

A lot of the text in this document is borrowed from [RFC5036]. The authors of the document are acknowledged. The authors also aknowledge the help of Manoj Dutta and Vividh Siddha. Thanks to Bob Thomas for providing critical feedback to improve this document early on.

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

10. References

10.1. Normative References

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

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

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC4291] Hinden, R. and S. Deering, "Internet Protocol Version 6 (IPv6) Addressing Architecture", RFC 3513, April 2003.

10.2. 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.

[IANA-IPv6] http://www.iana.org/assignments/ipv6-address-space.

Author's Addresses

Vishwas Manral
IP Infusion Inc.,
Bamankhola, Bansgali,
Almora, Uttarakhand 263601
Email: vishwas@ipinfusion.com

Rajiv Papneja
ISOCORE
12359 Sunrise Valley Dr, STE 100
Reston, VA 20190
Email: rpapneja@isocore.com

Rajiv Asati
Cisco Systems,
7025-6 Kit Creek Rd, RTP, NC, 27709-4987
Email: rajiva@cisco.com