PW Endpoint Fast Failure Protection

draft-shen-pwe3-endpoint-fast-protection-04

Yimin Shen (Juniper)
Rahul Aggarwal (Arktan Inc)
Wim Henderickx (Alcatel-Lucent)

IETF 87, Berlin

Update

- LDP “Egress Protection Capability” TLV
  - Advertised by a protector to a primary PE; also to a backup PE in centralized model.
  - Carried in Initialization and Capability message.

Motivation

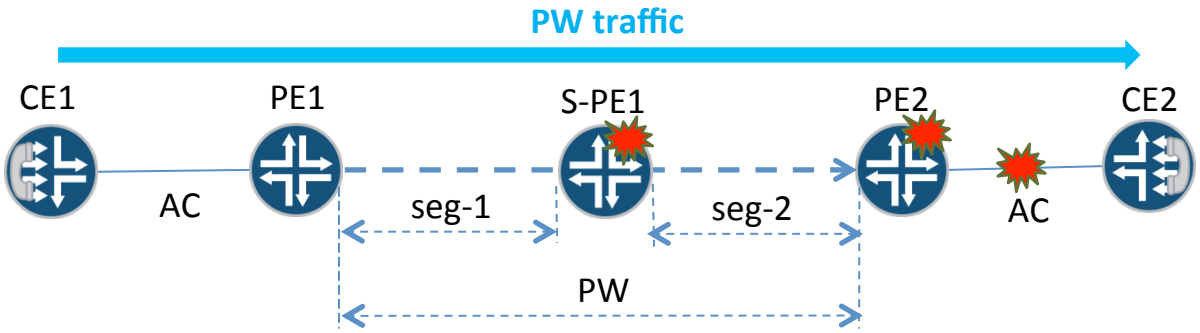

- To specify a local repair mechanism to protect egress AC, egress/terminating PE, and switching PE.
  - 50 ms restoration time.
  - The final piece for end-to-end fast restoration.
  - Can complement global repair.

