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

Presenter: Jeffrey Zhang (zzhang@juniper.net)

Authors: Yimin Shen (yshen@juniper.net)

Minto Jeyananth (minto@juniper.net)

What is MPLS Egress Protection ?

- MPLS egress failure – Failure of the egress node of an MPLS tunnel.
- MPLS egress protection – FRR for protecting an MPLS tunnel and the services carried by the tunnel against an egress failure.
  - Driven by local failure detection and local repair on penultimate hop router.
  - Equivalent to existing FRR for transit link and node failures, e.g. RSVP, LDP, LFA, etc.
  - complements global repair (i.e. end-to-end repair) and control plane convergence

Specific Nature of Egress Protection

- Egress failure must be considered at two levels:
  - Transport level – A failure of transport tunnel, for MPLS packets not being able to reach the egress router.
  - Service level – A failure of every service carried by the tunnel, for service packets not being able to reach the service instances on the egress router.
- Accordingly, egress protection must be provided at both levels.
  - Transport level – PLR redirects packets to a “protector” which acts as backup egress router.
  - Service level – Protector hosts backup service instances to forward service packets to ultimate service destinations.
- The protector and backup service instances are unique components in egress protection.

Goals of This Draft

- Build a unified framework with support for all tunnel protocols and all service types.
- Minimize complexity and impact on MPLS networks by avoiding extension to tunnel protocols.
- Provide guidelines for extensions to service protocols.
  - Should be addressed by separate drafts on a per-service-type basis.
- Serve as an informational document.

Requirements

- Must support P2P tunnels, as well as P2MP and MP2P tunnels by treating sub-LSPs as P2P.
- Must be independent of tunnel protocols, such as RSPV, LDP, BGP, SR.
- Must be generic to support all IP/MPLS services, including layer-2/3 VPNs.
- PLR must be agnostic on services and service labels, and maintain protecton state on a per-tunnel basis, rather than a per-service-label basis.
- PLR must be able to use routing and TE info to resolve path for bypass tunnel.
- Protector must be able to perform context label switching for rerouted MPLS service packets, and perform context IP forwarding for rerouted IP service packets.
- Must work seamlessly with transit link and node protection mechanisms

Basic Idea

- PLR is penultimate hop router.
  - Pre-establishes a bypass tunnel to protector, with UHP.
- Protector is a backup egress router.
  - Points bypass tunnel to special label table and IP forwarding table, corresponding to the label space and IP address space of egress router, respectively.
- Protection
  - PLR reroutes service packets to protector via bypass tunnel, with service label intact.
  - Protector forwards service packets to ultimate destinations, by using the label table and IP forwarding table indicated by bypass tunnel.

