Fast Reroute for Node Protection in LDP-based LSPs

draft-esale-mpls-ldp-node-frr-00

Raveendra Torvi (rtorvi@juniper.net)
Luay Jalil (luay.jalil@verizon.com)
Luyuan Fang (lufang@microsoft.com)
Santosh Esale (sesale@juniper.net)

IETF-93 (Prague)

speaker: Santosh Esale

Requirements

- Applicability/scope: LDP-signaled transport LSPs within single IGP (ISIS/OSPF) routing domain
- Local protection to minimize connectivity disruption
- Protection for both link and node failure
- No restrictions on the network topology – provide topology independent local protection
- Minimize additional provisioning/configuration required

Building blocks

- Point of Local Repair (PLR):
  - Node which notices failure of:
    - Downstream link/node
  - Begins forwarding traffic towards Merge Point using Bypass LSP when detects failure
- Merge Point (MPT)
  - Any router on LDP-signaled (multi-point to point) LSP, provided that the path from that router to the egress of that LSP is not affected by the failure of the protected link/node
- Bypass LSP
  - LSP from PLR to MPT bypassing the protected link/node
  - Established prior to failure
  - This presentation assumes use of RSVP-TE for establishing bypass LSPs
- Label mapping from MPT
  - PLR has to obtain label mapping from MPT
  - Label mapping obtained prior to failure
  - Once PLR detects failure, PLR swaps the incoming label with the label from MPT
    - Rather than with the label received from the next hop

Link Protection

Link Protection Building Blocks

- For a given LSP traversing a given (protected) link:
  - PLR: router at the upstream end of the link
    - With respect to the LSP
    - Acts as PLR for the downstream link
  - MPT: router at the downstream end of the link
    - With respect to the LSP
    - Next-hop from PLR's point of view
  - Bypass LSP: LSP created between the two routers at the end of the (protected) link
    - Bypasses the protected link
    - The same bypass LSP protects all LSPs traversing the protected link
  - Label mapping: the same as prior to link failure
    - Because MPT is the next hop

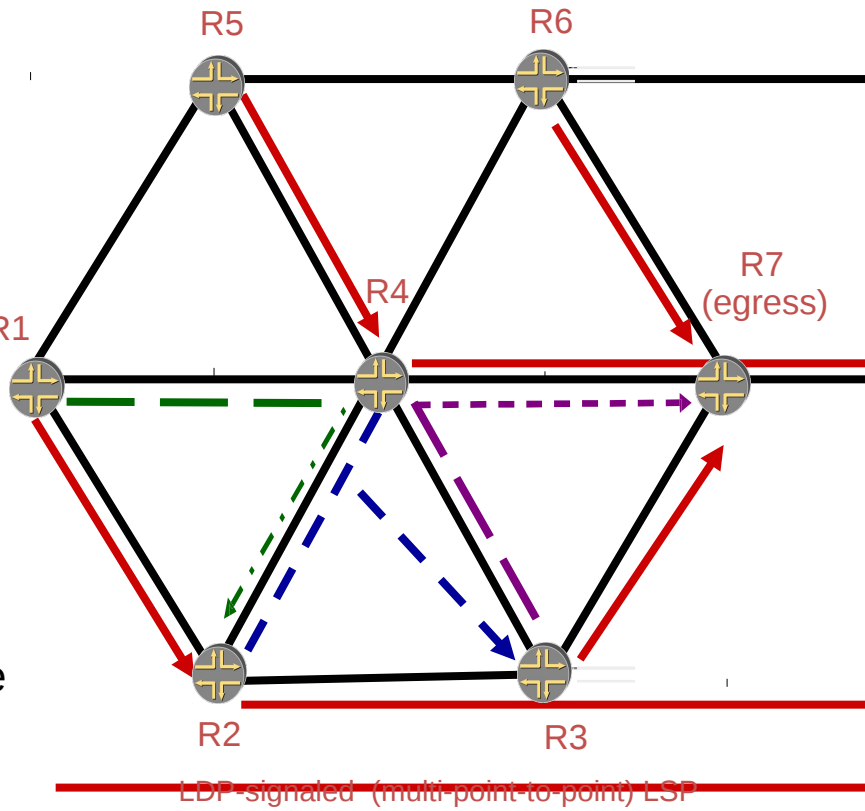

Protected link	PLR	MPT	Bypass LSP
R1-R2	R1	R2	<R1, R4, R2>
R2-R3	R2	R3	<R2, R4, R3>

