

PW Endpoint Fast Failure Protection

draft-shen-pwe3-endpoint-fast-protection-03

Yimin Shen (Juniper)

Rahul Aggarwal (Arktan Inc)

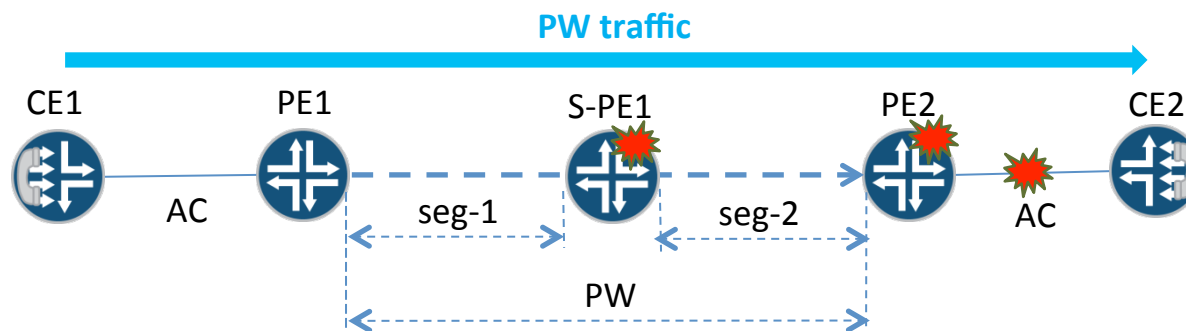
Wim Henderickx (Alcatel-Lucent)

IETF 86, Orlando

Update

- The new revision provides high-level guidelines for IGP advertisement and path computation for context IDs.
 - Allows a range of options.
 - Considers specific option as out of scope.
 - Provides flexibility to accommodate existing and future approaches, which may be generic to egress protection for all layer-2/3 VPNs.

Motivation



- It is important to have local repair (aka. fast reroute) mechanisms protect every link and node along a PW.
- This draft provides a such kind of mechanism for protecting egress AC, egress/terminating PE, and switching PE.
 - The final piece for a complete solution for end-to-end service restoration in 10s of milliseconds.
 - Complements global repair.

