

MPLS Working Group
INTERNET-DRAFT
Intended Status: Proposed Standard

Santosh Esale
Raveendra Torvi
Juniper Networks

Luyuan Fang
Microsoft

Luay Jalil
Verizon

March 24, 2015

Fast Reroute for Node Protection in LDP-based LSPs
draft-esale-mpls-ldp-node-frr-00

Abstract

This document describes procedures to support node protection for (unicast) Label Switched Paths (LSPs) established by LDP ("Label Distribution Protocol"). In order to protect a node N, the Point of Local Repair (PLR) of N must discover the Merge Points (MPTs) of node N such that traffic can be redirected to them in case node N fails. Redirecting the traffic around the failed node N depends on existing point-to-point LSPs originated from the PLR to the MPTs while bypassing the protected node N. The procedures described in this document are topology independent in a sense that they provide node protection in any topology.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>

INTERNET DRAFT

<[draft-esale-mpls-ldp-node-frr](#)>

March 24, 2015

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

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

1.	Terminology	3
2.	Introduction	3
3.	Merge Point (MPT) Discovery	3
4.	Constructing Bypass LSPs	4
5.	Obtaining Label Mapping from MPT	5
6.	Forwarding Considerations	5
7.	Synergy with node protection in mLDP	6
8.	Security Considerations	6
9.	IANA Considerations	6
10.	Acknowledgements	6
11.	Normative References	6
12.	Informative References	6
	Authors' Addresses	6

INTERNET DRAFT

<[draft-esale-mpls-ldp-node-frr](#)>

March 24, 2015

1. Terminology

- PLR: Point of Local Repair (the LSR that redirects the traffic to one or more Merge Point LSRs).
- MPT: Merge Point. Any LSR on the LDP-signaled (multi-point to point) LSP, provided that the path from that LSR to the egress of that LSP is not affected by the failure of the protected node
- tLDP: Targeted LDP session.

2. Introduction

This document describes procedures to support node protection for (unicast) Label Switched Paths (LSPs) established by LDP ("Label Distribution Protocol") [[RFC5036](#)]. In order to protect a node N, the Point of Local Repair (PLR) of N must discover the Merge Points (MPTs) of node N such that traffic can be redirected to them in case node N fails. Redirecting the traffic around the failed node N depends on existing Point-to-Point (P2P) LSPs originating from the PLR LSR to the MPTs while bypassing node N. The procedures to setup these P2P LSPs are outside the scope of this document, but one option is to use RSVP-TE based techniques [[RFC3209](#)] to accomplish this. Finally, sending traffic from the PLR to the MPTs requires the PLR to obtain FEC-label mappings from the MPTs. The procedures described in this document relies on Targeted LDP (tLDP) session [[RFC5036](#)] for the PLR to obtain such mappings. The procedures for node protection described in this document fall into the category of local protection. The procedures described in this document apply to LSPs bound to either an IPv4 or IPv6 Address Prefix FEC. The procedures described in this document are topology independent in a sense that they provide node protection in any topology. Thus these procedures provide topology independent fast reroute.

3. Merge Point (MPT) Discovery

For a given LSP that traverses the PLR, the protected node N, and a

particular neighbor of the protected node, we'll refer to this neighbor as the "next next-hop". Note that from the PLR's perspective the protected node N is the next hop for the FEC associated with that LSP. Likewise, from the protected node's perspective the next next-hop is the next hop for that FEC. If for a given <LSP, PLR, N> triplet the next next-hop is in the same IGP area as the PLR, then that next next-hop acts as the MPT for that triplet. For a given LSP traversing a PLR and the node protected by the PLR, the PLR discovers the next next-hops (MPTs) that are in the same IGP area as the PLR from either its Traffic Engineering database or Link State database. The Traffic Engineering database or Link State database is populated by either ISIS or OSPF. The discovery of the next next-hop (depending

on an implementation) may not involve any additional SPF, above and beyond what would be needed by ISIS/OSPF anyway, as the next next-hop, just like the next-hop, is a by-product of SPF computation. If for a given <LSP, PLR, N> triplet the node protected by the PLR is an Area Border Router (ABR), then the PLR and the next next-hop may end up in different IGP areas (this could happen when an LSP traversing the PLR and the protected node does not terminate in the same IGP area as the PLR). In this situation the PLR may not be able to determine the next next-hop from either its Traffic Engineering database or Link State database, and thus may not be able to use the next next-hop as the MPT. In this scenario the PLR uses an "alternative" ABR as the MPT, where an alternative ABR is defined as follows. For a given LSP that traverses the PLR and the (protected) ABR, an alternative ABR is defined as any ABR that advertises into PLR's own IGP area reachability to the FEC associated with the LSP. Note that even if a PLR protects an ABR, for some of the LSPs traversing the PLR and the ABR, the next next-hops may be in the same IGP area as the PLR, in which case these next next-hops act as MPTs for these LSPs. Note that even if the protected node is not an ABR, if an LSP traversing the PLR and the protected node does not terminate in the same IGP area as the PLR, then for this LSP the PLR MAY use an alternative ABR (as defined above), rather than the next next-hop as the MPT.