Link Protection - Example

Label handling and data flow during steady state

Data flow after link failure

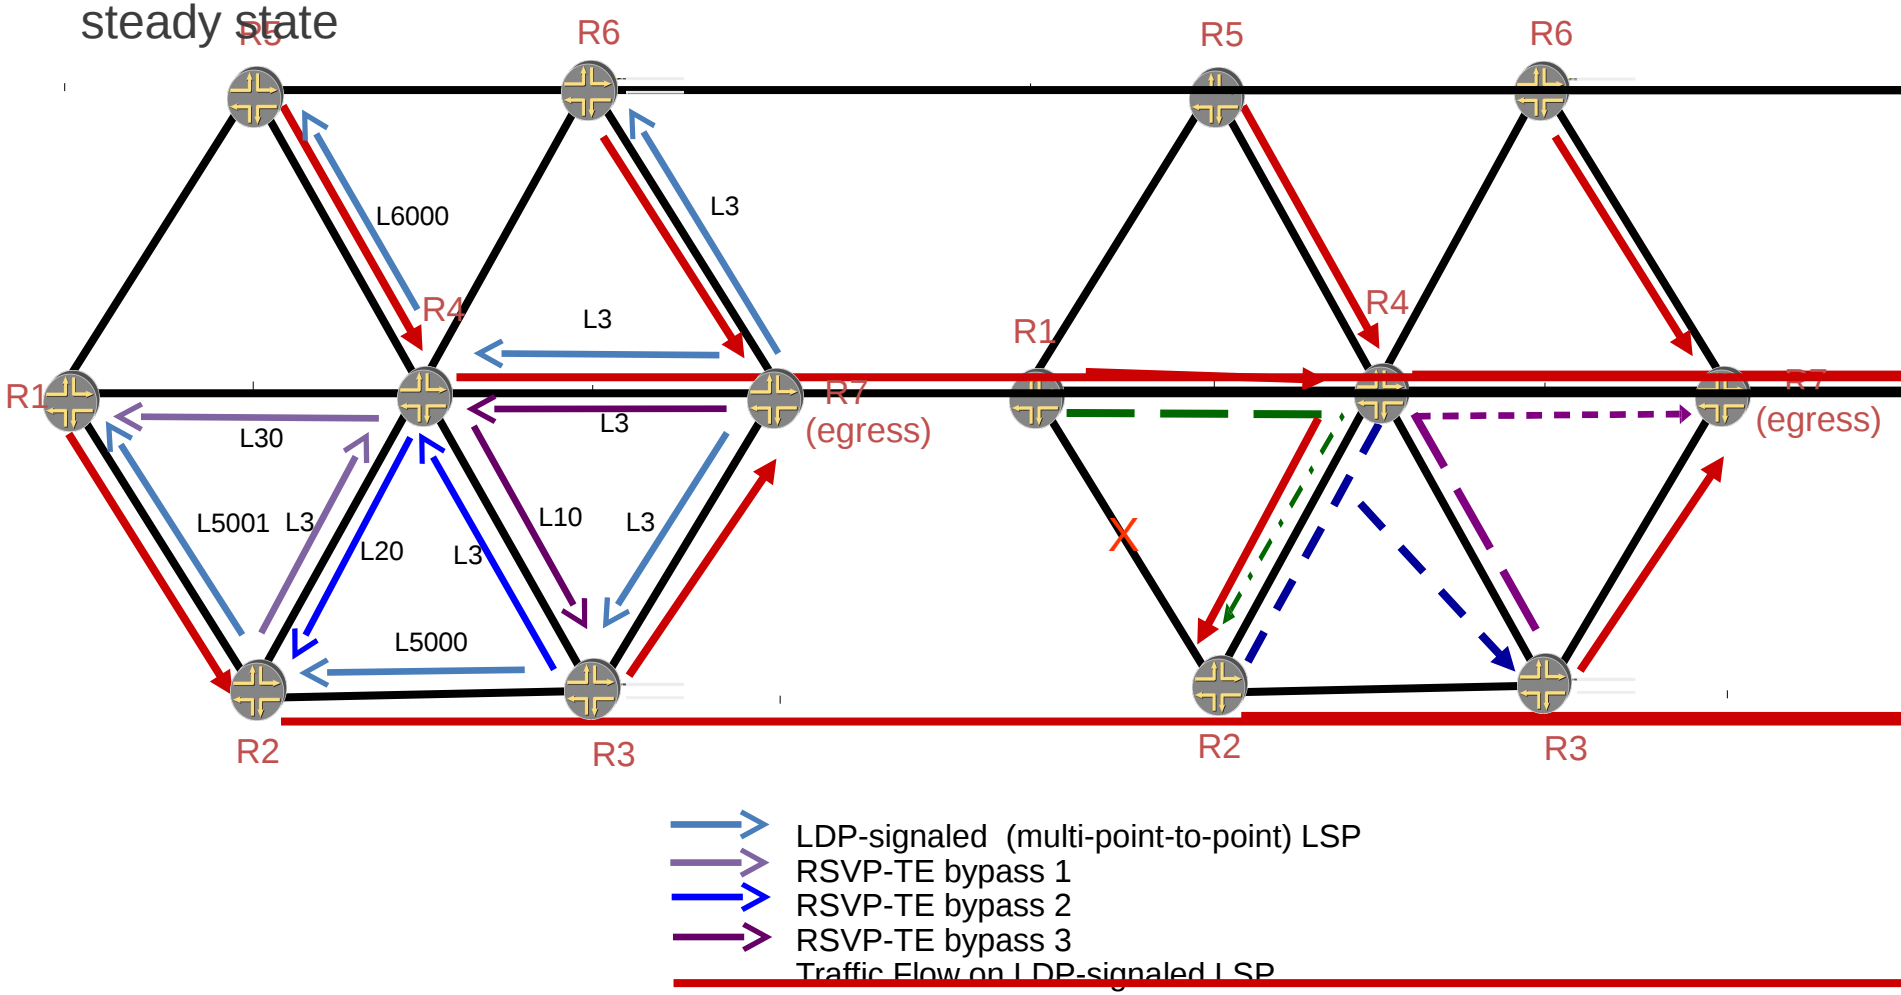

Node Protection

Node Protection Building Blocks

- For a given (multi-point to point) LSP traversing a given protected node:
  - PLR: router one hop upstream from the protected node
    - With respect to the LSP
    - Previous hop with respect to the protected node
  - MPT: Any router on the LSP, provided that the path from that router to the egress of the LSP is not affected by failure of the protected node
    - More on this in the next slides…
  - Bypass LSP: LSP created from PLR to MPT
    - Bypasses the protected node
    - The same bypass LSP is used to protect all LSPs traversing PLR, protected node, and MPT
  - Label mapping: obtained from MPT using Targeted LDP between PLR and MPT
    - The label from MPT may not be the same as the label from the next hop
    - Only labels for Address Prefix FECs with Prefix Lenght 32 (IPv4) or 128 (IPv6) should be exchanged
    - To acquire label mapping only for the FEC of this LSP PLR may use LDP Downstream on Demand
    - Same applies to every LSPs traversing PLR, protected node, and MPT

Node Protection – Determining MPT (1)

(protected node is **not** ABR)

Consider an LSP that traverses PLR, protected node, and particular neighbor of the protected node - we'll refer to this neighbor as the "*next next-hop*“

- From PLR’s perspective the protected node is the next hop for the FEC associated with that LSP
- From protected node’s perspective the next next-hop is the next hop for that FEC

When the protected node is not an Area Border Router (ABR), PLR can determine the next next-hop as a by-product of SPF required by ISIS/OSPF

- No additional SPF may be needed

When the protected node is not an ABR, PLR uses the next next-hop as MPT

- As path from the next next-hop to the egress is not affected by failure of the protected node

