Pseudowire Endpoint Fast Failure Protection

draft-shen-pwe3-endpoint-fast-protection-00

Rahul Aggarwal (rahul@juniper.net)
Yimin Shen (yshen@juniper.net)

IETF 81, Quebec City, Canada

Overview

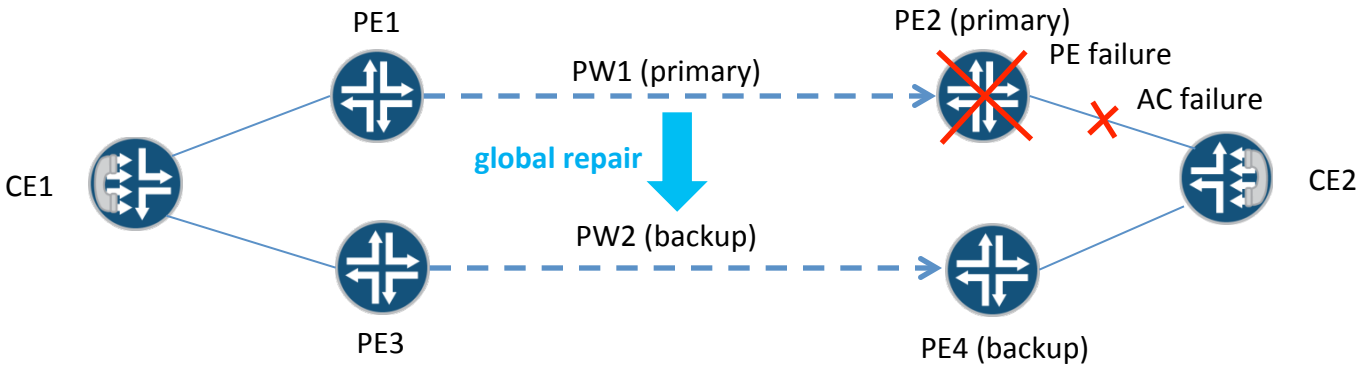

- PW egress endpoint failures:
  - AC
  - PE
- Global repair mechanisms rely on multi-homed CE and PW redundancy. Control plane convergence may be relatively slow.
- This draft proposes local repair to complement global repair.
  - Protection against both egress AC and PE failures.
  - Fast restoration in the order of 10s of milliseconds, comparable to FRR.
  - LDP and BGP signaled PWs.

Overview (cont.)

