

draft-malhotra-bess-evpn-unequal-lb-00

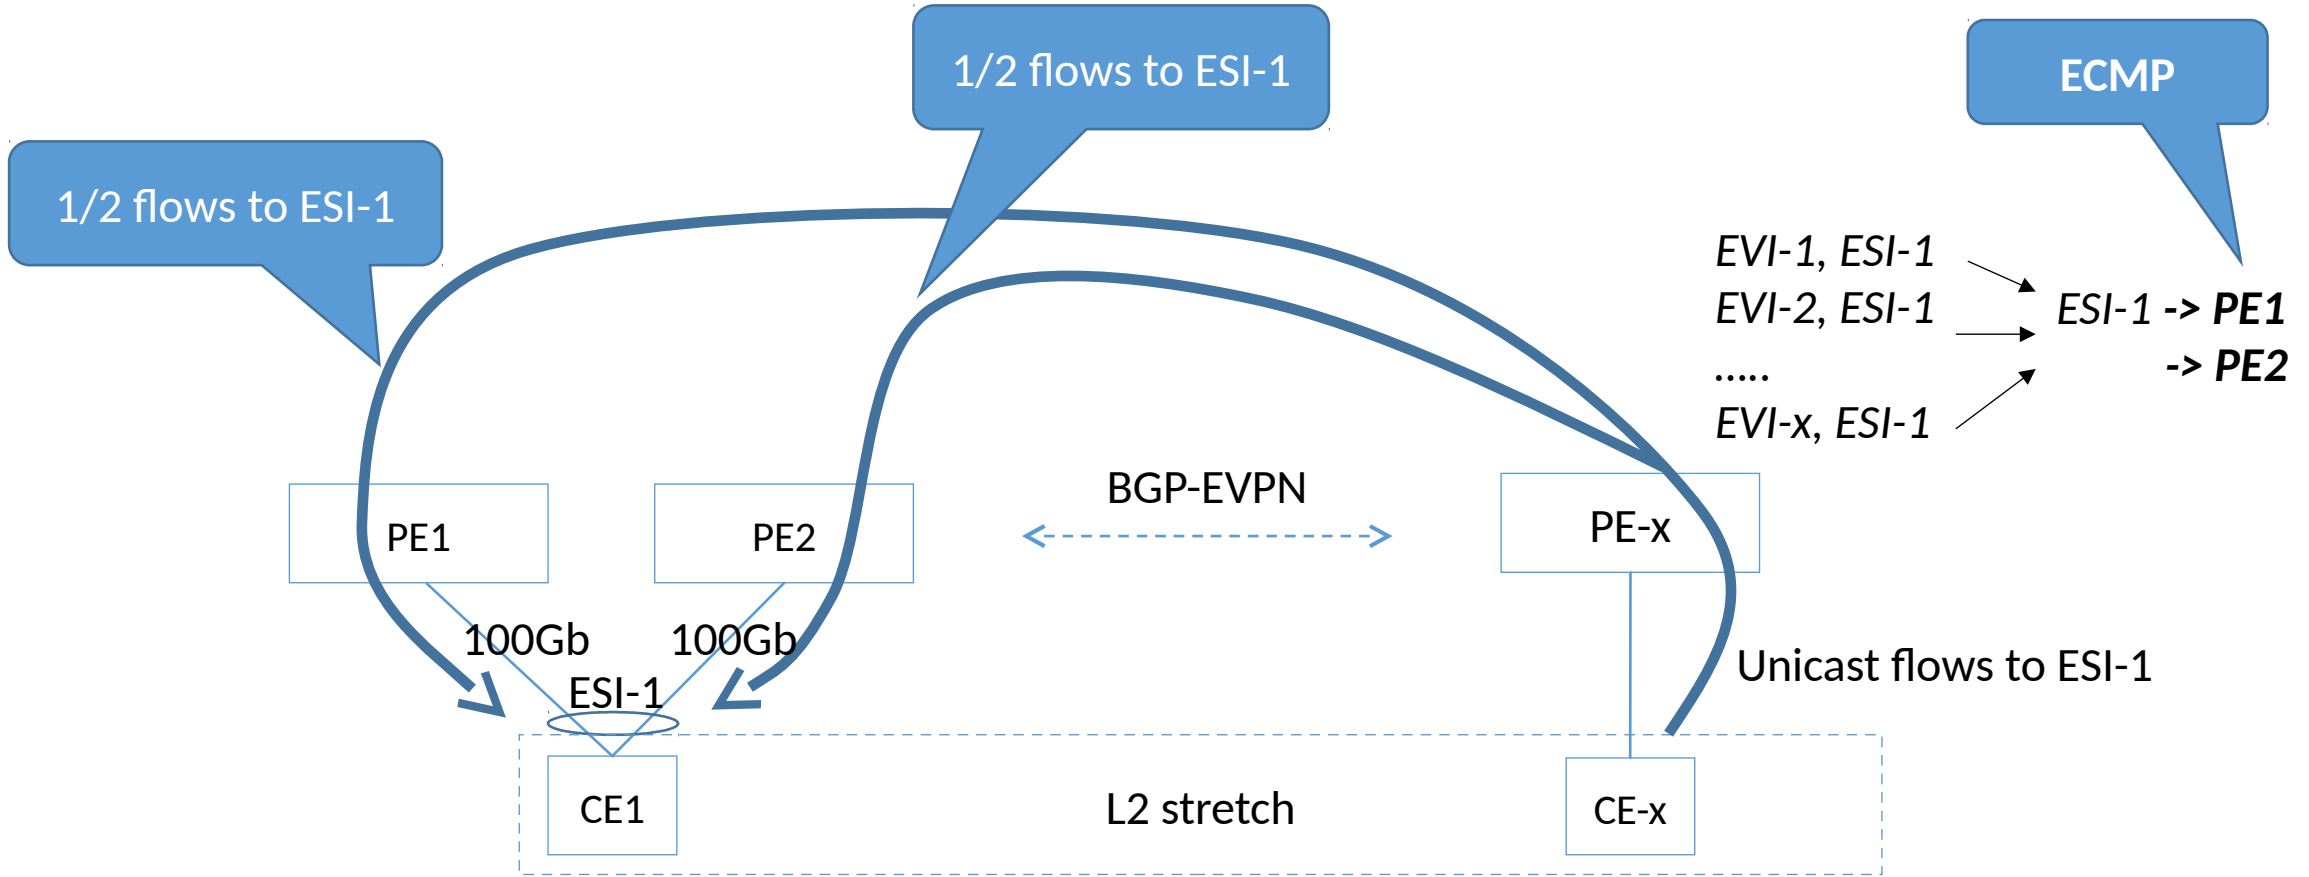
N. Malhotra (Cisco), S. Thoria (Cisco), A. Sajassi (Cisco),
A. Lingala (AT&T)

IETF 100, November 2017
Singapore

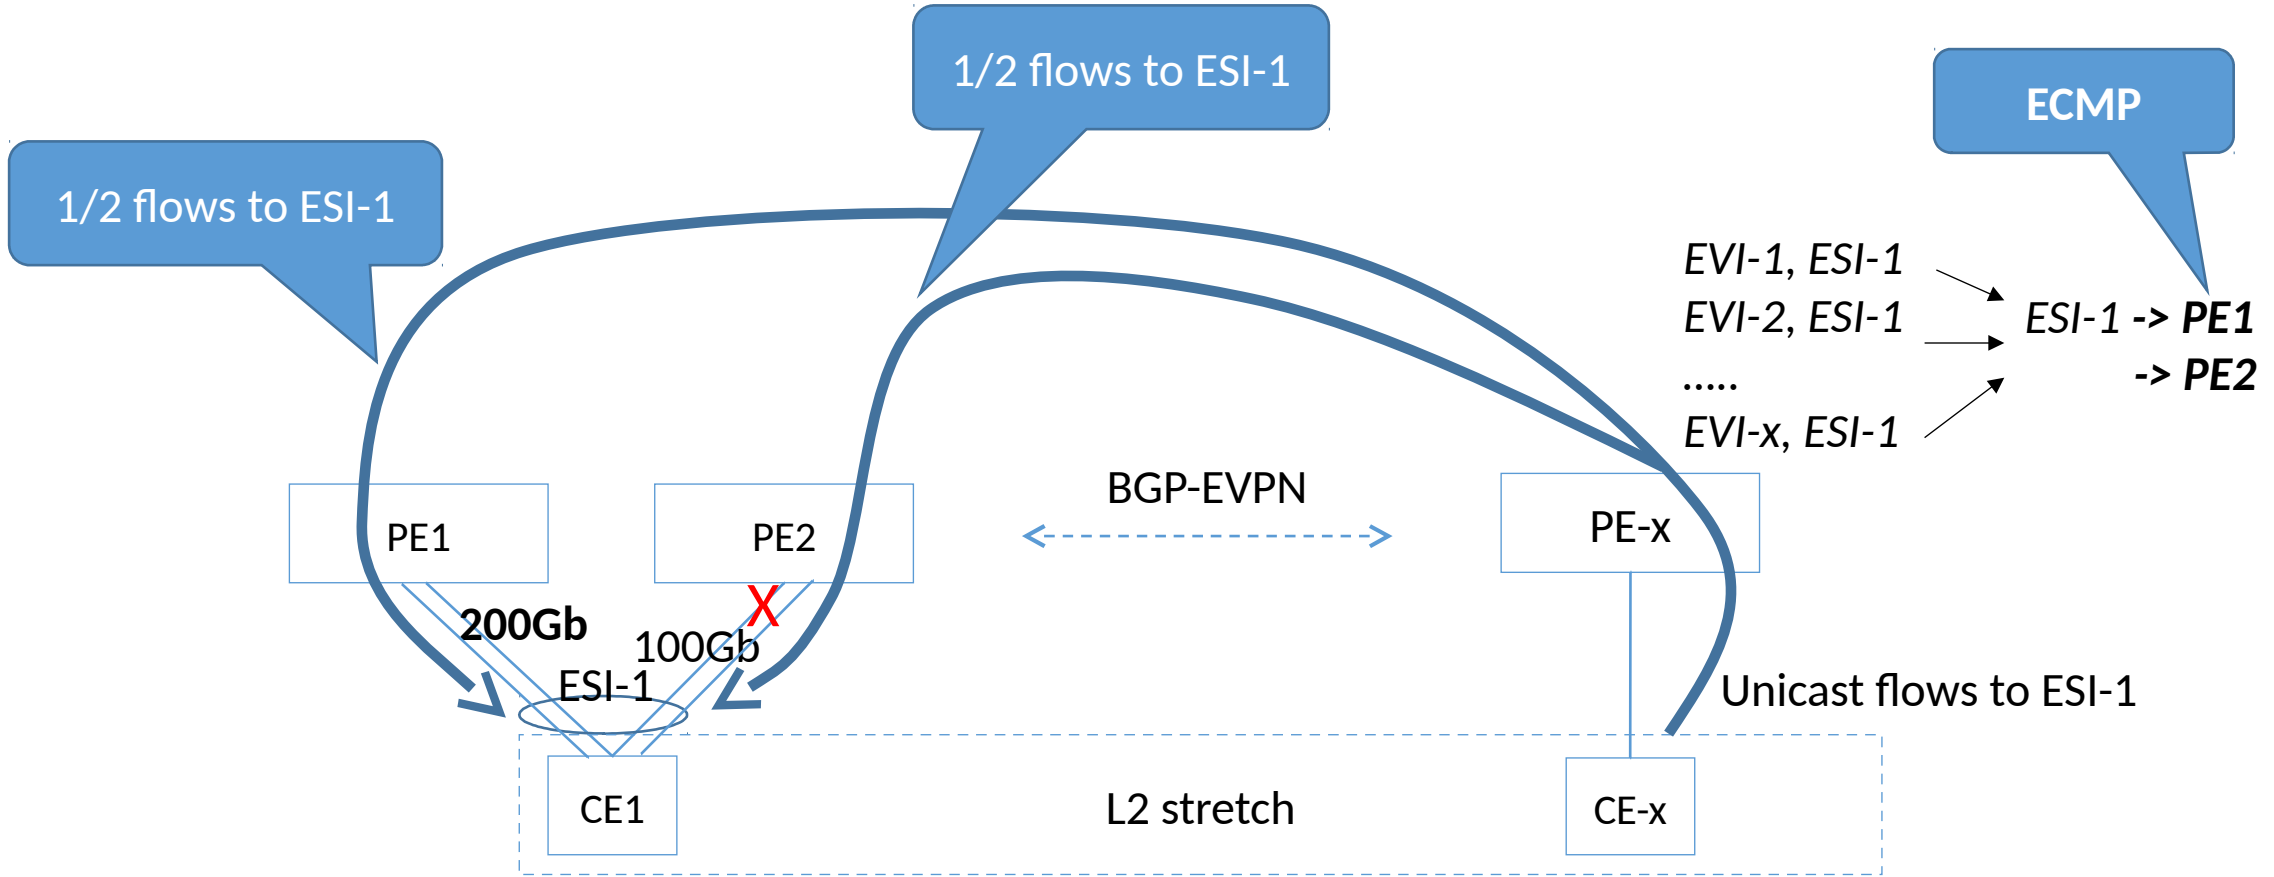
Prior Art

- RFC 7432 EVPN All-Active Multi-Path procedures (aliasing, mass withdraw)
