Hierarchical Labels in LDP

draft-kini-mpls-ldp-hierarchy-00

Author: Sriganesh Kini

IETF 80 Prague, March 27 – April 1, 2011

Problem statement

- Typical LDP deployments advertise a unique label per FEC
- When nexthop changes (e.g. due to failure), ILM and FTN entries need to be updated.
- Under scaled scenarios this leads to slower convergence
  - Detection time (Nexthop change)
  - Download time (Control to data plane)

Alternatives considered

- Scale IGP
- Use another protocol to distribute FECs (e.g. BGP)
- Targeted LDP

Solution

- Changes restricted to LDP
- Include “Egress LSR Address” for a FEC in LDP messages. The label is referred to as a **hierarchical label or H-Label**.
- Transport the label mapping with H-Label to all LSRs in the area. (Note: Label mapping to Egress LSR address is unchanged)
- Transport the metric info for a FEC in LDP.
- Install LSP to the shortest-path to the FEC

Solution (contd)

- Use label hierarchy by stacking a label to the Egress above the label of the FEC
- Under failure, downloading the changed nexthop to the egress results in traffic being restored

Solution characteristics

- Works with both link-state and non link-state IGPs
- Enables IGP to carry minimal info
  - Added side effect of IGP speed up due to carrying less state
- With PHP no extra labels required on data path

New/Changed TLVs

- Hierarchical Label TLV
- Metric TLV
- More Label TLV
- Capability Parameter TLV

Messages and procedures

- Supports all current label distribution/control and retention modes
  - Label Mappings received from the neighbor that has the nexthop to the Egress is chosen for advertisement to neighbors
- Metric comparison procedure
  - Allows preference between different metric types

Example

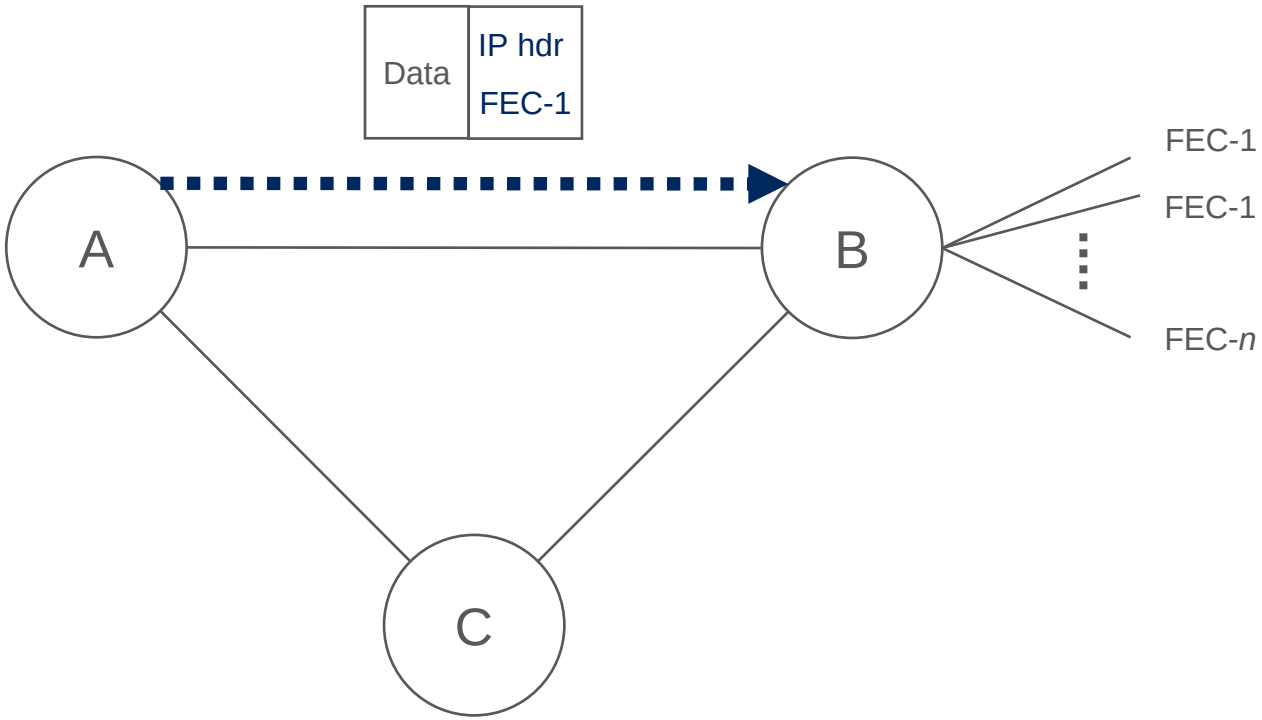

› Pre-failure packet flow (with PHP)

Example (contd)

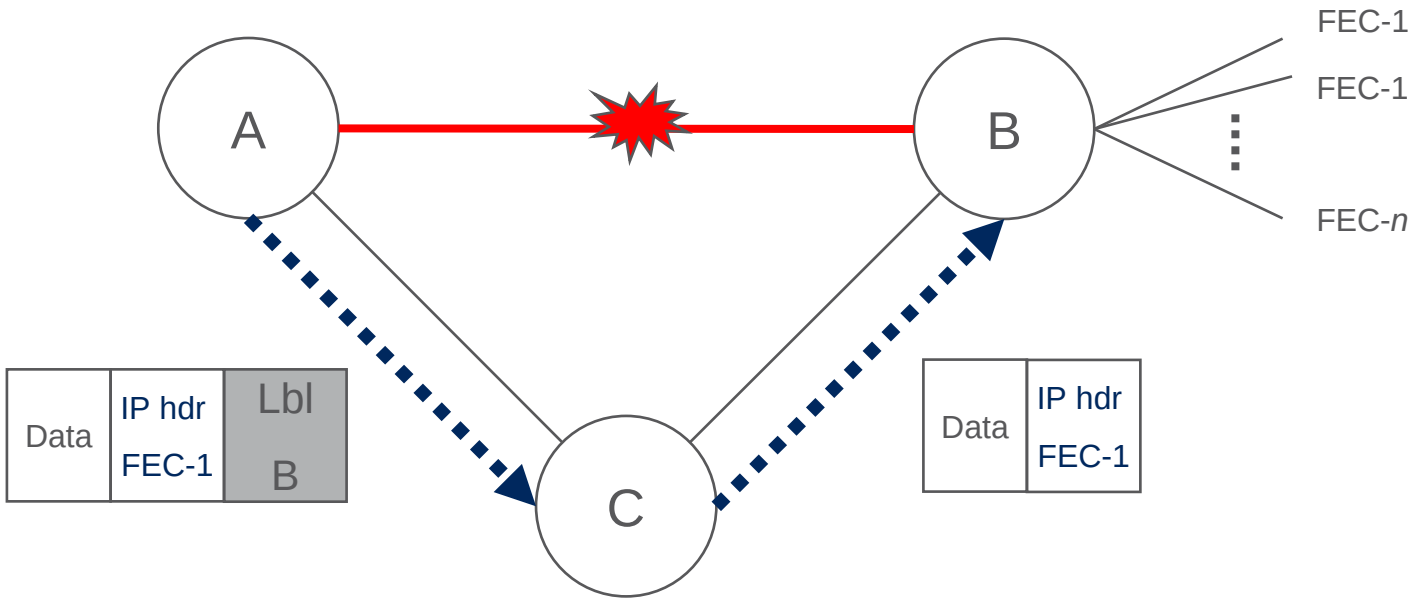

- On failure when A updates FTN/ILM entry for egress B
- Traffic to FEC-1 recovers

Comments/Questions