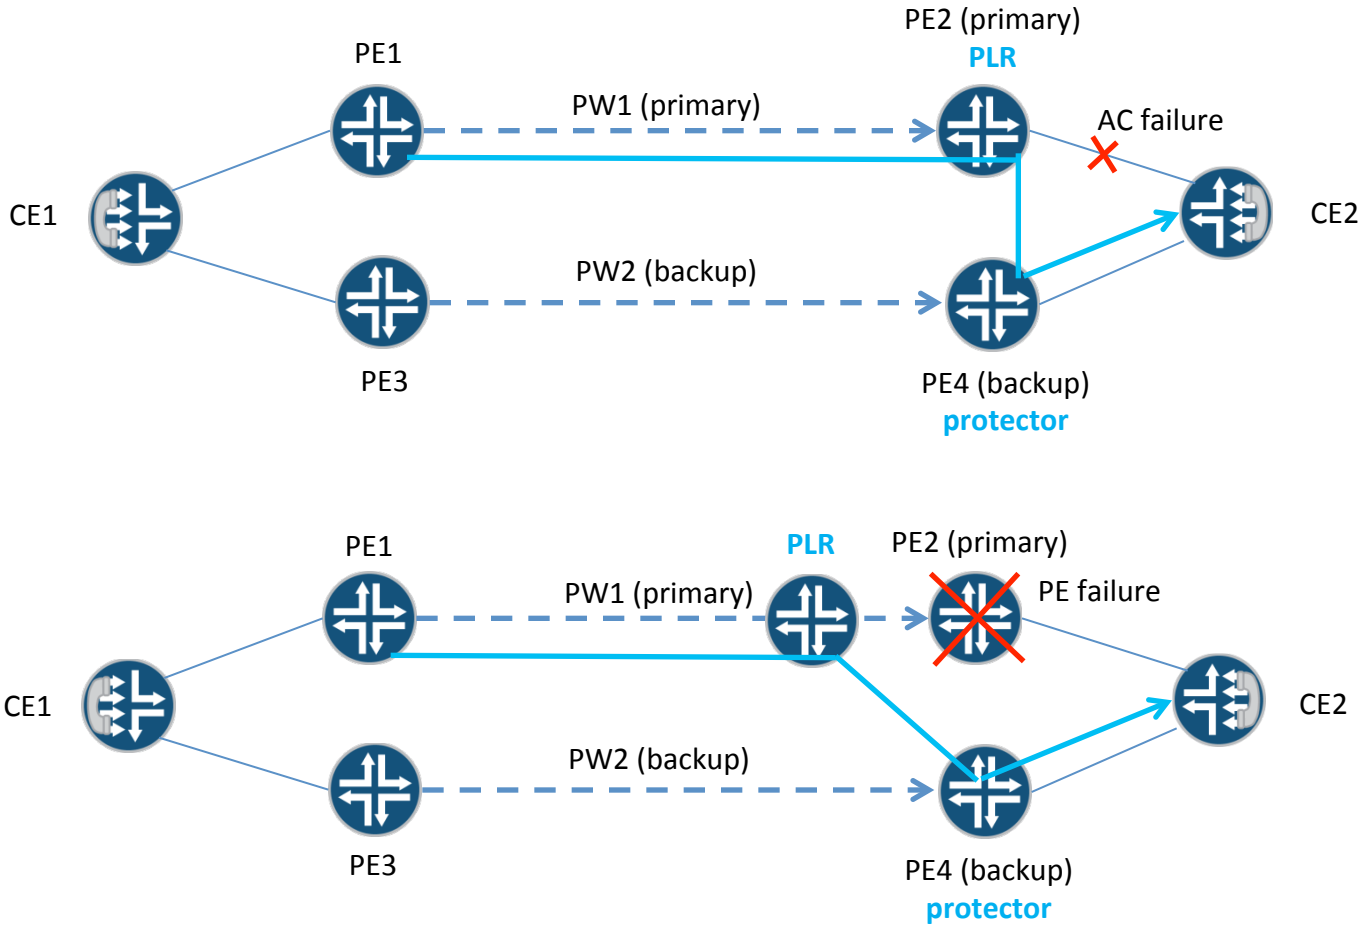

Terminology

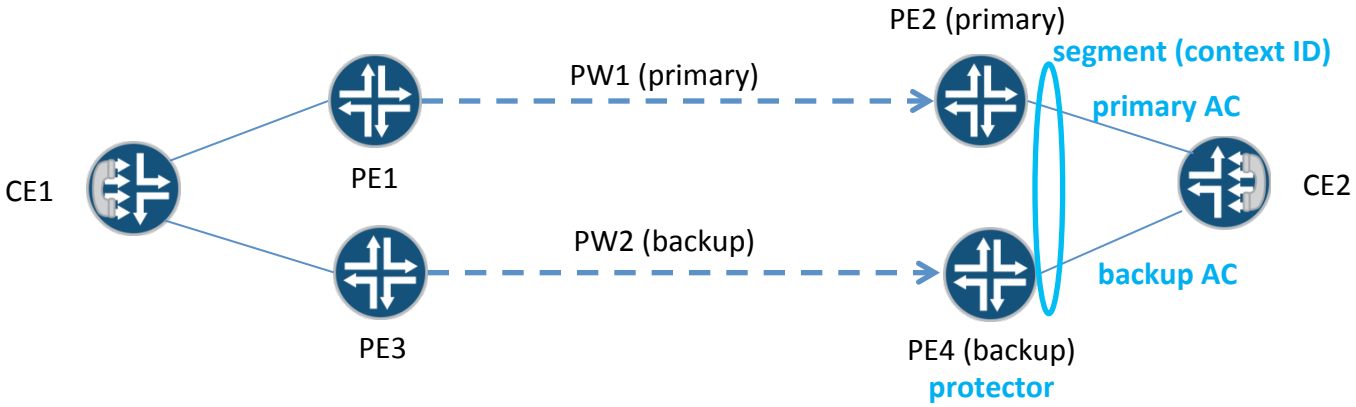

- **Segment**: The set of ACs that a multi-homed CE uses to connect to PEs for a given emulated service.
- **Context ID**: A unique IP-v4 address that is assigned to a segment. It is advertised by both primary PE and the protector PE in IGP as a stub link.
- **PLR (point of local repair)**: The router where a failure is detected and local repair is performed.
  - AC failure: primary PE.
  - PE failure: peultimate hop of PW transport LSP.
- **Protector PE**: A PE that protects a segment, and receives PW packets from PLR upon a failure.
- **Bypass LSP**: An LSP used by PLR to tunnel PW packets to a protector PE upon a failure.
  - Destination is context ID.
  - UHP (ultimate hop popping).