 - Enable overlay Equal Cost Multi-Path
 - Overlay flows load-balanced “equally” across a set of all-active multi-homing PEs
- RFC 7432 EVPN “per-service” DF election
 - Per-service DF role ”equally” distributed across a set of multi-homing PEs

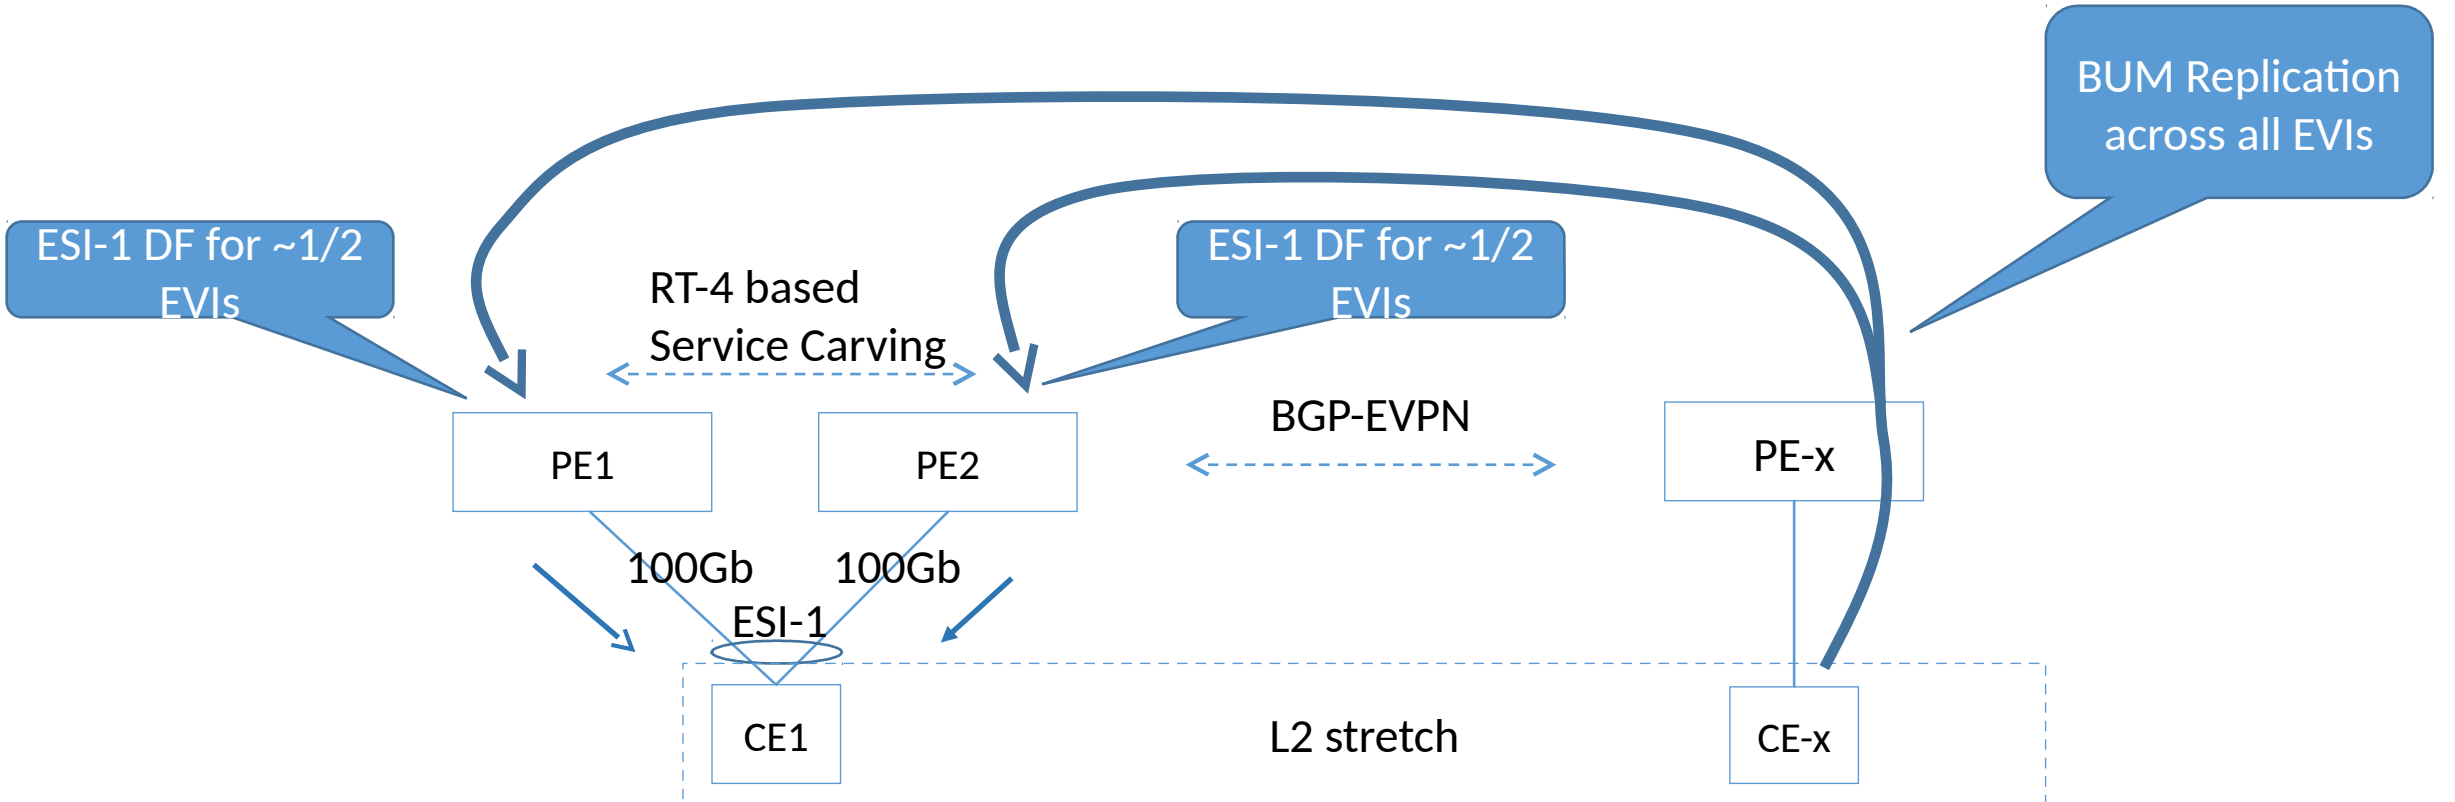
Unicast ECMP



