Provider Bridge “Unlearn” Signaling

Norman Finn

“802” Problem Statement (1)

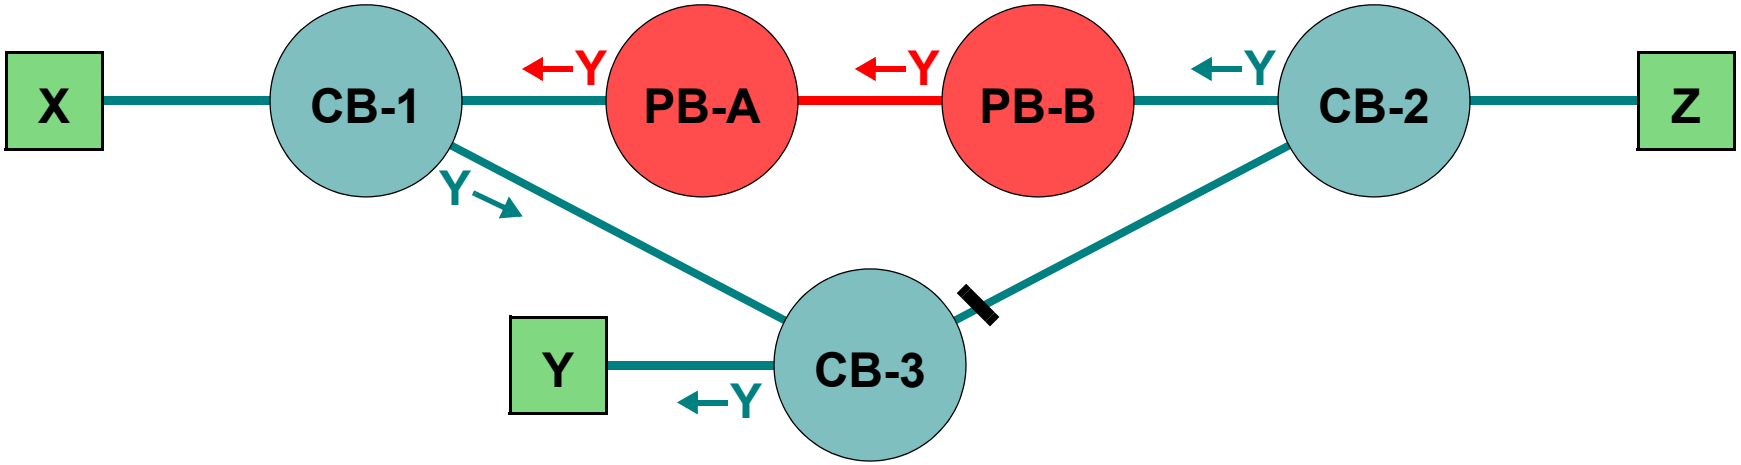

- Customer bridges 1-3 connected through provider bridges A and B, and also outside Provider’s network.
- Y is conversing with both X and Z, so all three MAC addresses are learned everywhere (only Y shown).
- PB-A will filter any frame from CB-1 addressed to Y.

“802” Problem Statement (2)

