

MPLS Fast Re-route using extensions to LDP

draft-kini-mpls-frr-ldp-01

Authors: Sriganesh Kini & Srikanth Narayanan

IETF 81 Quebec City, July 2011

Motivation and Goal

- › LDP LSPs are widely deployed.
- › Goal of sub 50msec recovery for traffic on routed paths (IGP shortest path)
- › Full coverage needed
- › Solution should be self-contained. It should be independent of other protocols and mechanisms such as IP-FRR, RSVP-TE, IGP convergence etc

Solution characteristics

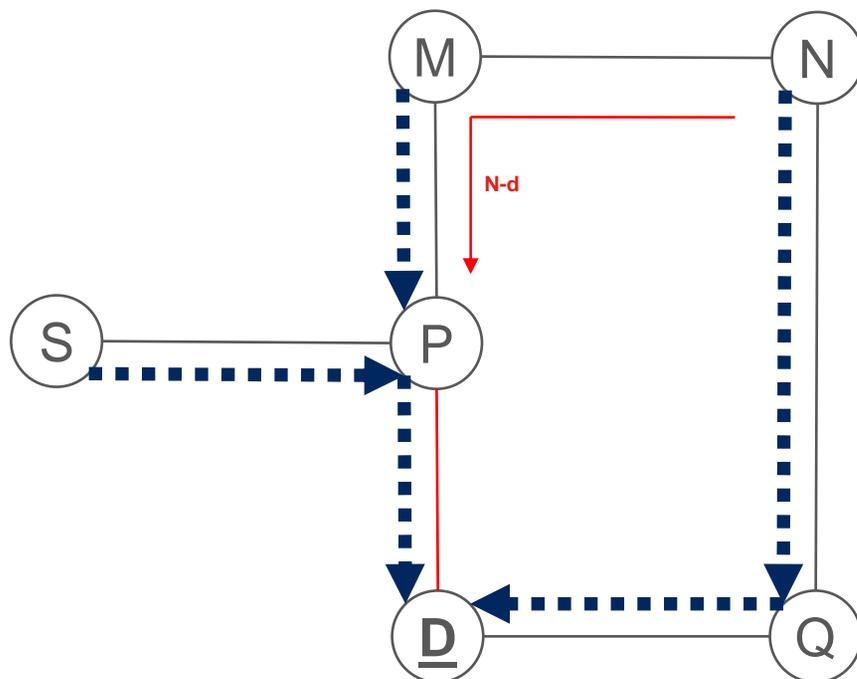
› Local repair mechanism

- Computation intensive tasks are performed much before the actual failure (during steady state).
- Only PLR reacts to the failure trigger to recover the traffic
- Actions at the PLR to recover the traffic are simple (and pre-computed)

Solution summary

- › Defined for link-state IGP. And for platform label space.
- › Backup shortest path (BSP) LDP LSP setup before failure whenever LFA does not exist
- › BSP LSP starts at PLR and merges into shortest path LDP LSP tree. Merge point referred to as BSP-MP.
- › Fast re-route action on detecting failure
 - PLR label switches to pre-selected BSP LDP LSP
 - Stack label to aggregate failures. Use shortest-path LSP from PLR to BSP MP whenever possible.