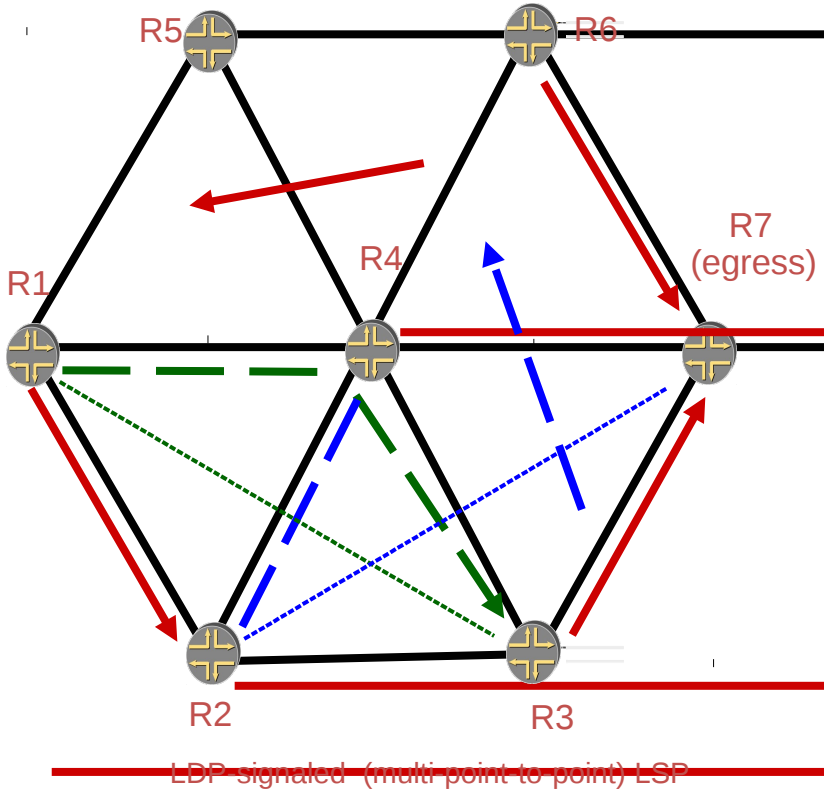

Protected node	PLR	MPT (next next-hop)	Bypass LSP
R2	R1	R3	<R1, R4, R3>

Node Protection – Determining MPT (2)

(protected node is ABR)

Consider an LSP that traverses PLR, protected node, and particular neighbor of the protected node - we'll refer to this neighbor as the "*next next-hop"*

When the protected node is an ABR, PLR may not be able to determine the next next-hop from its SPF

- As PLR and the next next-hop may end up in different IGP areas
- Yet in ISIS/OSPF scope of SPF is the IGP area of PLR

In this scenario PLR uses an "***alternative***" ABR as MPT

- For a given LSP that traverses PLR and 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

