

2547 egress PE Fast Failure Protection


draft-minto-2547-egress-node-fast-protection-01

Jeyanthan “Minto” Jeganathan minto@juniper.net

Hannes Gredler hannes@juniper.net


Problem statement.

- MPLS provides fast service restoration for link and node failure **except** egress node failure.
- This draft proposes rfc 2547 service restoration when egress node goes down.
- This proposal could be extended other mpls based services.



Egress PE node failure

Local-repair – requirements




Requires protecting service (or LSP) endpoint on primary node (PE3)

- But PLR (P3) does not hold any service state
- Protecting failed node requires PLR to divert the outer (transport) LSP to another backup node (PE4)
- Backup edge node needs to be able to interpret labels allocated by the primary edge node

Egress PE node failure

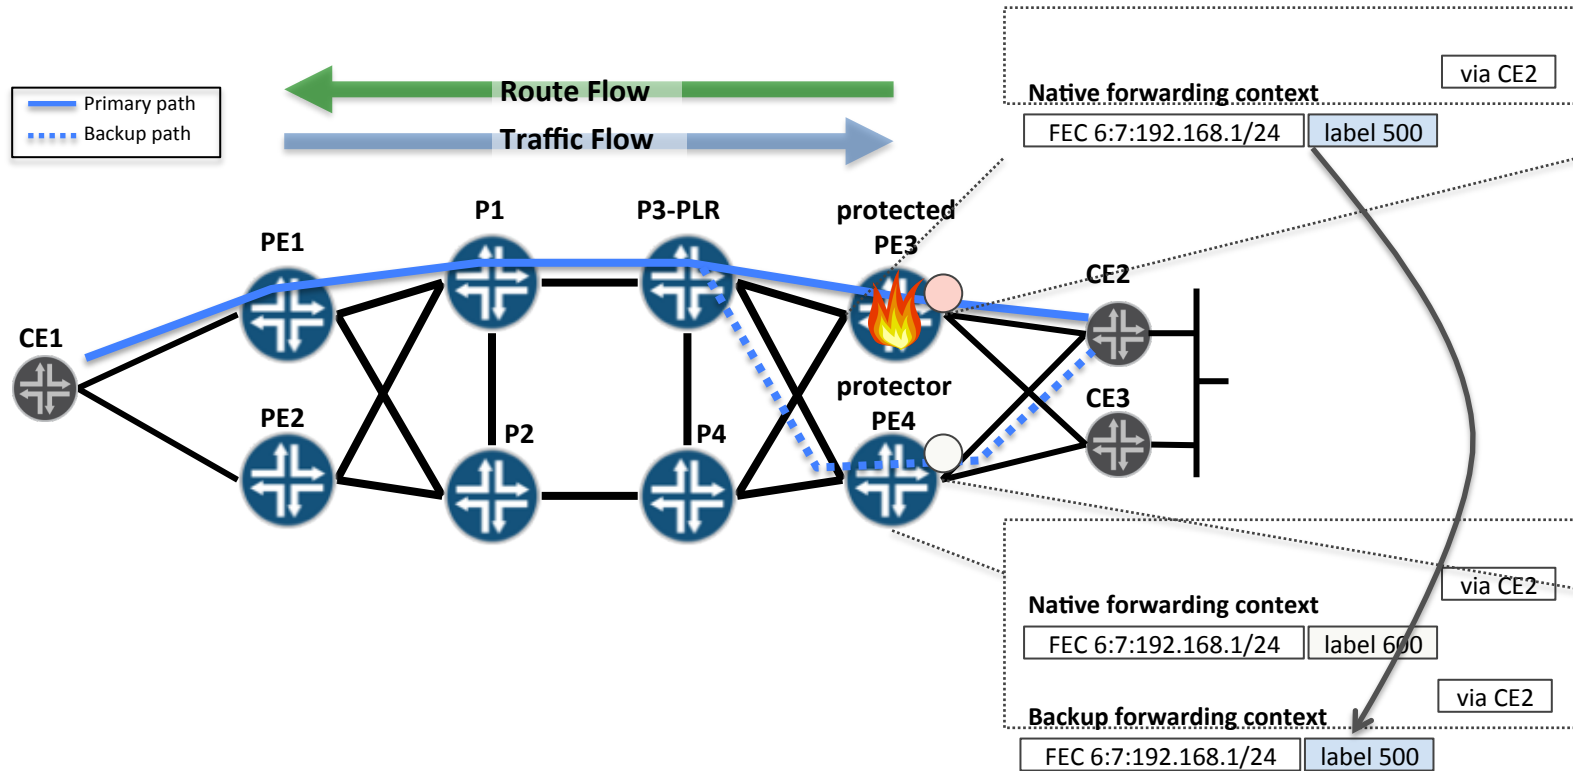
Local-repair – LSP tailend protection



- Protector PE4 maintains a “mirror image” of the protected PE3 service label table
- a context specific label space identified by a context-id (an IP address) present on both protected and protector PEs
- Protected PE3 “owns” the context-id address, advertising it in the BGP Next_Hop attribute (context-id is never used for control plane peerings)