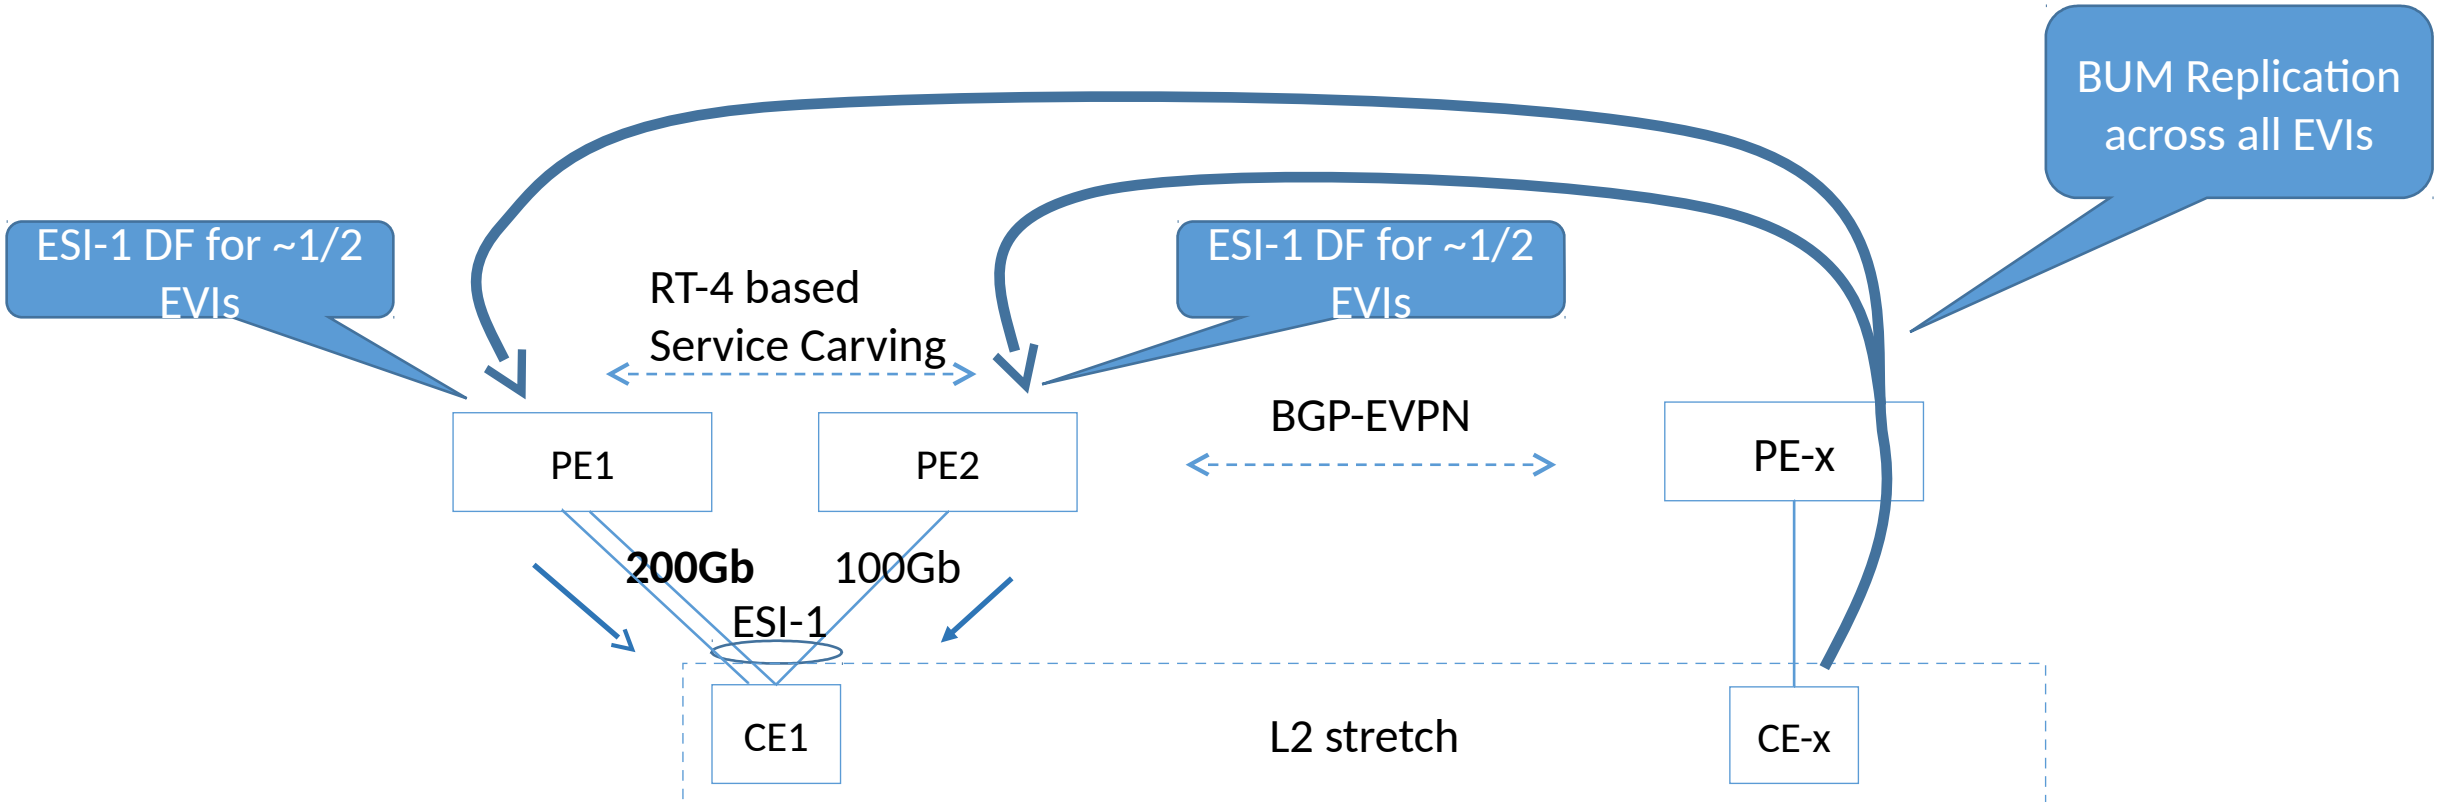
Sub-optimal Unicast ECMP – asymmetric access BW distribution