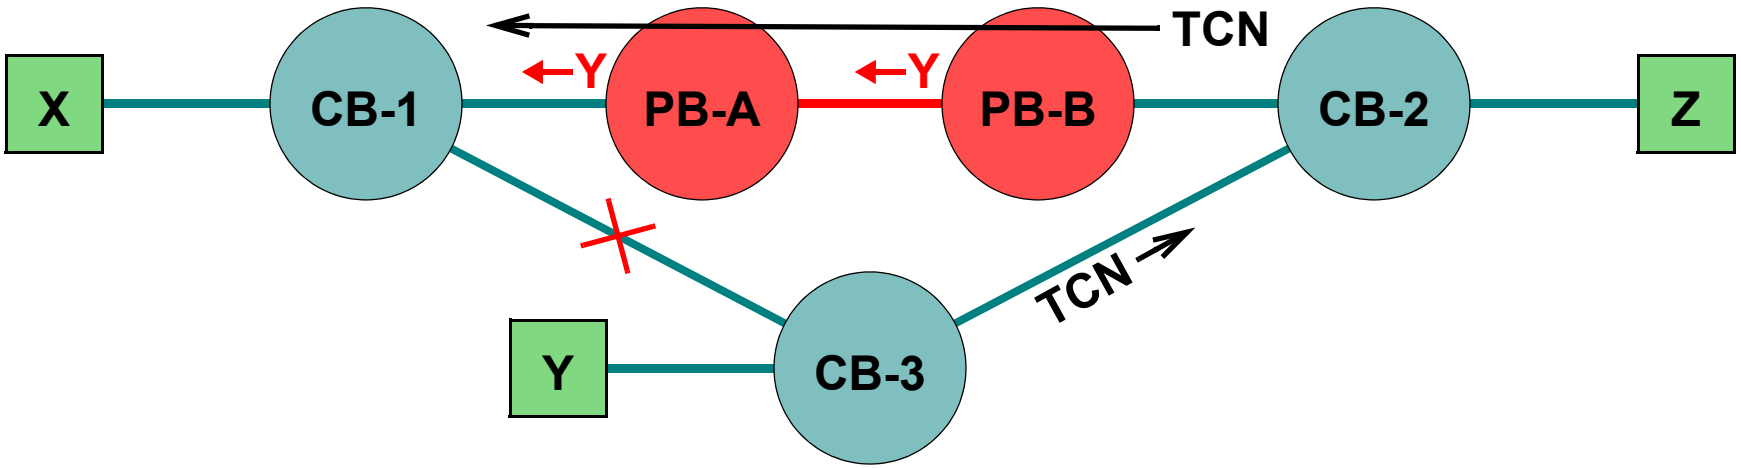

- Customer’s link fails. Customer’s RSTP reconfigures.
- Customer’s bridges see TCN and forget Y.
- Provider’s bridges do *not* see TCN or forget Y.

“802” Problem Statement (3)

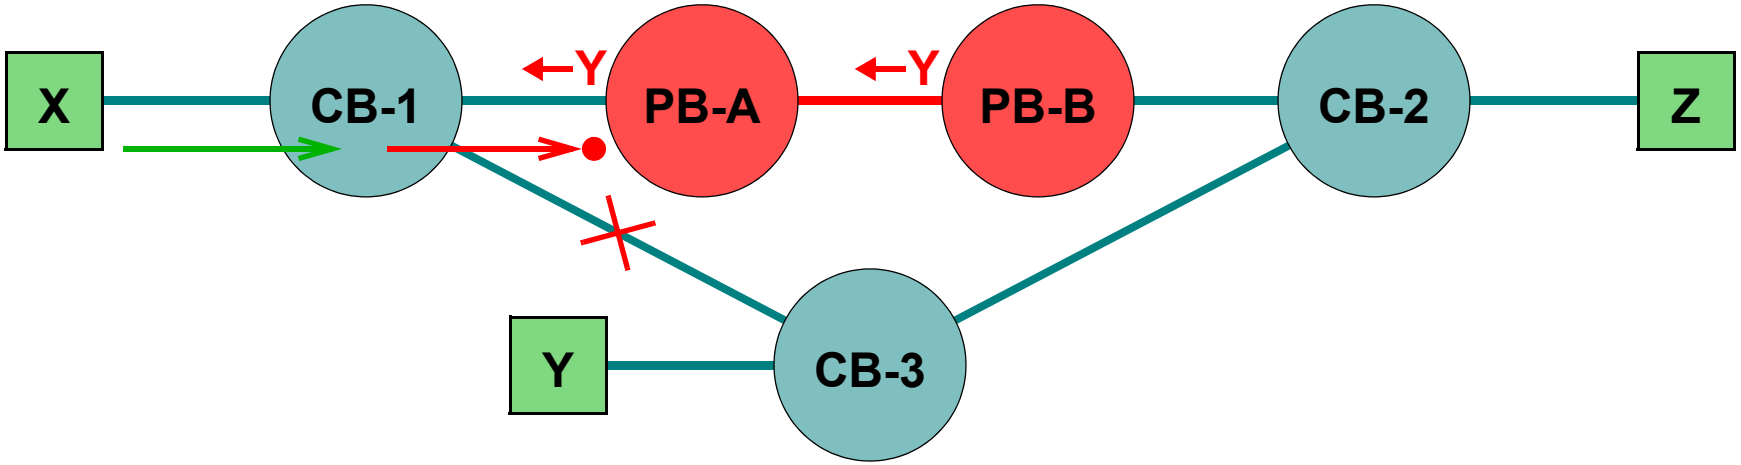

- Station X now transmits to Y. CB-1 does not know where Y is, so it floods frame to PB-A.
- PB-A thinks it knows where Y is, so it filters the frame.
- X cannot transmit to Y until PB-A’s FDB entry for Y times out.