Diagram

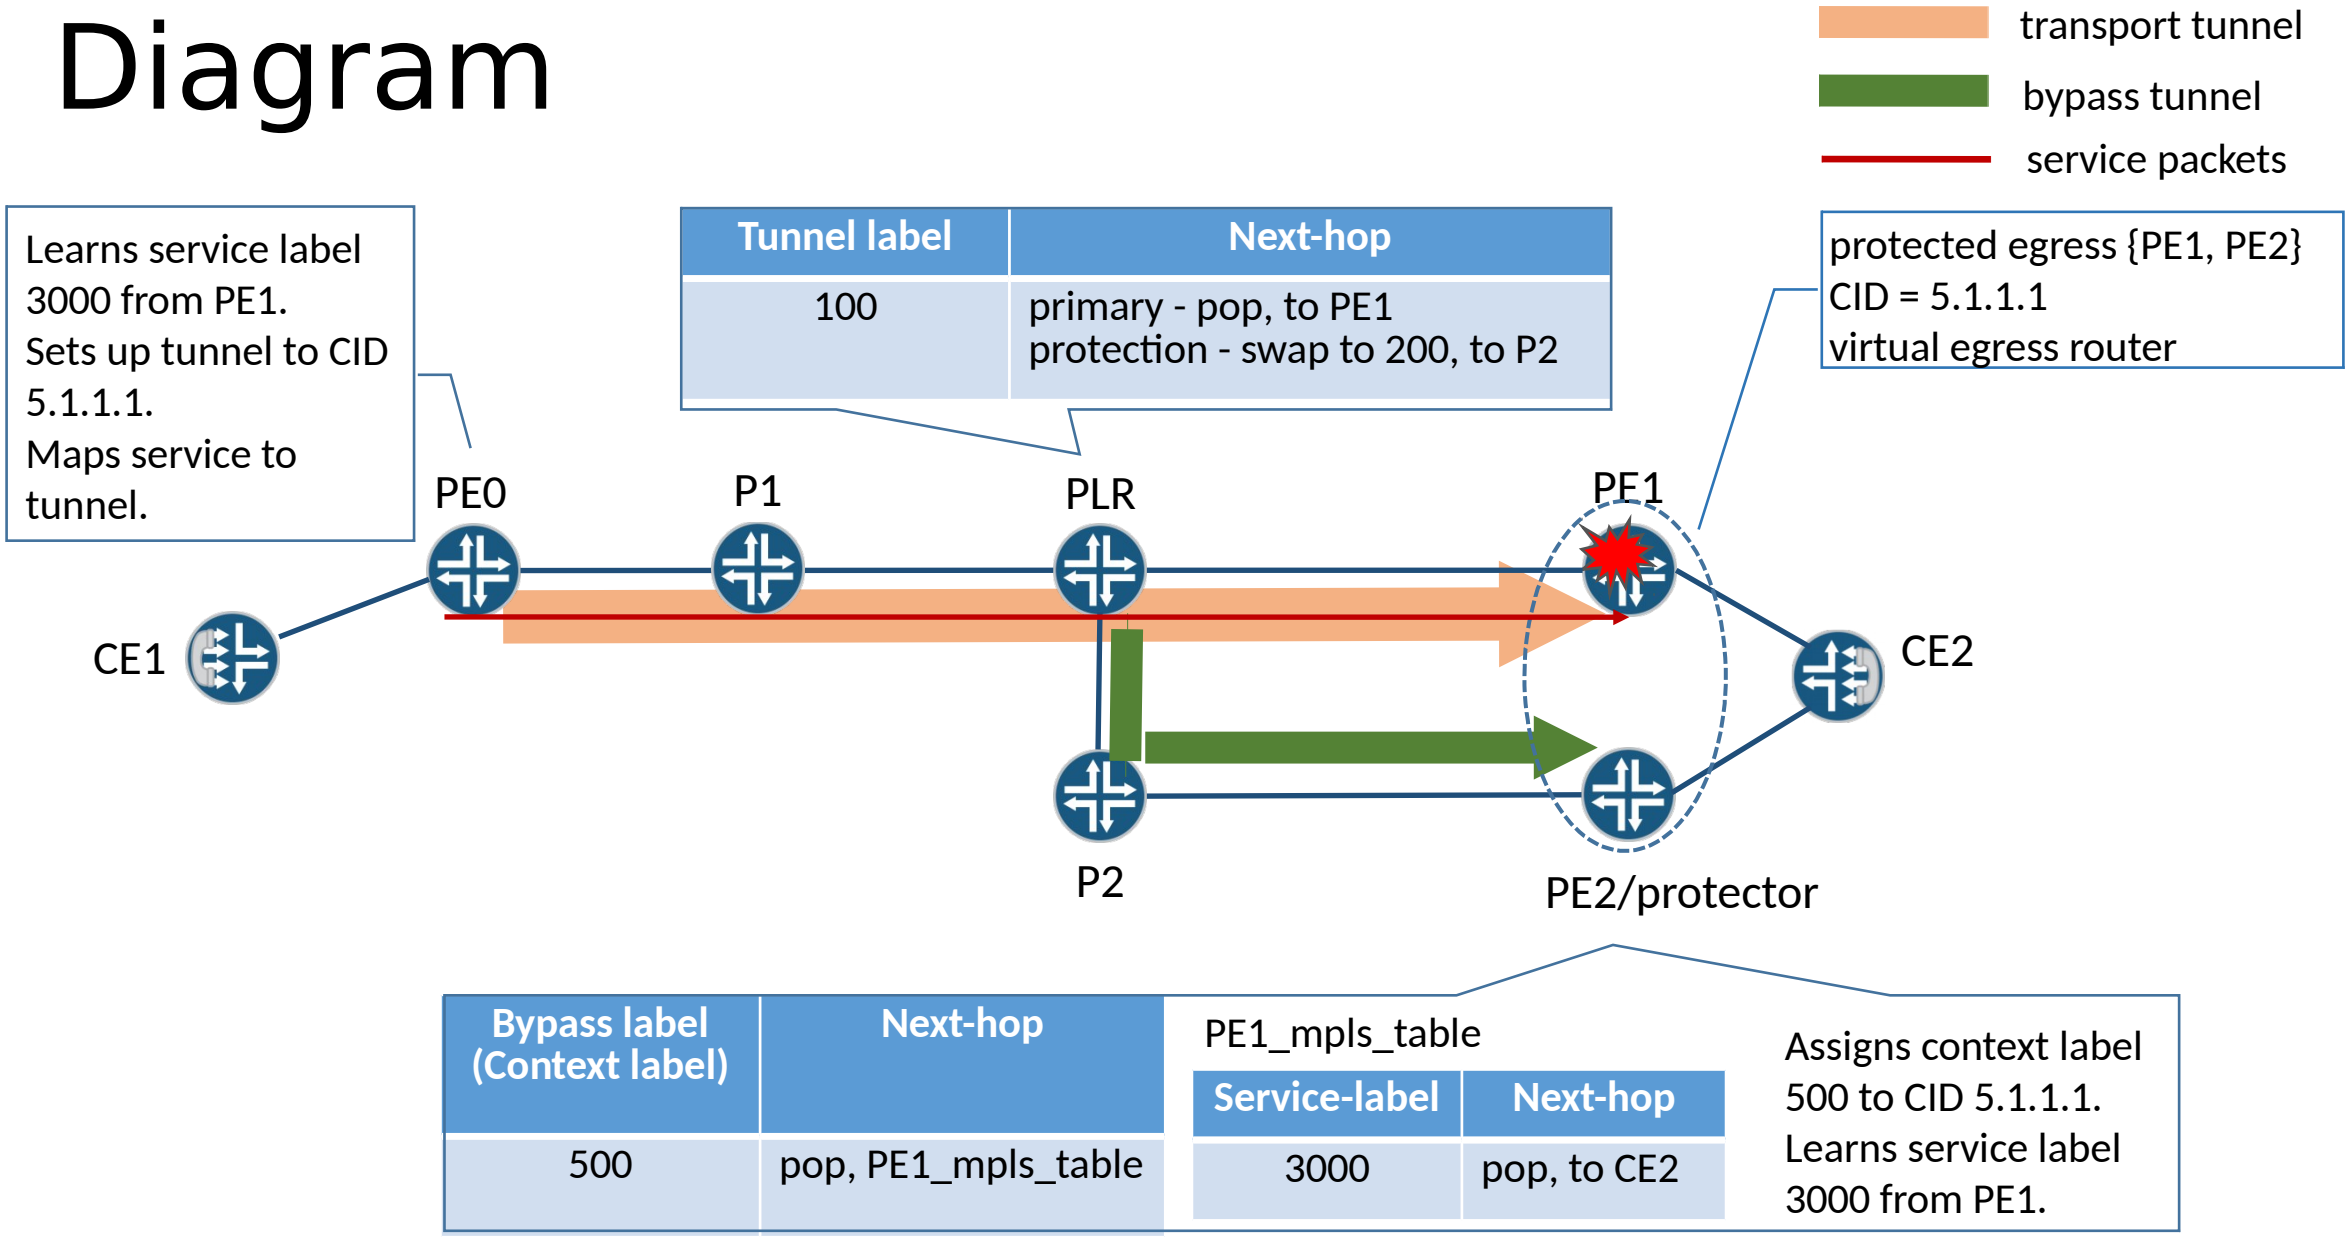
transport tunnel
bypass tunnel
service packets
Learns service label 3000 from PE1.
Sets up tunnel to CID 5.1.1.1.
Maps service to tunnel.
Tunnel label
Next-hop
100
primary - pop, to PE1
protection - swap to 200, to P2
protected egress {PE1, PE2}
CID = 5.1.1.1
virtual egress router
PE0
P1
PLR
PE1
CE1
CE2
P2
PE2/protector
Bypass label (Context label)
Next-hop
500
pop, PE1_mpls_table
PE1_mpls_table
Service-label
Next-hop
3000
pop, to CE2
Assigns context label 500 to CID 5.1.1.1.
Learns service label 3000 from PE1.