PLR discovers an alternative ABR from the IGP database

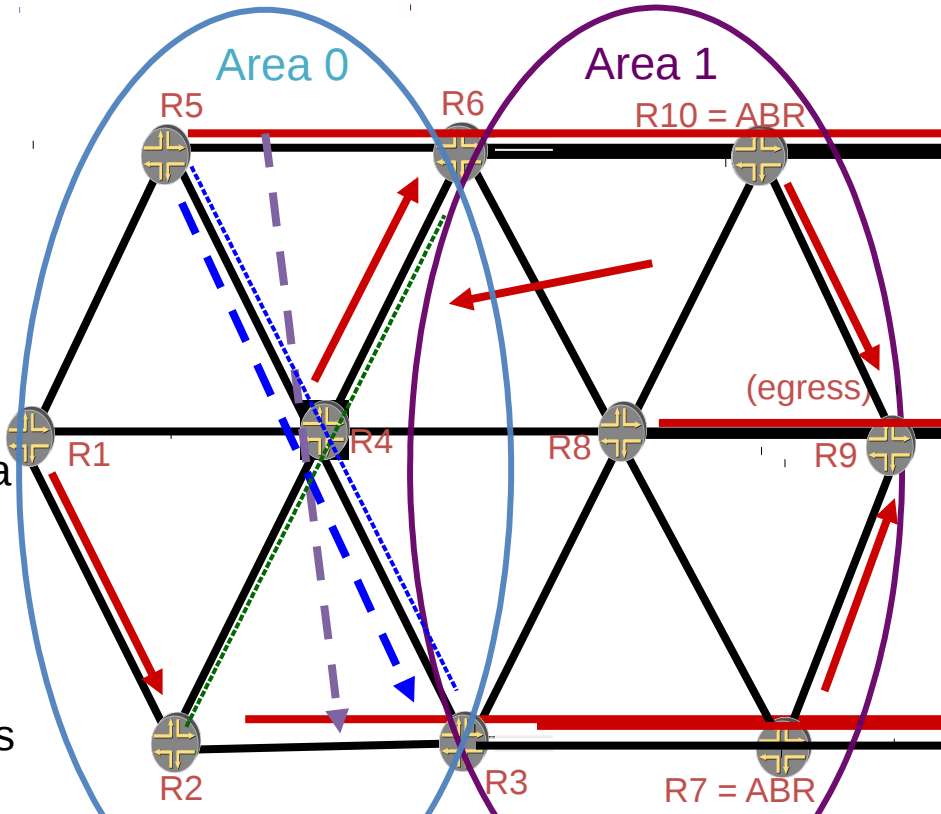

Protected node	PLR	MPT (Alternative ABR)	Bypass LSP
R7	R3	R10	<R3, R8, R10>

Node Protection - Example

(protected node is **not** ABR)