Link failure protection example

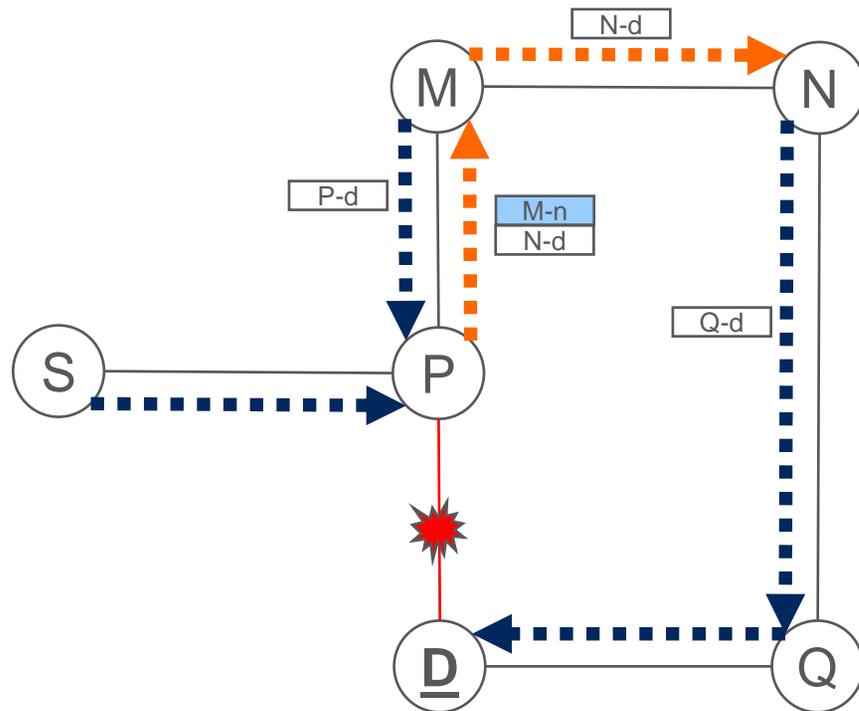


- › Protect link P-D failure
- › For Destination D
 - P is PLR
 - N is merge point
 - N advertises label **N-d** to P for the backup shortest-path LSP
 - **N-d** is the shortest-path LDP LSP label at N for D
 - P uses shortest-path LSP from P to N to tunnel label **N-d**



Traffic flow over shortest path LSP

Link failure protection fast re-routed traffic



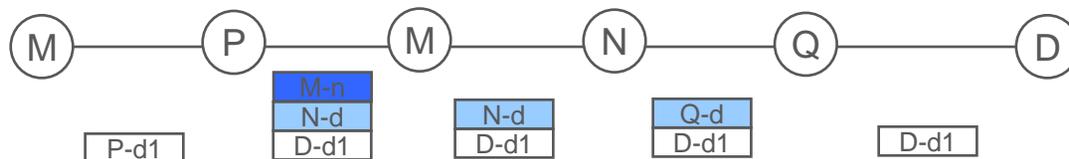
FRR traffic paths to D when link P-D fails

- › P, M, N, Q, D
- › S, P, M, N, Q, D
- › M, P, M, N, Q, D

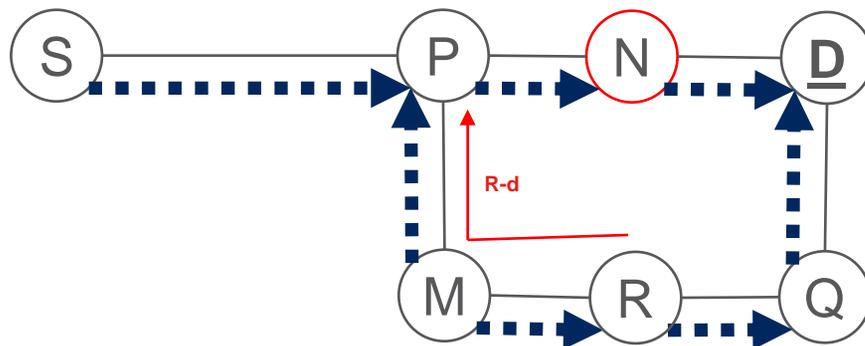
For entire network

- › No 'new' labels needed in the network
- › 12 additional label advertisements needed

Fast re-routed traffic



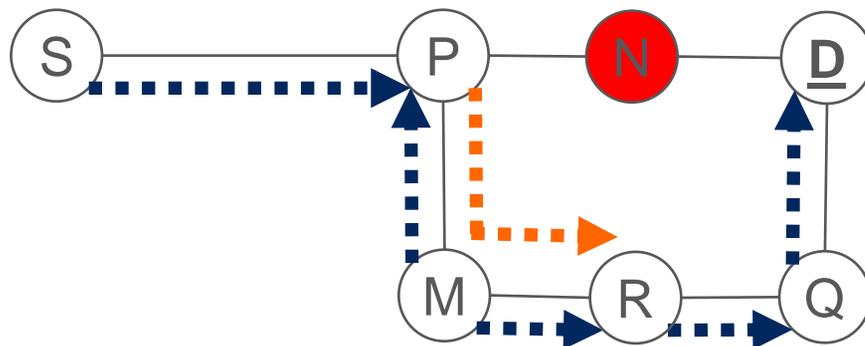
Node failure protection example



- › Node N failure
- › Destination D
- › P is PLR
- › R is merge point
- › R advertises label **R-d** to P for the backup shortest-path LSP