Procedure for AC Failure

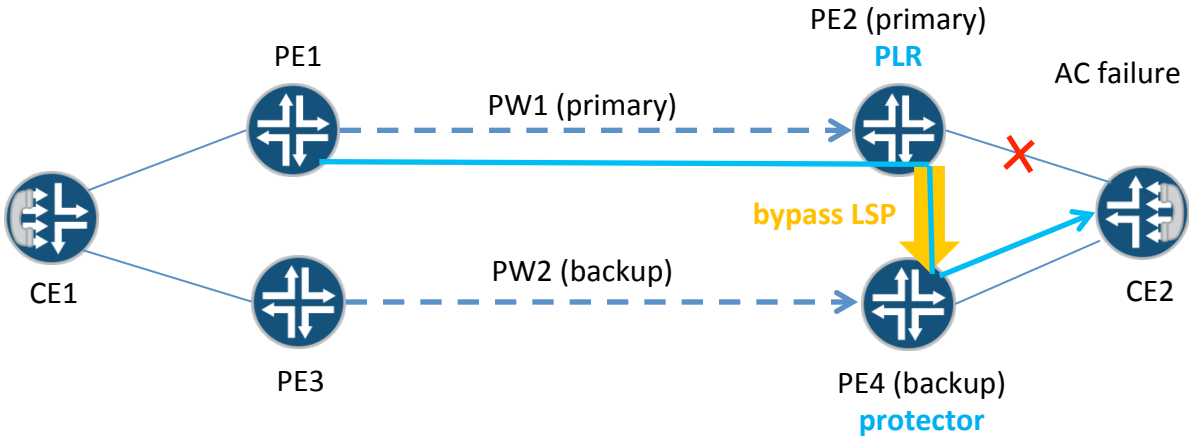

PE2 (PLR) sets up local protection.

- A bypass LSP is established to PE4 (protector PE).
  - Destination = context ID.
  - Requires UHP.
- Upon a failure, PE2 tunnels PW packets to PE4 via bypass LSP.
  - PW label remains unchanged.

Procedure for AC Failure (cont.)

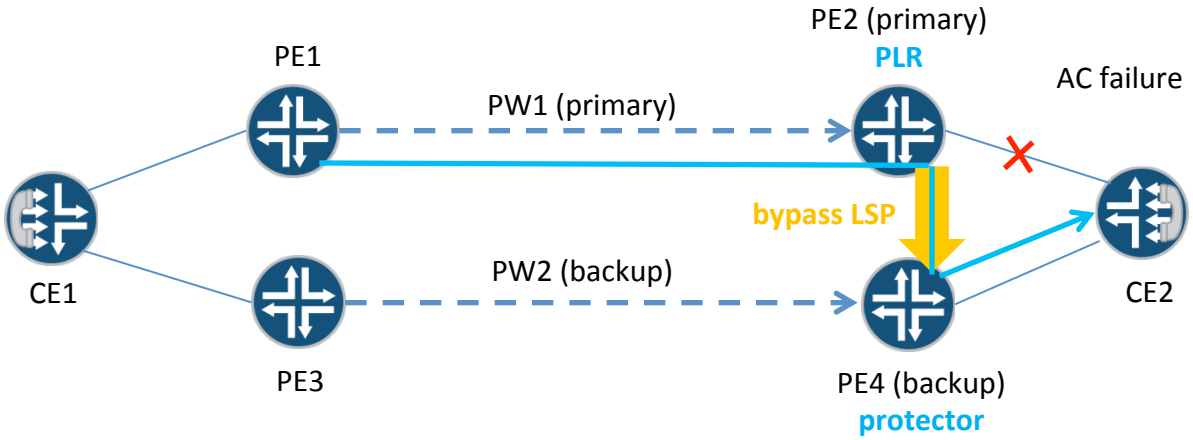

PE4 (protector PE) sends packets to CE2 based on PW label (i.e. upstream label) assigned by PE2.

- Learns {PWid, label, context ID} from PE2 via LDP.
  - PWid: newly defined Protection FEC Element.
  - Label: Upstream Assigned Label TLV.
  - Context ID: Logical Interface ID in IF_ID TLV.
- Installs forwarding state for PW label in a context-specific label space indicated by context ID.
- Allocates a non-reserved label for bypass LSP, pointing to this label space.
- Performs 2 label lookups during forwarding:
  - 1st lookup based on bypass LSP's label finds context-specific label space.
  - 2nd lookup based on PW label in context-specific label space results in sending packet to CE2.

