MPLS Egress Protection Framework
draft-shen-mpls-egress-protection-framework-02

Yimin Shen (yshen@juniper.net)

Minto Jeyananth (minto@juniper.net)

Bruno Decraene (bruno.decraene@orange.com)

Updates

- New co-author.
- Editorial changes.

Egress failure

- Failure of the egress router of an MPLS tunnel.
- Two-level failure
  - Transport – Packets can no longer reach the egress router.
  - Service – Packets can no longer reach service instances on the egress router.
- Traffic repair is possible, if a service destination is reachable via an alternative path.

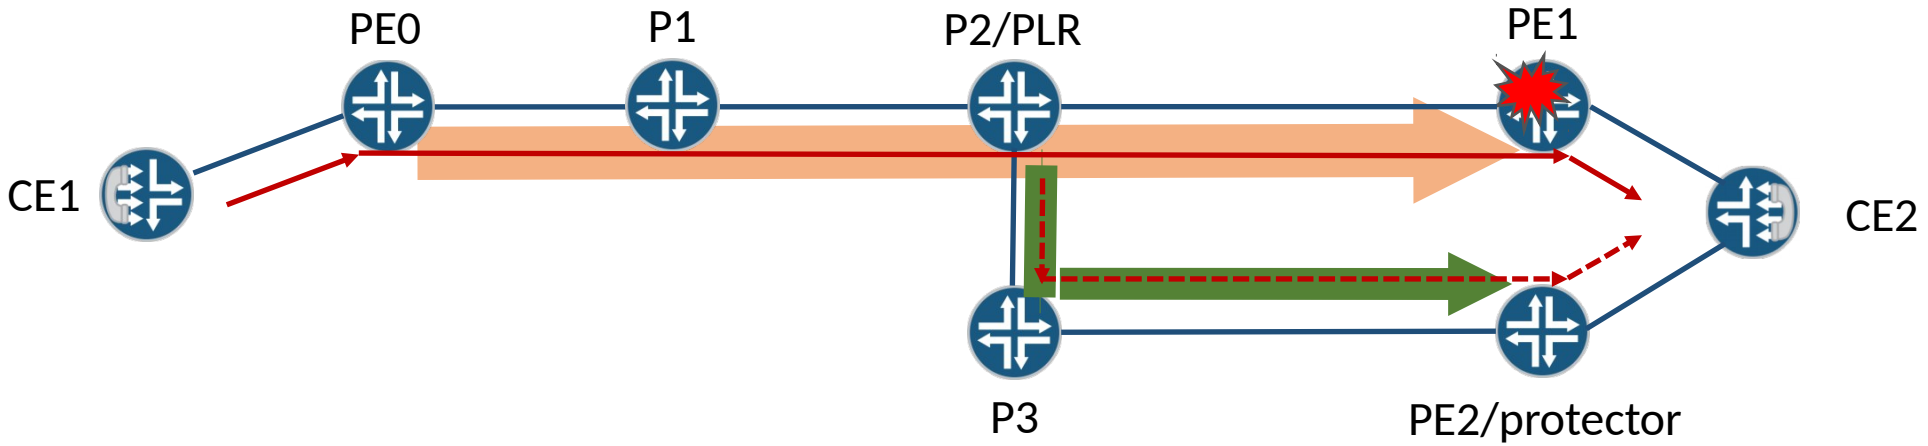

Egress Protection

- Fast reroute for protecting an MPLS tunnel and the services carried by the tunnel against an egress failure.
  - Penultimate-hop router (as PLR) – Local failure detection and local repair.
  - A “protector” - Hosts backup service instances, and forwards rerouted traffic to service destinations.
  - Bypass tunnel from PLR to protector.
- Two-level protection - transport and service.
- Equivalent to the traditional FRR of transit links/routers.
- Complements the traditional FRR.

Goals of This Draft

- Provide a unified framework with a holistic approach for egress protection.
  - Service types – L2/3 VPNs, hierarchical transport, etc.
  - Tunnel types – RSVP, LDP, BGP-LU, SR, etc.
  - Tunnel topologies - P2P, P2MP and MP2P.
- Minimize complexity and impact.
  - Work seamlessly with the traditional FRR.
  - Avoid extensions for tunnel protocols.
  - Provide guidelines for extensions to service protocols.
    - Specific details should be addressed by separate drafts on a per-service basis.