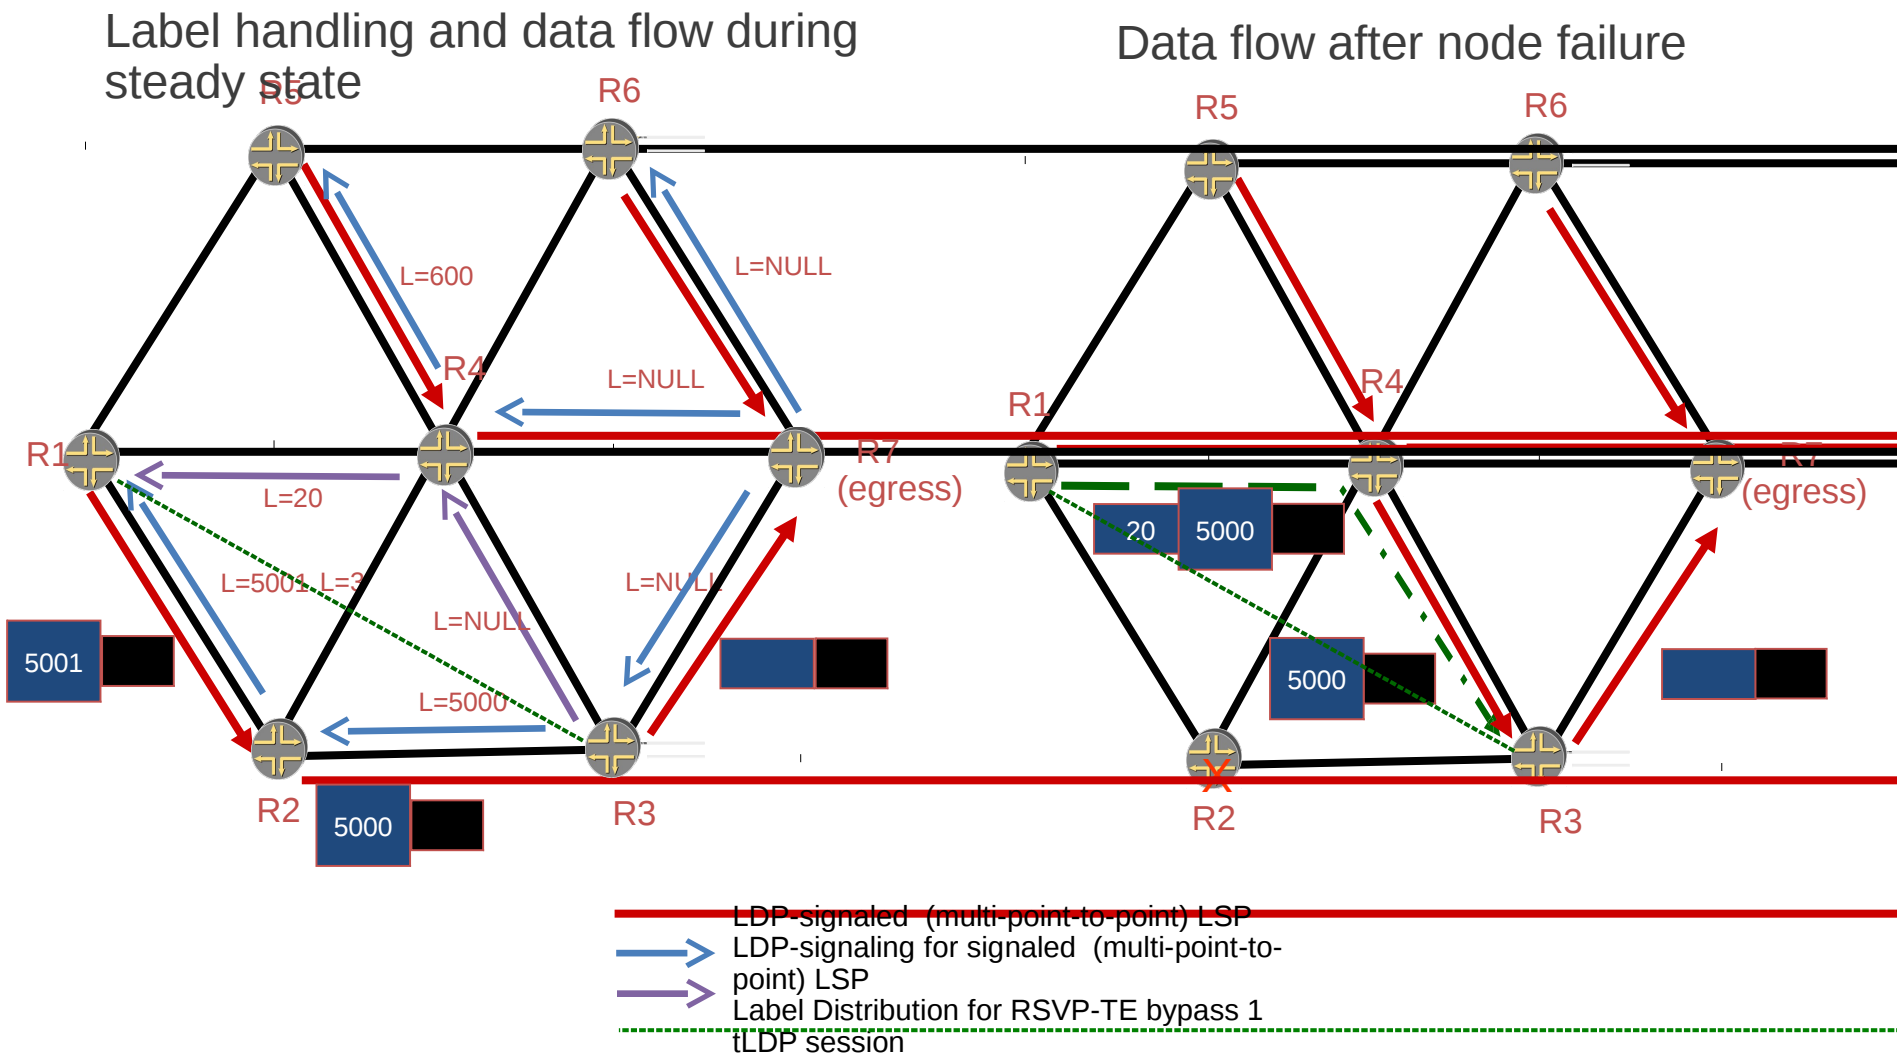

In conclusion

- Local link/node protection for LDP based transport LSPs using RSVP-TE bypasses
- No restrictions on the network topology – provides **topology independent local protection**
- Additional provisioning/configuration required could be fairly small
  - Depends on implementation
  - bypass LSPs from PLR to MPT and Targeted LDP between PLR and MPT can be established automatically
- Relies on the existing IETF standards
  - RSVP-TE for establishing bypass LSPs
  - Targeted LDP to obtain label mapping from MPT
    - Needed only for node protection
- Synergy with link/node protection for mLDP-signaled LSPs