“IETF” Problem Statement (1)

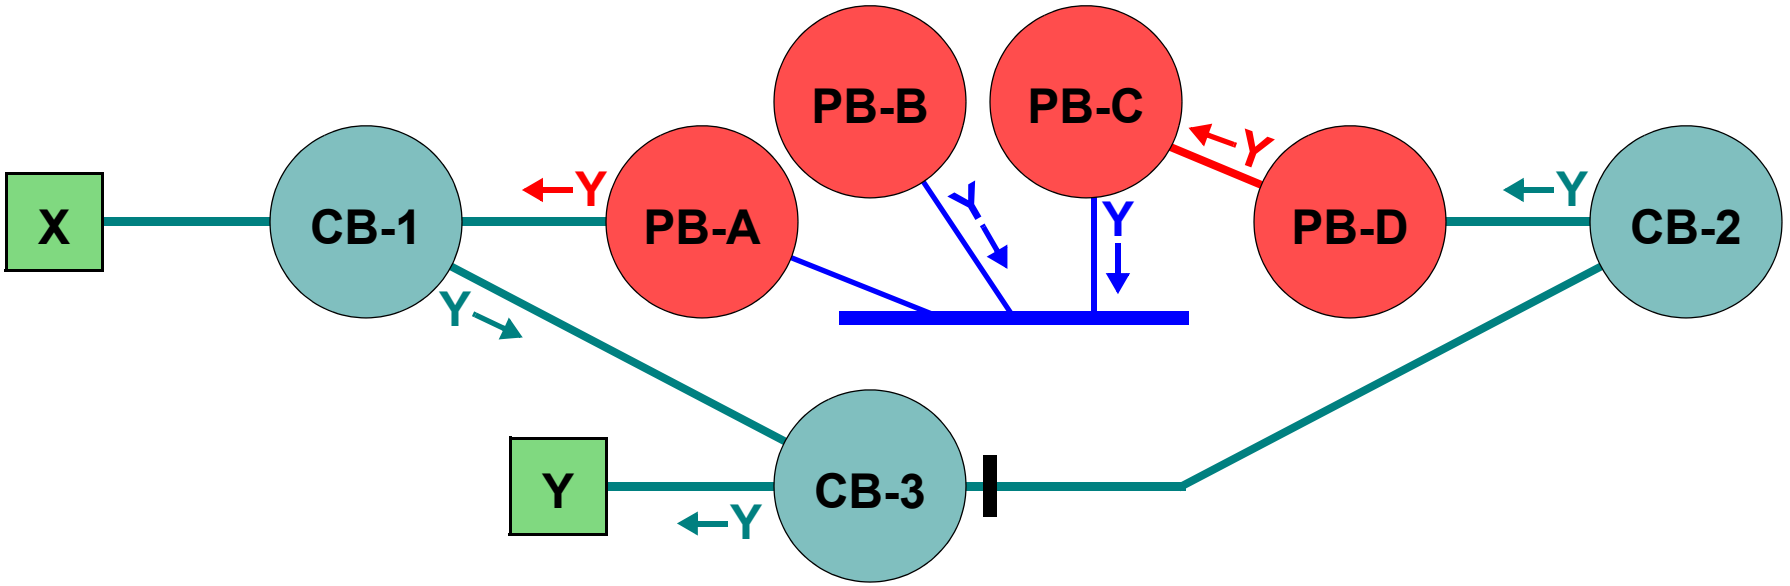

- **Same problem, though some Provider bridges may be interconnected via “Emulated Ethernet Over MPLS”.**
- **EEoMPLS utilizes a full mesh of Pseudo-Wires.**
- **Provider bridges’ Emulation Functions learn and forget {MAC Address, Pseudo-Wire} associations.**

“IETF” Problem Statement (2)