Building Blocks

- Protected egress {E, P} , where E = egress router, P = protector.
  - Serves as a virtual egress node for both MPLS tunnel and services.
  - Key strategy to include protector in the schema.
- Context ID (CID, aka. proxy ID)
  - A unique IP address representing a protected egress {E, P}.
  - Reachable via both E and P in routing and TE domains.
- Context label switching and IP forwording on P
  - P assigns an unreserved label (i.e. context label) to CID, to indicate label table and IP forwarding table corresponding to E's label space and IP address space, respectively.
  - P populates the label table with service labels learned from E.
  - P uses the context label as in-label for bypass tunnel.
  - P forwards services packets received on bypass tunnel to ultimate destinations, based on the above tables.

Protection Establishment

- CID is advertised by IGP and IGP-TE.
- E tags service label advertisements with CID.
- Ingress router establishes transport tunnel to E, by using CID as destination. It then maps services to the tunnel.
- PLR establishes bypass tunnel to P, by using CID as destination and avoiding E.
- Protector assigns a context label to CID, and points the label to label table and IP forwarding table corresponding to E's label space and IP address space, respectively.
- P uses the context label as in-label for bypass tunnel.
- E distributes service labels to P
  - All the service labels which E has advertised to ingress router(s), tagged with CID.
  - P installs the service labels in the label table corresponding to E. Next-hops are based on P's own connectivity to service destinations.

Next Steps

- Seek comments and feedbacks.
- Seek WG adoption.