.....➔ Traffic flow over shortest path LSP

Node failure protection fast re-routed traffic



FRR traffic paths to D
when node N fails

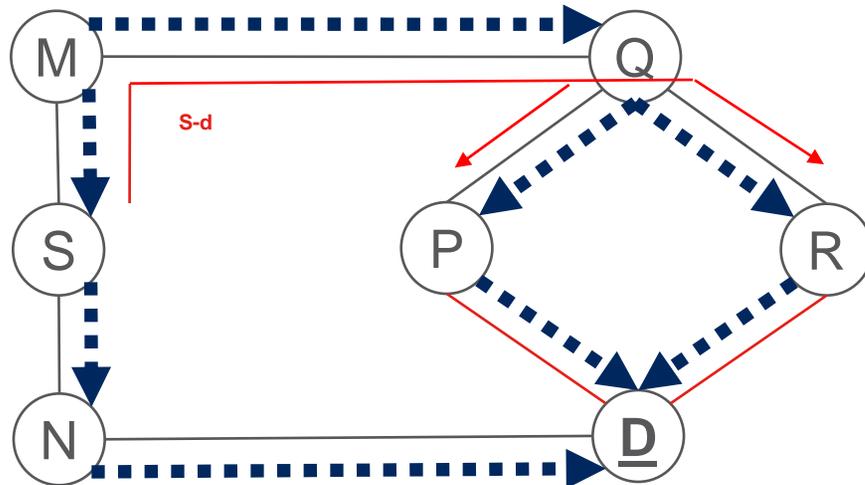
- › P, M, R, Q, D
- › S, P, M, R, Q, D
- › M, P, M, R, Q, D

For entire network

- › No 'new' labels needed in the network
- › 6 additional label advertisements needed

.....➔ Fast re-routed traffic

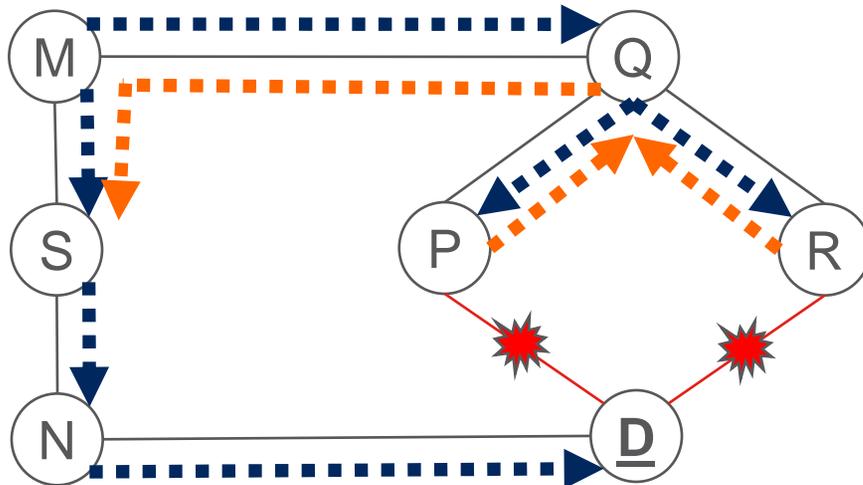
SRLG failure protection example



- › SRLG (link P-D, link R-D) failure
- › Destination D
- › P, R are PLRs
- › S is merge point
- › S advertises its shortest path LSP label (S-d) to P and R for failure against SRLG

.....➔ Traffic flow over shortest path LSP

SRLG failure protection fast re-routed traffic



FRR traffic paths to D
when SRLG fails

- › P, Q, M, S, N, D
- › Q, P, Q, M, S, N, D
- › Q, R, Q, M, S, N, D
- › M, Q, P, Q, M, S, N, D
- › M, Q, R, Q, M, S, N, D

.....➔ Fast re-routed traffic

Operational details

- › Per-nexthop protection can reduce number of BSP LSPs

- › What happens when a shortest-path LSP is not available for tunneling ?
 - Explicit routing for BSP LSP using extensions to LDP

- › Protocol Extensions
 - Failure Element TLV
 - Tunneled FEC TLV (when label stacking not used)
 - Backup Path Vector TLV

Comparison with other approaches

› LDP over RSVP

- Less OpEx (managing one less protocol). Simplicity.
- Less protocol state
- Multi-path on backup

› LFA & Not-via

- Full coverage
- Re-uses MPLS FRR infrastructure
- No IP address management issues

Future Work

- › Analyze applicability

Questions/Comments