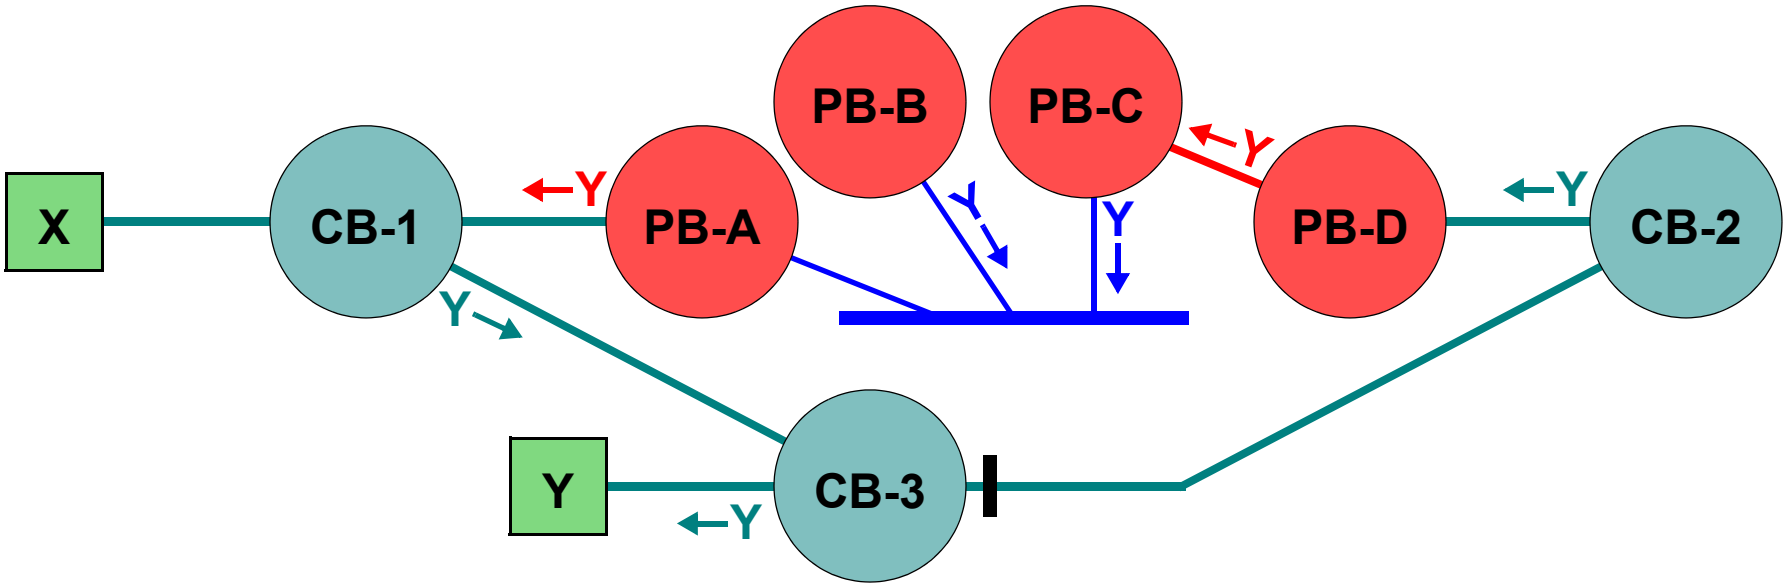

- **(Yes, the IETF more-or-less understands that an “Emulation Function” must snoop on the Provider’s BPDUs, so that it can forget {MAC Address, Pseudo-Wire} associations, when appropriate.)**