- In case of protected PE3 failure, P3-PLR diverts the traffic destined to the context-id address to the protector PE4 using TE FRR or IP FRR procedures
- Protector PE4 looks up received packets in the context-specific label table for PE3 (identified by the label associated with PE3 context-id), and forwards packets to the right destination

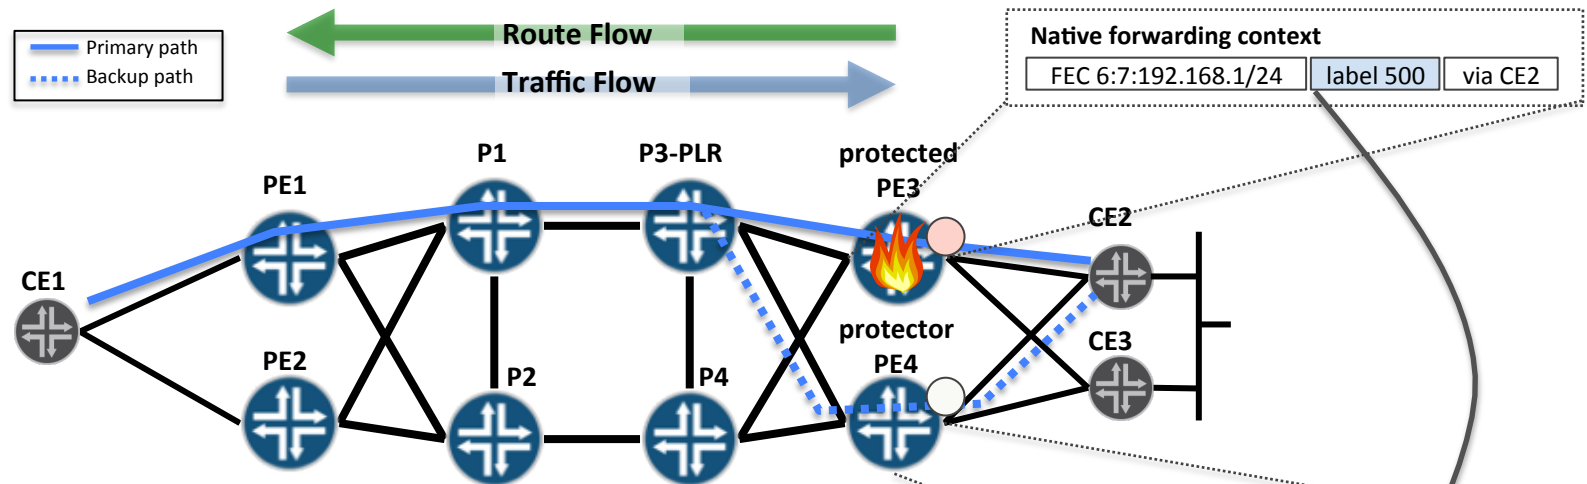
Egress PE node failure

Local-repair – LSP tailend protection





Egress PE node failure

Local-repair – LSP tailend protection



Upon receiving a packet with the outer label equal to the context label, protector PE4 pops the outer label and performs a recursive MPLS lookup in the context specific label space identified by the context label



Protector model

- Co-located protector

The alternate PE and protector is same.

- Centralized protector

One protector protecting more than one PE.

Service mirroring

The backup nexthop selection uses following rules

1. Exact matching route-target set (RD may be different)
2. Exact matching layer-3 Prefix part (excluding RD)

PLR transport FRR

- RSVP

Uses [rsvp-egress-frr](#)

- LDP

Uses LFA with tunneling to protector.

What next ?

- Generalize to other Prefix based MPLS services (EVPN, BGP-LU)
- Questions and comments?