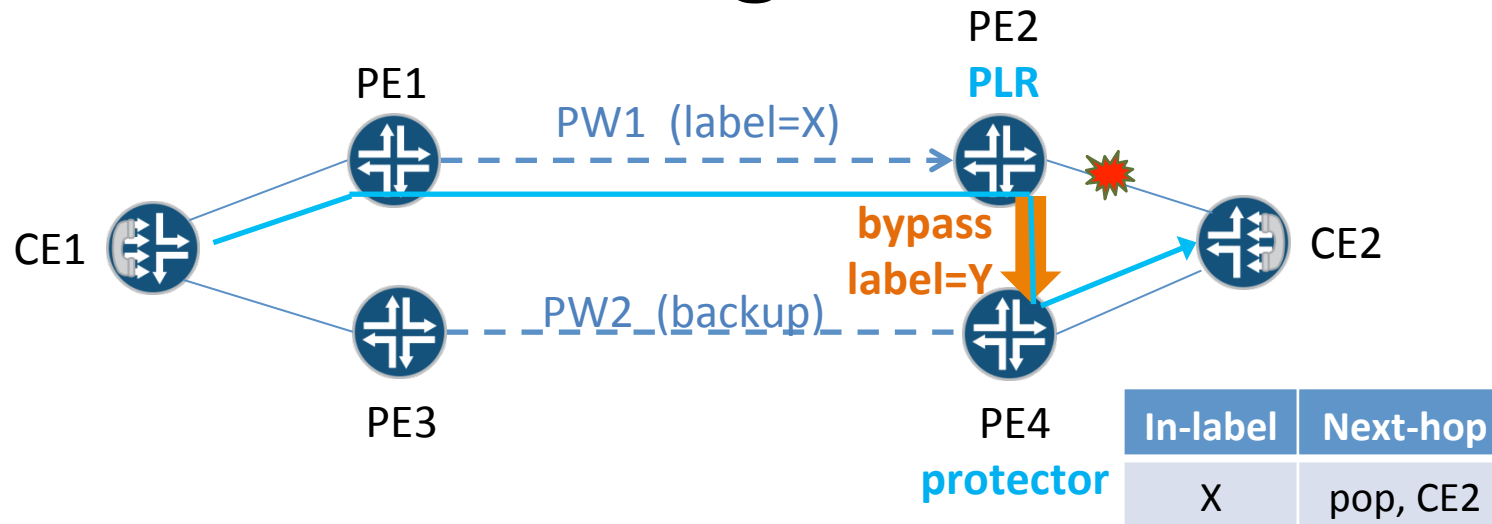
Highlights

- Single- and multi-segment PWs.
- RSVP, LDP and IP transport tunnels and bypass tunnels.
- Upstream label assignment and context label switching.
- LDP Extension
 - Protection FEC Element TLV, for FEC 128 and 129.
- Efficiency and scalability
 - Co-located protector model
 - Centralized protector model

How it works

- PE distributes PW label to a protector via LDP extension.
- Protector installs PW label in a context label table.
- Upon detecting a failure, PLR redirects PW packets to protector via a pre-established bypass tunnel, with PW label intact.
- Protector looks up PW label in context label table, and sends packets to target CE.
- Context ID
 - Identifies a pair of <primary PE, protector>.
 - Serves as a forwarding context, indicating the primary PE's label space to the protector.
 - Serves as destination for transport tunnel and bypass tunnel, to facilitate path computation.

Scenario 1: Egress AC failure



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

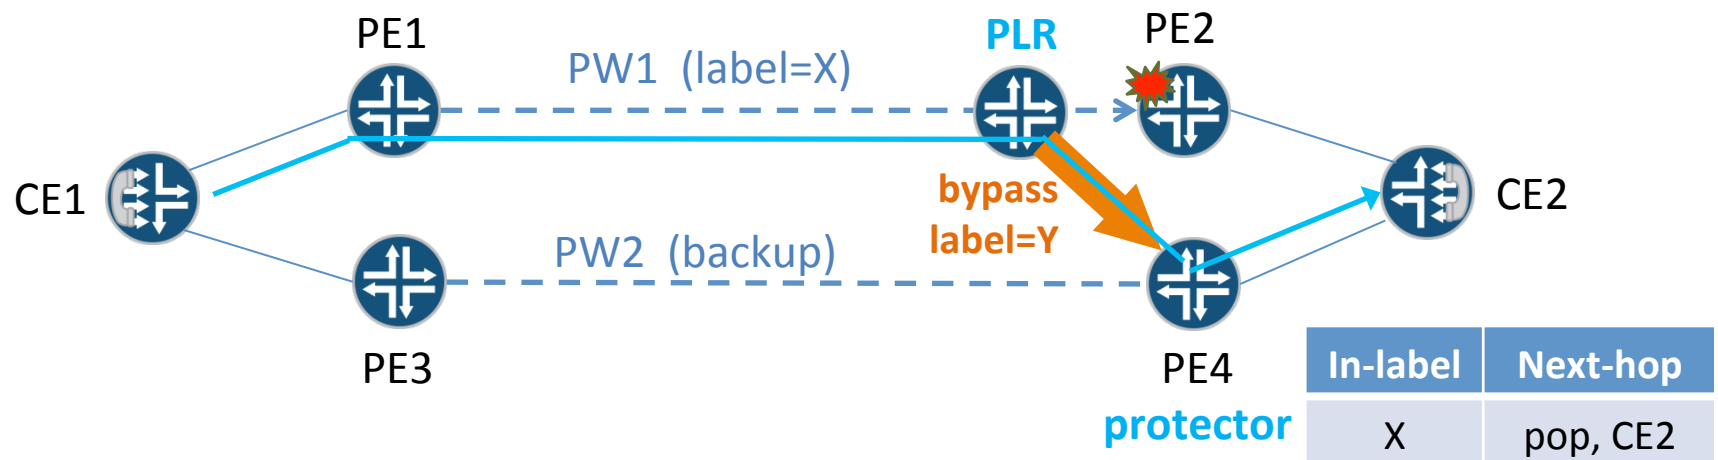
Protector (PE4)

- Learns PW1's label X from PE2 via the LDP extension.
- Installs a FIB entry for label X in a context label table. Next-hop 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



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. Next-hop 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



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. Next-hop 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.

Next steps

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