Customer BPDU "Snooping": Plan A

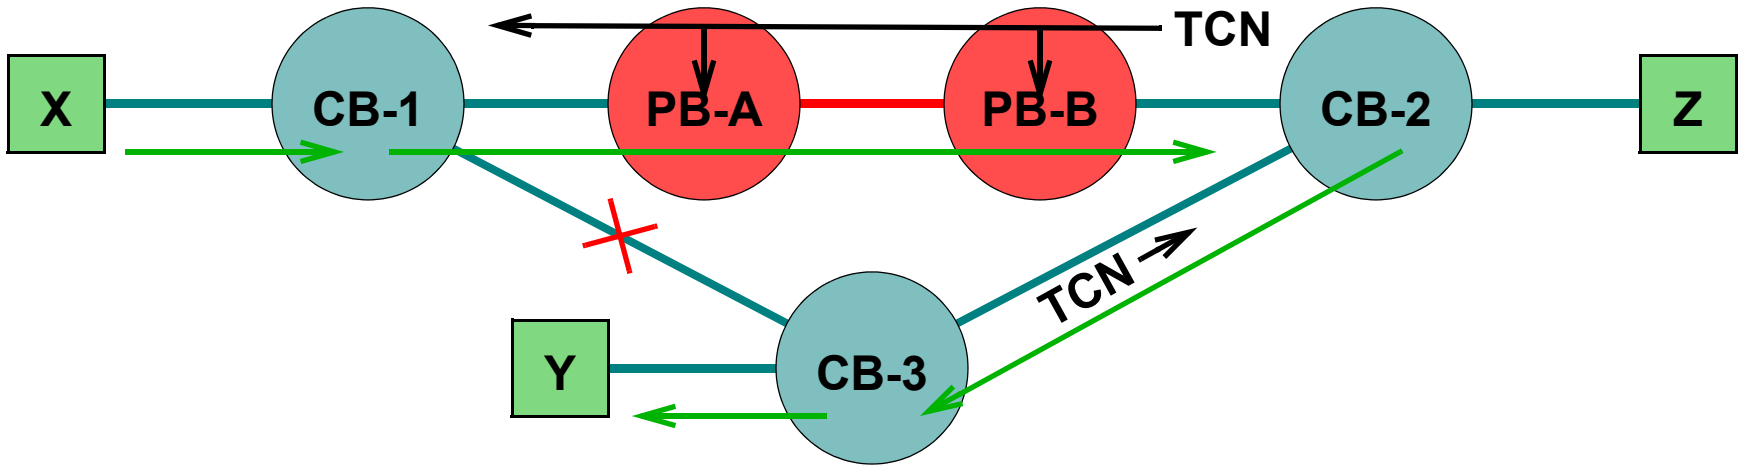

- If all Provider bridges "snoop" on Customer TCNs, and treat them just the same as Provider TCNs, then PB-A and PB-B would forget MAC address Y, and next frame from X to Y would be flooded properly.
- Of course, Provider ignores rest of Customer's BPDU.

Problems with Plan A

- **Customer may run any number of standard and/or proprietary spanning tree algorithms, all of which must be understood by all Provider bridges.**
- **Customer BPDUs must be snooped by every Provider bridge on at least every P-VLAN that allows the "outside" connections that cause the problem.**
- **If Provider is running two MSTP instances for 4K customers, it is very wasteful to forget 2K customers' MAC addresses because one customer issued a TCN.**
- **Provider bridges in different "Islands" do not see each others' Provider BPDUs. Even if they did, they would have to forget *all* customers' MAC addresses.**

Customer BPDU Snooping: Plan B

- **Same as Plan A, but BPDU Snooping is done only by "UNI Wart" at the edge of the Provider network.**
- **In addition to transmitting the Customer BPDU, the Provider Edge Bridge generates a Provider TCN, in effect "promoting" the Customer TCN to the Provider's network.**
- **This isolates odd customer protocols, and eliminates snooping on non-UNI ports, so scales better.**
- **However, one customer's TCN still causes 2K customers' MAC addresses to be forgotten.**
- **The "Inter-Island" problem remains.**

Customer BPDU Snooping: Plan C

- **Same as Plan B, but run 4K MSTP instances, using 64 MSTP BPDUs per transmission.**
- **This allows one Customer BPDU to trigger "unlearning" in just one Provider VLAN.**
- **It's expensive: BPDU load is multiplied by 64 over having just a few MSTIs.**
- **The "Inter-Island" problem remains.**

Snooping + Unlearn Message: Plan D

- **Provider Edge Bridge does the Snooping, as in Plan C, but instead of generating a Provider TCN, Edge Bridge generates an "Unlearn Message".**
- **The Unlearn Message is an in-band packet, using the (Provider's) "All Bridges" multicast MAC address, tagged with the appropriate P-VLAN. It stays within the Provider's network; it is not emitted from a UNI.**
- **The Unlearn Message causes each receiving Provider bridge to forget the same MAC addresses learned from the same interfaces as would a Provider TCN, but *only* for the one P-VLAN, *not* for all P-VLANs on the MSTI.**

All Problems Solved?

- **Snooping is isolated to the UNI Warts.**
- **MAC address forgetting is done on a per-P-VLAN basis, so is not wasteful.**
- **We do not need thousands of MSTIs.**
- **A Multicast Unlearn Message propagates through Provider's network much faster, and with less burden, than a hop-by-hop Provider TCN.**
- **An Unlearn Message is meaningful across Inter-Island boundaries, so it solves the "Inter-Island" problem.**
- **The Provider Bridges' "Emulation Functions" must, of course, snoop on Unlearn Messages as well as P-TCNs.**