Scenario 1: Egress AC failure

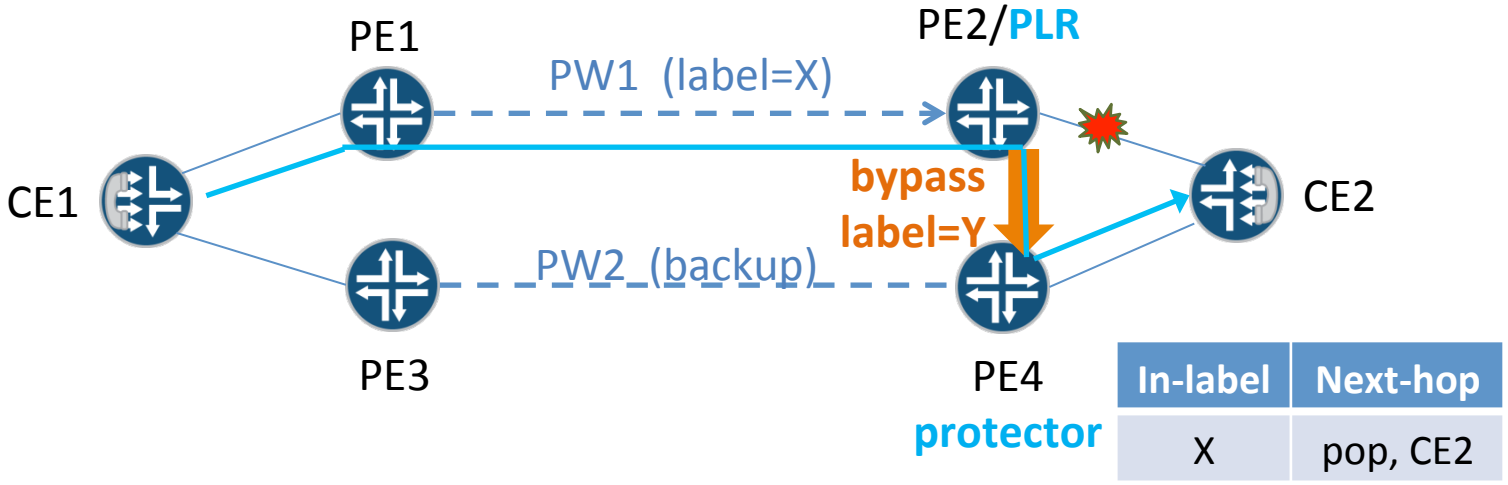

In-label	Next-hop
X	pop, CE2

PLR is the primary PE (PE2). It has established a bypass to protector (PE4).

Protector (PE4):

- Learns PW1's label X from PE2 via the LDP extension.
- Installs a FIB entry for label X in a context label table. Nexthop is CE2.
- Assigns a UHP label Y for the bypass, pointing to this context label table.

Local repair:

- PE2 redirects PW packets through the bypass.
- PE4 receives packets with outer label Y and inner label X.
- PE4 looks up label X in the context label table, and sends packets to CE2.

Scenario 2: Egress (T-)PE failure

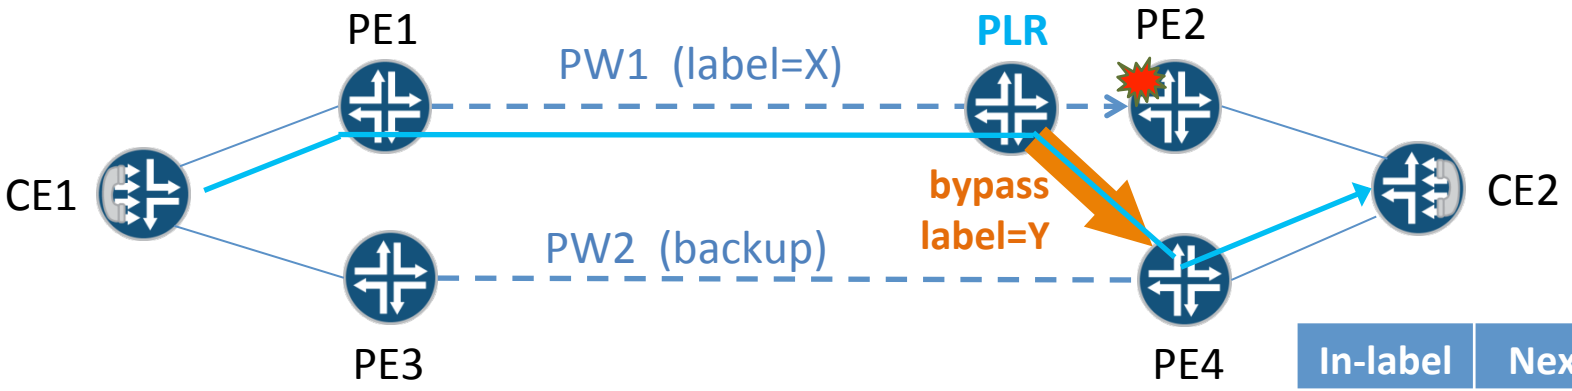

In-label	Next-hop
X	pop, CE2

PLR is the penultimate hop router. It has established a bypass to protector (PE4).

Protector (PE4)

- Learns PW1's label X from PE2 via the LDP extension.
- Installs a FIB entry for label X in a context label space. Nexthop is CE2.
- Assigns a UHP label Y for the bypass, pointing to this context label table.

Local repair:

- PLR redirects PW packets through the bypass.
- PE4 receives packets with outer label Y and inner label X.
- PE4 looks up label X in the context label table, and sends packets to CE2.

