Network Internet-Draft Updates: <u>5036</u> (if approved) Intended status: Standards Track Expires: March 6, 2020

LDP behaviour on link-shut scenarios draft-aa-mpls-ldp-link-shut-00

Abstract

This document is intended as clarification of LDP behaviour in linkdown scenarios. Base LDP <u>RFC5036</u> lacks sufficient clarity on what an LDP enabled node should be doing when a link down event is received, and the only LDP adjacency for an LDP peer is over this link. Different vendors have handled this scenario differently, with some immediately resetting tcp session with neighbor and some waiting for igp recovergence instead of reacting directly to link events. With this document we intend to clarify the expected behaviour explicitly so that any interop issues can be avoided.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <u>https://datatracker.ietf.org/drafts/current/</u>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on March 6, 2020.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>https://trustee.ietf.org/license-info</u>) in effect on the date of publication of this document. Please review these documents

Anush & Anupkumar

Expires March 6, 2020

draft-aa-ldp-link-shut-00

carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

<u>1</u> .	Introduction										<u>2</u>
<u>2</u> .	Terminology										<u>2</u>
<u>3</u> .	Problem Description										<u>2</u>
<u>4</u> .	Solutions										<u>3</u>
<u>5</u> .	Security Considerations										<u>4</u>
<u>6</u> .	IANA Considerationss .										<u>4</u>
<u>7</u> .	Acknowledgments										<u>4</u>
<u>8</u> .	Normative References .										<u>4</u>
Autl	nors' Addresses										4

1. Introduction

[RFC5036] details LDP specification and procedures to be followed by LDP implementations. However, for some scenarios like link down, the rfc isn't particularly clear as to what an implementation is supposed to do. This could lead to interop issues when routers from different vendors are part of the network. More details are given in the problem description section below. A possible solution is also suggested in the subsequent sections.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in <u>BCP</u> <u>14</u> [<u>RFC2119</u>] [<u>RFC8174</u>] when, and only when, they appear in all capitals, as shown here.

- o LDP: Label distribution protocol.
- o GR: Graceful-restart.

<u>3</u>. Problem Description

Consider the following topology:

[Page 2]

E | A ---- B | | C ---- D

Figure 1: Example topology for LDP

All the nodes here are LDP enabled and also support GR. We have an lsp from C-E via C-A-E path(this is igp bestpath). IGP has also computed LFA backup for this primary-lsp via C-D-B-A-E path and we have LDP lfa backup as well (taking this path).

Now, lets bring down A-C link. Node A has detected link-shut event and since this link is the only adjacency to LDP-neighbor C, it resets the LDP session and sends shutdown to neighbour C.

At C, the link-down event is detected bit late and subsequently the IGP update is also delayed. Meanwhile, C has received shutdown from peer A, and it results in C flushing all labels received from A. Since the primary-label for C-E lsp is no longer available (from A), the lsp itself is deleted by LDP, as LDP can't be congruent with IGP. This LDP-lsp flap can in turn impact l3vpn/l2vpn traffic which are dependent on this LSP.

We can definitely reduce traffic-loss by running BFD and switching traffic to lfa backup in forwarding, but the intention above is to highlight that IGP updates and subsequently LDP updates would be asynchronous at nodes A and C, which may be more prominent if there are routers with different capabilities (and maybe from different vendors) in the network. So even if traffic has moved to lfa-backup lsp in forwarding, the primary-lsp itself could be deleted by the shutdown message (which is a fatal error).

4. Solutions

When a node has LDP adjacency to its neighbor (With GR [RFC3478] enabled on both the node and its neighbor) over a 'single' directly connected link and that link goes down, the node MAY reset the tcp session with neighbor. However, it MUST NOT send shutdown message, which flushes advertised labels at neighbor immediately.

The neighbor itself could have different backup mechanisms (ldp-lfa, rsvp-bypass etc) to ensure minimal traffic loss in forwarding for lsps having this node as active(primary)-path. Transmitting shutdown

[Page 3]

message immediately could result in neighbor prematurely deleting LSPs instead of letting IGP recoverge.

Another approach could be to avoid reacting immediately to link down events. Instead, let hello timeout bringdown the session and update LSP-paths as soon as IGP reconverges.

Both approaches can help to avoid traffic loss by accounting for asynchronous ordering of events in LDP-peering routers.

5. Security Considerations

The security considerations described in $\underline{\text{RFC5036}}$ apply to this document.

<u>6</u>. IANA Considerationss

7. Acknowledgments

.

8. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <<u>https://www.rfc-editor.org/info/rfc2119</u>>.
- [RFC3478] Leelanivas, M., Rekhter, Y., and R. Aggarwal, "Graceful Restart Mechanism for Label Distribution Protocol", <u>RFC 3478</u>, DOI 10.17487/RFC3478, February 2003, <<u>https://www.rfc-editor.org/info/rfc3478</u>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", <u>RFC 5036</u>, DOI 10.17487/RFC5036, October 2007, <<u>https://www.rfc-editor.org/info/rfc5036</u>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.

Authors' Addresses

Anush & Anupkumar Expires March 6, 2020

[Page 4]

Anush Mohan Ericsson Bangalore India

Internet-Draft

Email: anush.mohan@ericsson.com

Anupkumar T Ericsson Bangalore India

Email: anupkumar.t@ericsson.com

Anush & Anupkumar Expires March 6, 2020

[Page 5]