Procedure for PE Failure

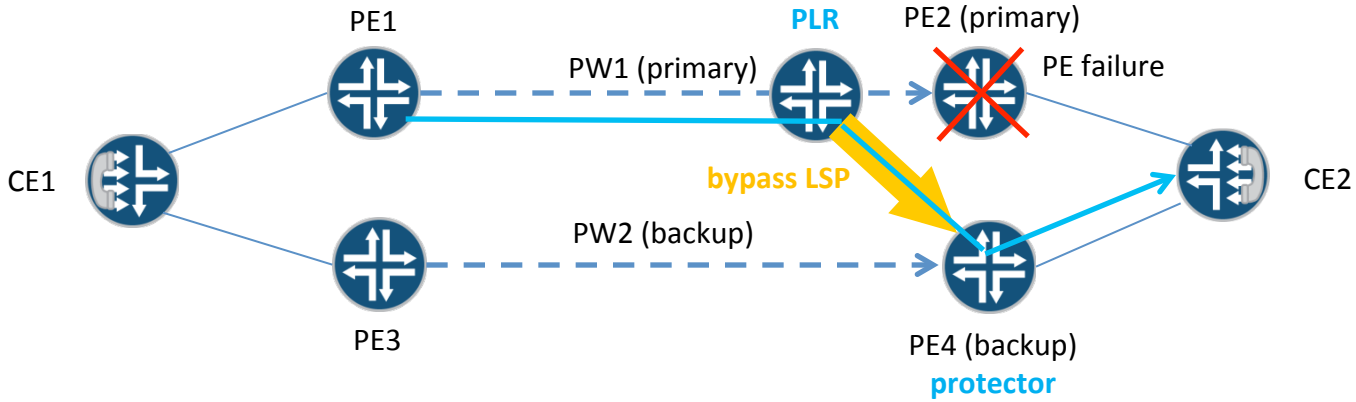

- PLR is the penultimate hop of PW transport LSP.
- If RSVP is the signaling protocol, PLR should
  - Signal bypass LSP based on "node protection desired" flag in Path message of transport LSP.
  - Report node protection available and in-use status in RRO of Resv message.
  - Node protection FRR for LSP egress.

Protection Models

1. Co-located protector PE
   - Protector PE is a backup PE, co-located with primary PE on a segment.
   - Has a direct connection to CE via a backup AC.
   - PW label popping.
   - Context ID allocation may typically be per {primary PE, protector PE}.

Protection Models (cont.)

2. Centralized protector PE
   - One PE protects all segments for all primary PEs.
   - May not be a backup PE on all segments.
   - For a segment with no direct connection, it must:
     - Learn {Protection FEC, label, context ID} from a backup PE via LDP.
     - Swap primary PE's PW label to backup PE's PW label, and send packet to backup PE.
   - Only one context ID is needed per primary PE.

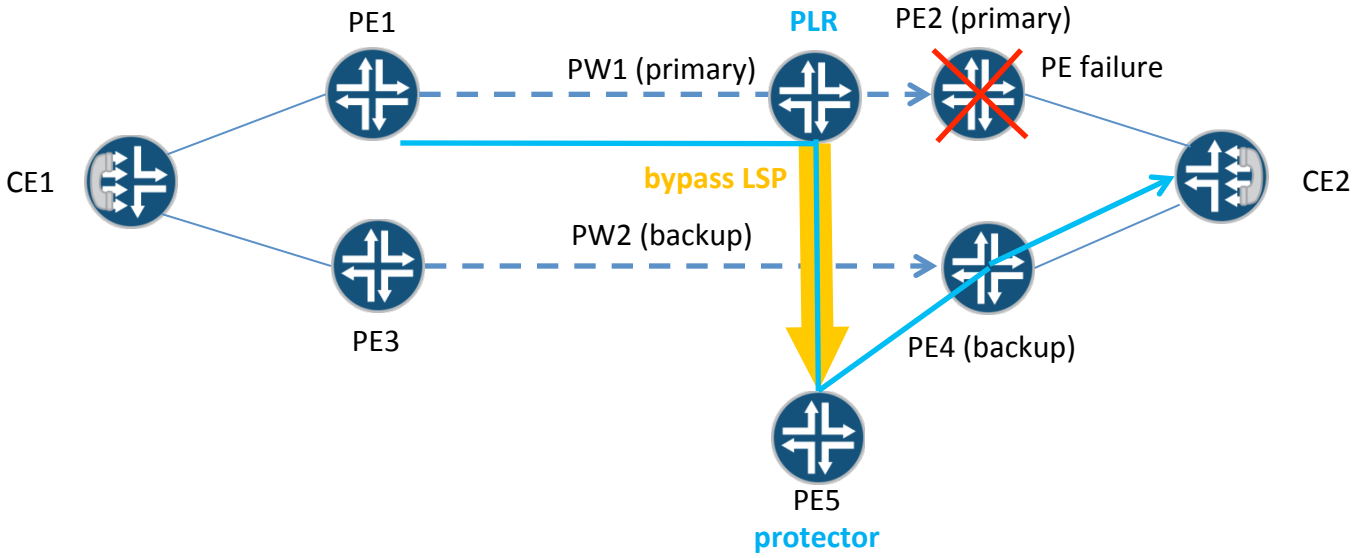

Next Steps

- Questions and comments?
- WG adoption?