Basic Procedures

- PLR is the penultimate hop router.
  - Pre-establishes a bypass tunnel to protector, in UHP manner.
- Protector
  - Hosts backup service instances.
  - Points the bypass tunnel to a “protection label table” corresponding to the label space of the egress router.
  - Populates the table with service labels learned from the egress router. Sets nexthops based on own connectivity to service destinations.
- Protection
  - PLR detects an egress failure.
  - PLR reroutes packets to the protector via the bypass tunnel, with service label intact.
  - Protector forwards packets to service destinations, based on lookups in the protection label table.

Building Blocks

- Protected egress {E, P} , where E = egress router, P = protector.
  - Serves as a virtual egress node for both MPLS tunnel and services.
- Context ID (CID)
  - A unique IP address assigned to a protected egress {E, P} in routing and TE domains.
  - Every egress protection advertisement or signaling message is with CID.
  - Ingress router, egress router, PLR and protector coordinate based on CID.
- Capability of context label switching on P
  - P uses a context label to indicate a protection label table, i.e. label table corresponding to E's label space.
  - P learns service labels from E, and populates the protection label table.
  - P uses the context label as in-label for bypass tunnel.
  - P forwards services packets received on bypass tunnel to service destinations, based on lookups in the protection label table.

Example

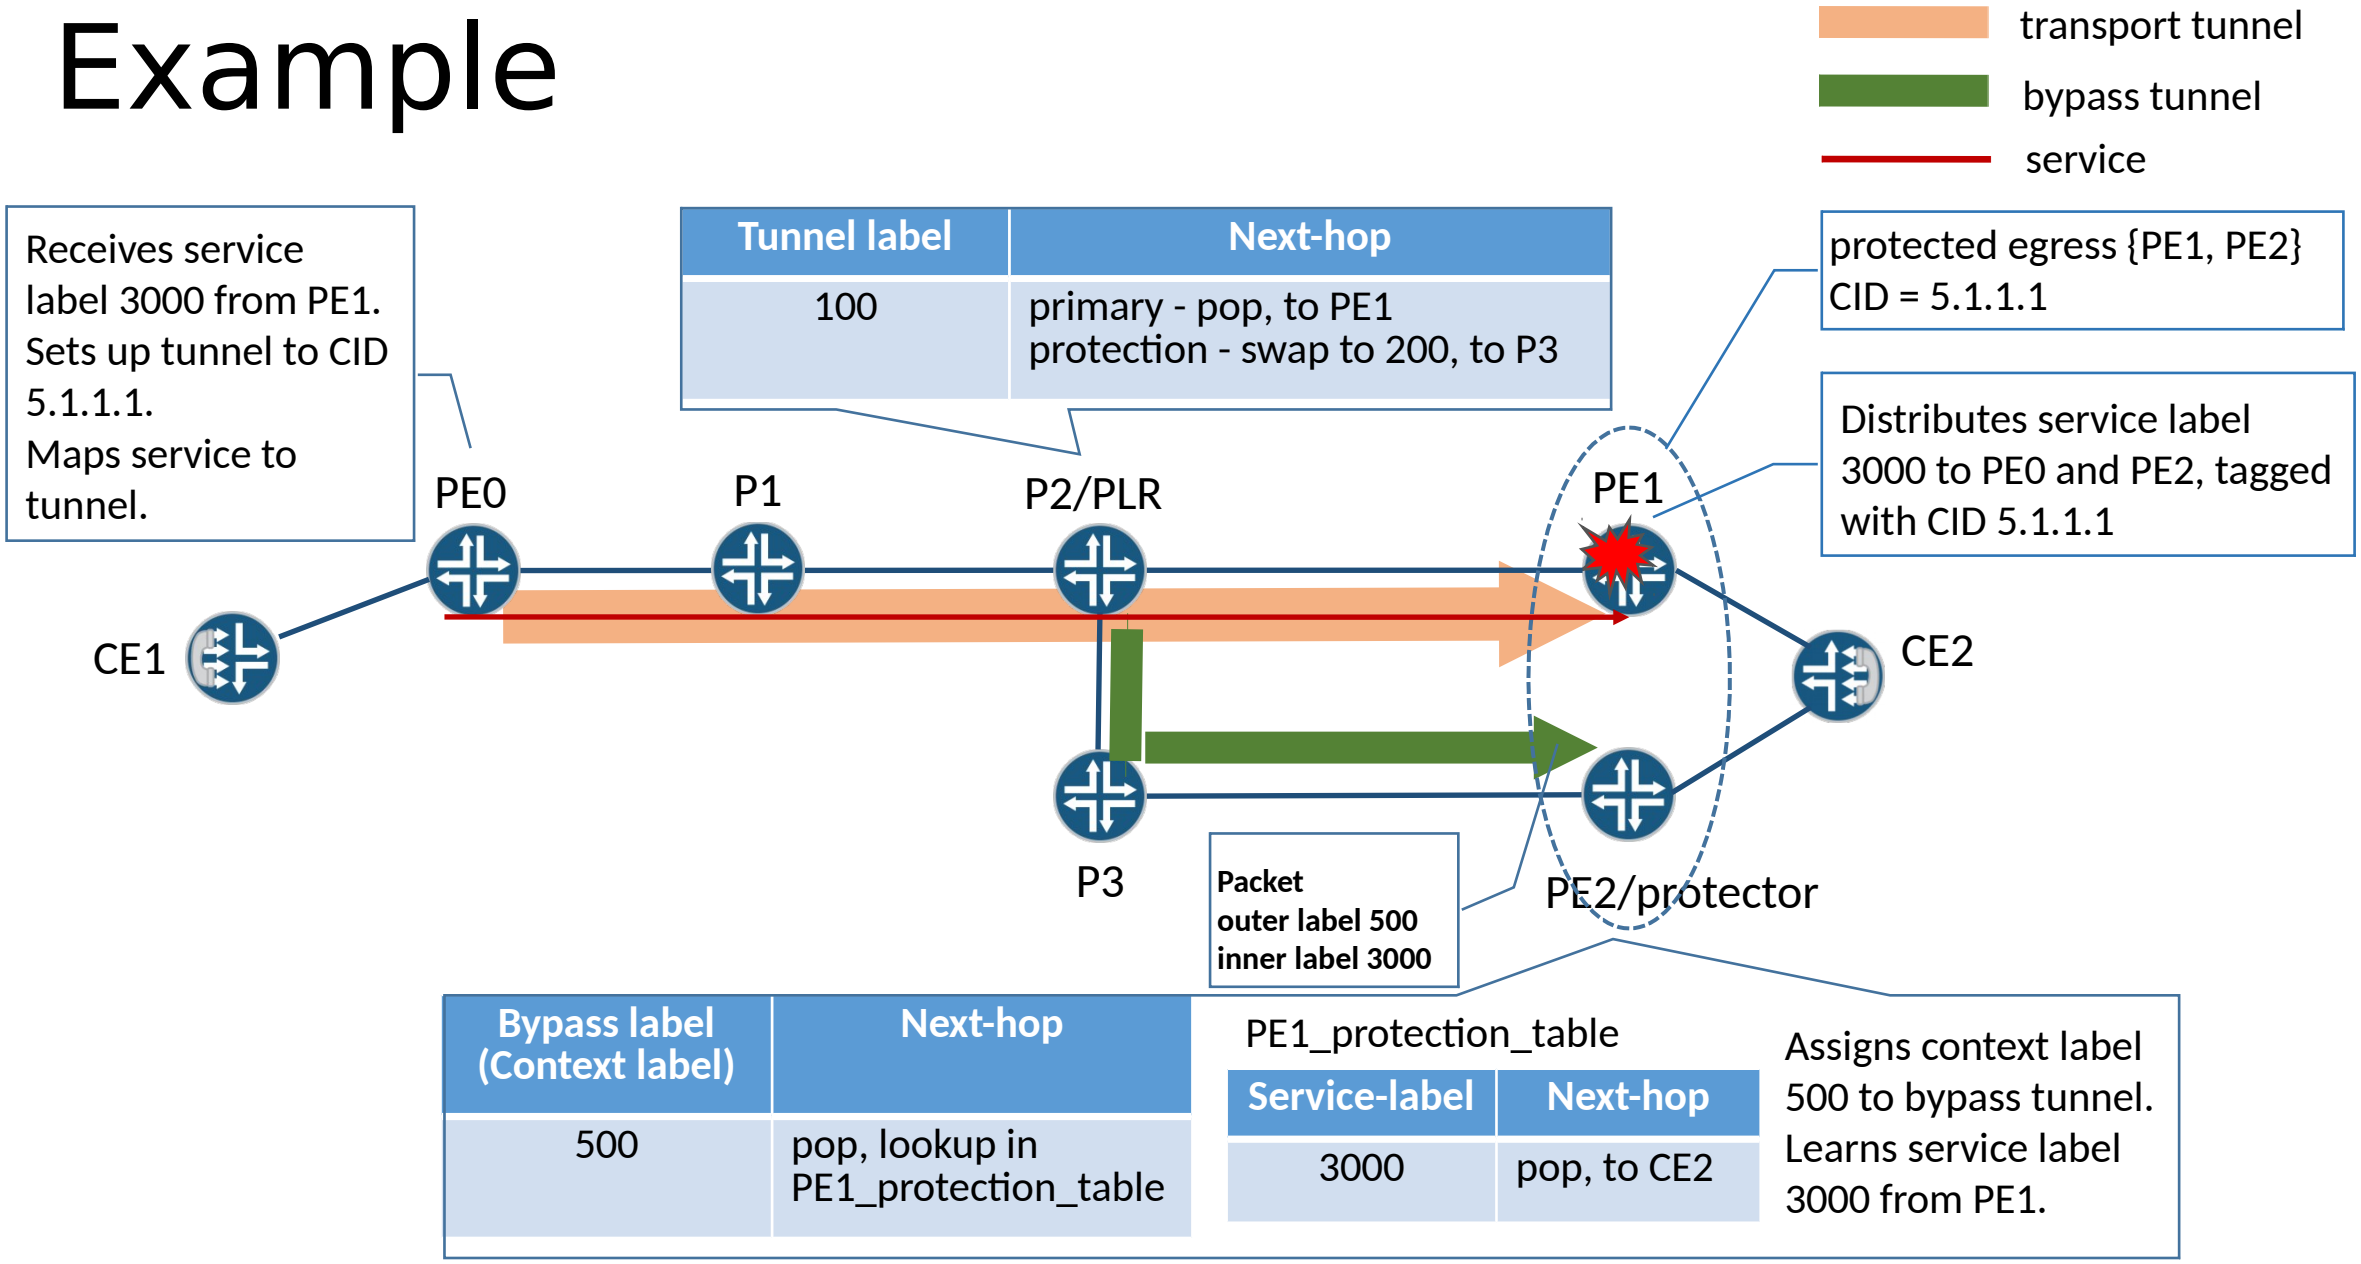
transport tunnel
bypass tunnel
service
Receives service label 3000 from PE1. Sets up tunnel to CID 5.1.1.1.
Maps service to tunnel.
Tunnel label
Next-hop
100
primary - pop, to PE1
protection - swap to 200, to P3
protected egress {PE1, PE2}
CID = 5.1.1.1
Distributes service label 3000 to PE0 and PE2, tagged with CID 5.1.1.1
PE0
P1
P2/PLR
PE1
CE1
CE2
P3
PE2/protector
Packet
outer label 500
inner label 3000
Bypass label (Context label)
Next-hop
500
pop, lookup in PE1_protection_table
PE1_protection_table
Service-label
Next-hop
3000
pop, to CE2
Assigns context label 500 to bypass tunnel.
Learns service label 3000 from PE1.

Protection Establishment

- CID is advertised by IGP.
  - Proxy mode – E and P advertise CID as a proxy node connected to both routers.
  - Alias mode – E advertises CID as regular address. P advertises CID and context label binding by using the “mirroring context segment” defined in SR.
- E tags service label advertisement with CID.
- Ingress router establishes a tunnel to E (CID as destination), and maps service to tunnel.
- P allocates context label for CID, and points context label to E’s protection tables.
- PLR establishes bypass tunnel to P, avoiding E.
- Bypass tunnel is established in a manner that context label is the incoming label on P.
- E distributes service label to P, tagged with CID.
- P installs service label in E’s protection label table. Next-hop is set to P’s own connectivity to service destination.

Next Steps

- Seek comments and feedbacks.
- Seek WG adoption.