MPLS Working Group   Vishwas Manral
Internet Draft   IPInfusion Inc.
Updates: 5036 (if approved)
Intended status: Standards Track   Rajiv Papneja
Expires: September 2011   Isocore

Rajiv Asati
Cisco

Carlos Pignataro
Cisco

May 17, 2011

Updates to LDP for IPv6
draft-ietf-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, 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 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 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 Label exchange: No rule for advertising IPv4 or/and IPv6 FEC-label bindings over an LDP session

6) LDP TTL security: No rule for built-in Generalized TTL Security Mechanism (GTSM) in LDP

This document addresses the above deficiencies by specifying the desired behavior/rules/details. It also clarifies the topology scenarios in section 1.1.

Note that this document updates RFC5036.

1.1. Topology Scenarios

The following scenarios in which the LSRs may be inter-connected via one or more dual-stack interfaces (figure 1), or two or more single-stack interfaces (figure 2 and figure 3) become quite relevant to consider while addressing the deficiencies highlighted in section 1.

```
                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

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 as well as IPv6 LDP, even though the IPv4 LDP session may already be established between the LSRs.

The topology scenario illustrated in figure 2 also covers the case of two routers getting connected via an additional single-stack interface (IPv6, say), even though the IPv4 LDP session may already be established between the LSRs over the existing interface.

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

LDPv4   - LDP for enabling IPv4 MPLS forwarding

LDPv6   - LDP for enabling IPv6 MPLS forwarding

LDPoIPv4 - LDP over IPv4 transport session

LDPoIPv6 - LDP over IPv6 transport session

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."

While the above rule mentions 'Address Prefix FEC', it is also applicable to 'Typed WildCard prefix FEC' [RFC5918].

Additionally, it is desirable that a packet is forwarded to an LSP of an egress router, only if LSP's address-family matches with that of the LDP hello adjacency on the next-hop interface.

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). Additionally, section 2.5.2 of [RFC5036] implicitly prohibits using the same label space for both IPv4 and IPv6 FEC-label bindings.

The first four octets of the LDP identifier, the 32-bit LSR Id, identify the LSR and is a globally unique value. This is regardless of the address family used for the LDP session. In other words, this document preserves the usage of 32-bit LSR Id on an IPv6 only LSR.

> 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 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.

The first sentence of last paragraph of Section 2.5.2 of [RFC5036] is qualified per address family and therefore updated to the following: "For a given address family over which a Hello is sent, and a given label space, an LSR MUST advertise the same transport address." This rightly enables the per-platform label space to be shared between IPv4 and IPv6.

In summary, this document not only allows the usage of a common LDP identifier i.e. same LSR-Id, but also the common Label space id for both IPv4 and IPv6 on a dual-stack LSR.

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 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 for IPv6 LDP Link Hellos. An LDP Hello packet received on any of the other addresses must be dropped.

Also, the 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 specified in Generalized TTL Security Mechanism (GTSM)[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. enabled with both IPv4 and IPv6 LDP), then the LSR must periodically send both IPv4 and IPv6 LDP Link Hellos (using the same LDP Identifier per section 4) and must separately maintain the Hello adjacency for IPv4 and IPv6 on that interface.

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 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-sections discuss the LDP consideration for IPv6 and/or dual-stacking in the context of session establishment and maintenance.

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 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 Hello adjacencies for both IPv4 and IPv6 whether over a single interface or multiple interfaces.

This document specifies that:

- An LSR should not send a Hello containing both IPv4 and IPv6 transport address optional objects. In other words, there should 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 must send separate Hellos (each containing either IPv4 or IPv6 transport address optional object) for each IP address-family, if LDP was enabled for both IP address-families.

- 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 should be established, even if there are two Hello adjacencies (one for IPv4 and another for IPv6). This is independent of whether the Hello Adjacencies are created over a single interface (scenarios 1 in section 1.1) or multiple interfaces (scenario 2 in section 1.1) between two LSRs.

- An LSR should prefer the LDP/TCP connection over IPv6 for a new LDP session with a remote LSR, if it has both IPv4 and IPv6 hello adjacencies for the same LDP Identifier (over a dual-stack interface, or two or more single-stack IPv4 and IPv6 interfaces). This applies to the section 2.5.2 of RFC5036.

- An LSR should prefer the LDP/TCP connection over IPv6 for a new LDP session with a remote LSR, if they attempted two TCP connections using IPv4 and IPv6 transport addresses simultaneously.

This document allows an implementation to provide a configuration to override the above stated preference from IPv6 to IPv4 on a per-peer basis. Suffice to say that, such preference must be set on both LSRs.

This document also specifies that the LDP/TCP transport connection over IPv6 must follow the GTSM procedures (Section 3 of [RFC5082]) by default, if the LDP/TCP transport connection is being established between the adjacent LSRs (using Basic Discovery, as described in section 5.1). This means that the IP Hop Limit field is set to 255 upon sending, and checked to be 255 upon receipt. The built-in inclusion of GTSM automatically protects IPv6 LDP peering session from off-link attacks.

6.2. Maintaining Hello Adjacencies

As outlined in section 2.5.5 of RFC5036, this draft suggests that if an LSR has a dual-stack interface, which is 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 on that interface.

This ensures successful labeled IPv4 and labeled IPv6 traffic forwarding on a dual-stacked interface, as well as successful LDP peering using the appropriate transport on a multi-access interface (even if there are IPv4-only, IPv6-only and dual-stack LSRs connected to that multi-access interface).

6.3. Maintaining LDP Sessions

Two LSRs maintain a single LDP session between them, as described in section 6.1, whether they are connected via a dual-stack LDP enabled interface or via two single-stack LDP enabled interfaces. This is also true when a single-stack interface is converted to a dual-stack interface, or when another interface is added between two LSRs.

On the other hand, if a dual-stack interface is converted to a single-stack interface (by disabling IPv4 or IPv6 routing), then the LDP session should be torn down ONLY if the disabled IP version was the same as that of the transport connection. 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 and RFC5036.

7. Label Distribution

This document specifies that an LSR should advertise and receive both IPv4 and IPv6 label bindings from and to the peer, only if it has valid IPv4 and IPv6 Hello Adjacencies for that peer, as specified in section 6.2.

This means that the LSR must not advertise any IPv6 label bindings to a peer over an IPv4 LDP session, if no IPv6 Hello Adjacency existed for that peer (and vice versa).

8. IANA Considerations

None.

9. 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].

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].

10. Acknowledgments

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

Thanks to Bob Thomas for providing critical feedback to improve this document early on. Thanks to Kamran Raza, Eric Rosen, Lizhong Jin, Bin Mo, Mach Chen, and Kishore Tiruveedhula for reviewing this document. The authors also acknowledge the help of Manoj Dutta and Vividh Siddha.

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

11. References

11.1. 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.

11.2. 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.

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

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

Author's Addresses

Vishwas Manral
IP Infusion Inc.,
1188 E. Arques Ave,
Sunnyvale, CA, 94089
Email: vishwas@ipinfusion.com

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