Scenario 3: S-PE failure

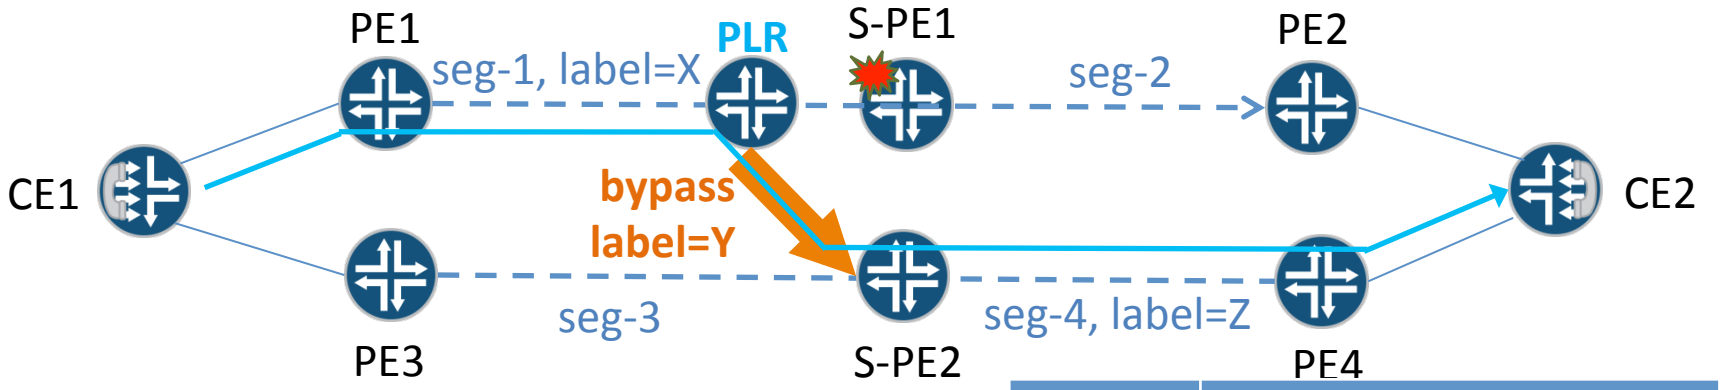

In-label	Next-hop
X	swap to Z; push transport LSP label to PE4

PLR is the penultimate hop router of seg-1. It has a bypass to protector (S-PE2).

Protector (S-PE2)

- Learns PW1's label X from S-PE1 via the LDP extension.
- Installs a FIB entry for label X in a context label table. Nexthop is seg-4.
- Assigns a UHP label Y for the bypass, pointing to this context label table.

Local repair:

- PLR redirects PW packets through the bypass.
- S-PE2 receives packets with outer label Y and inner label X.
- S-PE2 looks up label X in context label table, swaps to label Z, pushes a transport label to PE4.

Context ID

- A unique IP address that identifies and associates a pair of <primary PE, protector>.
- Advertised by IGP or IGP-TE to be reachable via both routers.
  - Via primary PE, by default.
  - Via protector, when bypassing primary PE.

Local protection on PLR

- Transport tunnel's destination = context ID.
  - The tunnel is routed or signaled to primary PE.
  - PLRs are automatically provided with the identity of protector.
- PLRs set up bypass tunnels to protector, by avoiding primary PE.
- PLRs install bypass routes in forwarding plane, in anticipation of failures.

Forwarding state on protector

- Protector learns PW labels from primary PE.
  - Targeted LDP session.
  - Primary PE tags each PW label with context ID.
- Protector installs the PW labels in a context label table, identified by the context ID.
- Protector points incoming bypass tunnels (destination = context ID) to this context label table.
- During local repair, PW packets received on a bypass tunnel are looked up in this context label table, and forwarded to target CE.

Summary

- Single- and multi-segment PWs.
- MPLS and IP transport tunnels and bypass tunnels.
- LDP Extensions:
  - Egress Protection Capability TLV.
  - Protection FEC Element TLV, for FEC 128 and 129.
- Protection models:
  - Co-located protector – simplicity.
  - Centralized protector – scalability.

Next steps

- The document has been stable.
- Would like to ask for WG adoption.