[4. Constructing Bypass LSPs](#)

As we mentioned before, redirecting traffic around the failed node N depends on existing Point-to-Point (P2P) LSPs originating from the PLR to the MPTs while bypassing node N. We'll refer to these LSPs as

"bypass LSPs". While the procedures to setup these bypass LSPs are outside the scope of this document, this document assumes use of RSVP-TE LSPs [[RFC3209](#)] to accomplish this. Once a PLR that protects a given node N discovers the set of MPTs associated with itself and the protected node, (at the minimum) the PLR MUST (automatically) establish bypass LSPs to all these MPTs. The bypass LSPs MUST be established before the failure of the protected node. One could observe that if the protected node is not an ABR and the PLR does not use alternative ABR(s) as MPT(s), then the set of all the IGP neighbors of the protected node forms a superset of the MPTs. Thus it would be sufficient for the PLR to establish bypass LSPs with all the IGP neighbors of the protected node, even though some of these neighbors may not be MPTs for any of the LSPs traversing the PLR and the protected node. The bypass LSPs MUST avoid traversing the protected node, which means that the bypass LSPs are explicitly routed LSPs (of course, using RSVP-TE to establish bypass LSPs allows these LSPs to be explicitly routed). As a given router may act as an MPT for more than one LSP traversing the PLR, the protected node, and the MPT, the same bypass LSP will be used to protect all these LSPs.

5. Obtaining Label Mapping from MPT

As we mentioned before, sending traffic from the PLR to the MPTs requires the PLR to obtain FEC-label mappings from the MPTs. The solution described in this document relies on Targeted LDP (tLDP) session [[RFC5036](#)] for the PLR to obtain such mappings. Specifically, for a given PLR and the node protected by this PLR, at the minimum the PLR MUST (automatically) establish tLDP with all the MPTs associated with this PLR and the protected node. These tLDP sessions MUST be established before the failure of the protected node. One could observe that if the protected node is not an ABR and the PLR does not use alternative ABR(s) as MPT(s), then the set of all the IGP neighbors of the protected node forms a superset of the MPTs. Thus it would be sufficient for the PLR to (automatically) establish tLDP with all the IGP neighbors of the protected node that are in the same area as the PLR, even though some of these neighbors may not be MPTs for any of the LSPs traversing the PLR and the protected node. At the minimum for a given tLDP peer the PLR MUST obtain FEC-label mapping for the FEC(s) for which the peer acts as an MPT. The PLR MUST obtain this mapping before the failure of the protected node. To obtain this mapping for only these FECs (and no other FECs that the peer may maintain) the PLR MAY rely on the LDP Downstream on Demand

(DoD) procedures [[RFC5036](#)]. Otherwise, without relying on the DoD procedures, the PLR may end up receiving from a given tLDP peer FEC-label mappings for all the FECs maintained by the peer, even if the peer does not act as an MPT for some of these FECs. If the LDP DoD procedures are not used, then for the purpose of the procedures specified in this draft the only label mappings that SHOULD be exchanged are for the Address Prefix FECs whose PreLen value is either 32 (IPv4), or 128 (IPv6); label mappings for the Address Prefix FECs with any other PreLen value SHOULD NOT be exchanged. When a PLR has one or more ABRs acting as MPTs, the PLR MAY use the procedures specified in [[draft-ietf-mpls-app-aware-tldp](#)] to limit the set of FEC-label mappings received from non-ABR MPTs to only the mappings for the FECs associated with the LSPs that terminate in the PLR's own IGP area.

[6.](#) Forwarding Considerations

When a PLR detects failure of the protected node then rather than swapping an incoming label with a label that the PLR received from the protected node, the PLR swap the incoming label with the label that the PLR receives from the MPT, and then pushes the label associated with the bypass LSP to that MPT. To minimize micro-loop during the IGP global convergence PLR may continue to use the bypass LSP during network convergence by adding small delay before switching to a new path.

[7.](#) Synergy with node protection in mLDP

Both the bypass LSPs and tLDP sessions described in this document could also be used for the purpose of mLDP node protection, as described in [[draft-ietf-mpls-mldp-node-protection](#)].

[8.](#) Security Considerations

The same security considerations apply as those for the base LDP specification, as described in [[RFC5036](#)].

[9.](#) IANA Considerations

This document introduces no new IANA Considerations.

10. Acknowledgements

The authors would like to thank Yakov Rekhter, Hannes Gredler, Aman Kapoor, Minto Jeyananth and Eric Rosen for their contributions to this document.

11. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3209] D. Awduche, et al., "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC3209](#), Decembet 2001
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [[draft-ietf-mpls-app-aware-tldp](#)] Esale, S., et al., "Application-aware Targeted LDP", [draft-esale-mpls-app-aware-tldp](#), work in progress

12. Informative References

- [[draft-ietf-mpls-mldp-node-protection](#)], IJ. Wijnands, et al., "mLDP Node Protection", [draft-ietf-mpls-mldp-node-protection](#), work in progress

Authors' Addresses

Santosh Esale
Juniper Networks
Email: sesale@juniper.net

Esale, et al. Expires September 25, 2015 [Page 6]

INTERNET DRAFT <[draft-esale-mpls-ldp-node-frr](#)> March 24, 2015

Yakov Rekhter
Juniper Networks
Email: yakov@juniper.net

Raveendra Torvi
Juniper Networks
Email: rtorvi@juniper.net

Luyuan Fang
Microsoft
Email: lufang@microsoft.com

Luay Jalil
Verizon
Email: luay.jalil@verizon.com