BUM Flows – DF Service Carving



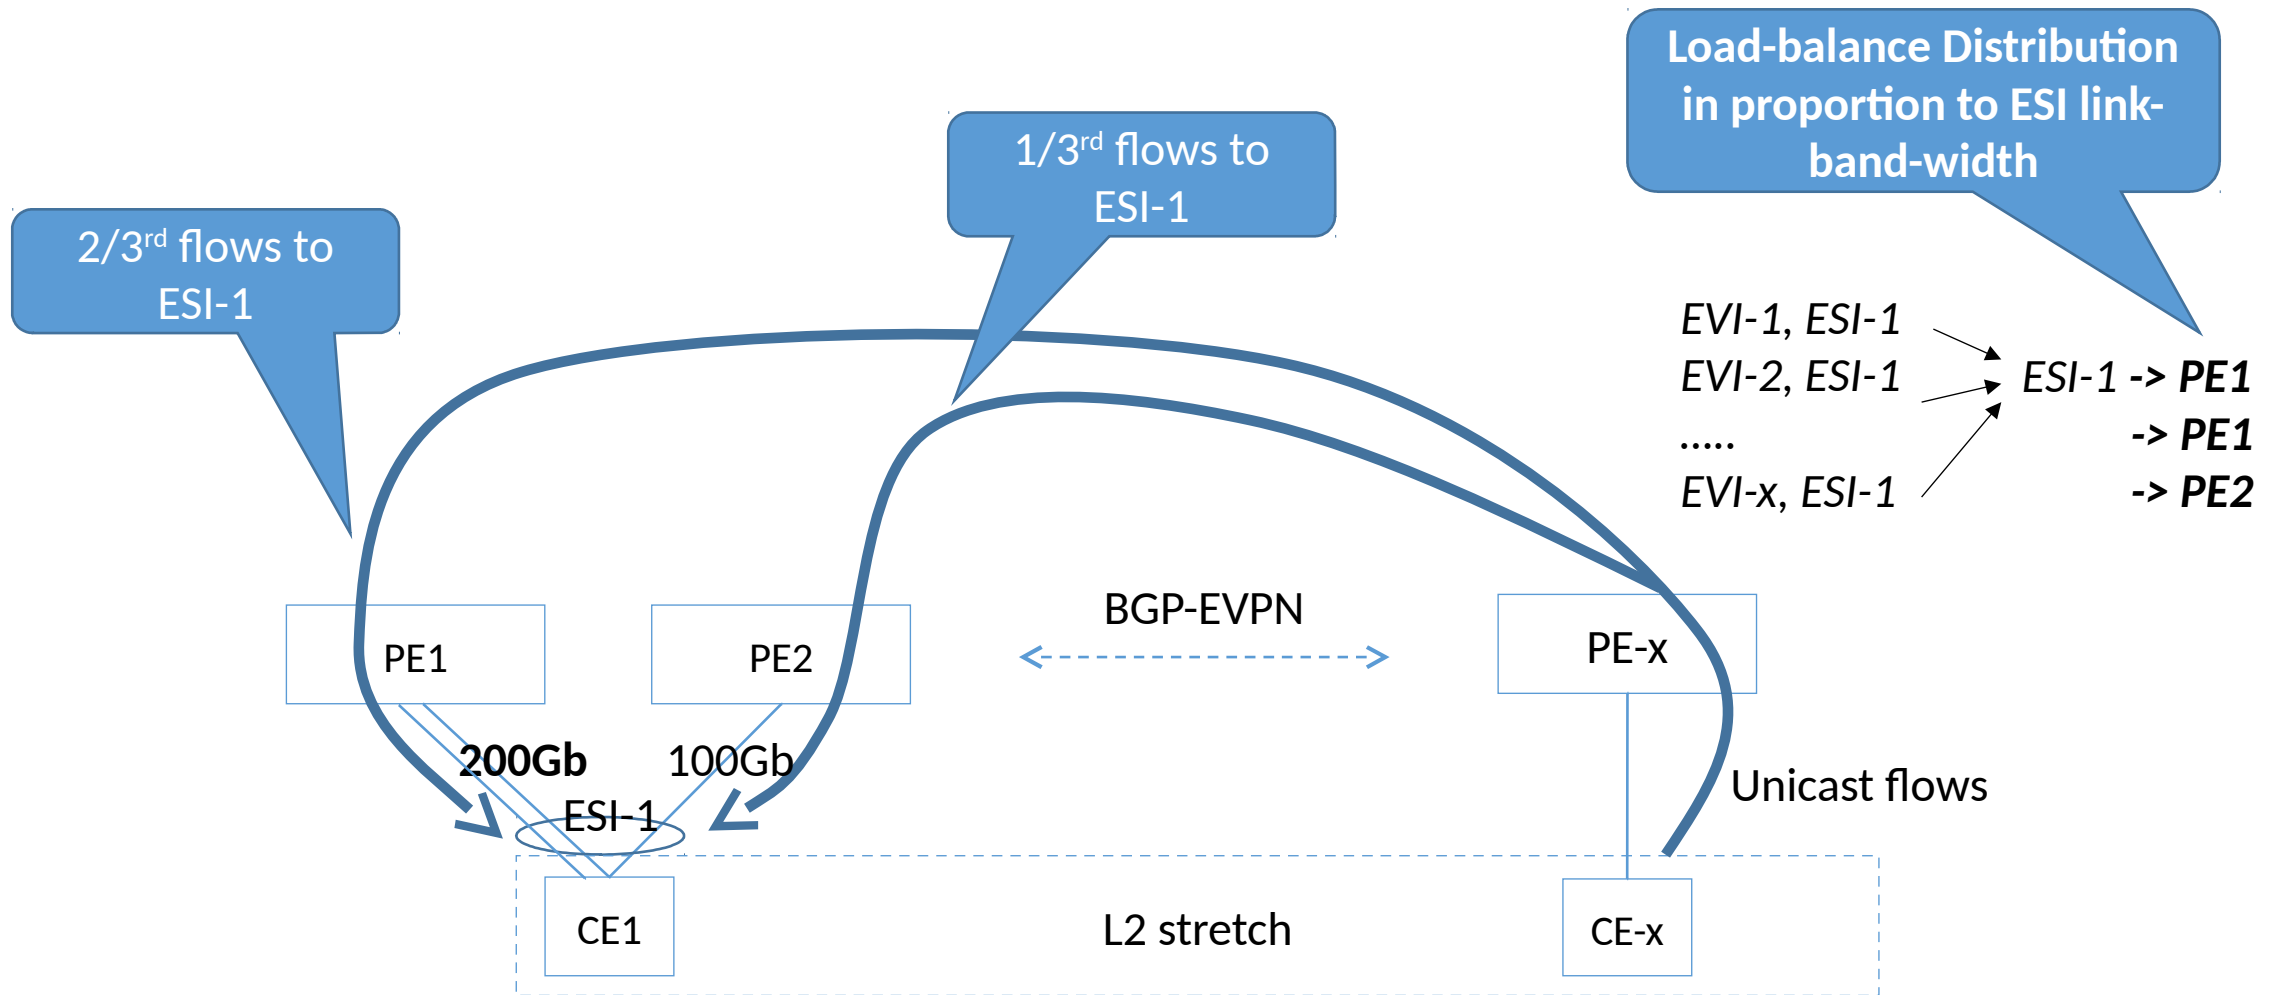
Sub-optimal BUM Flows – DF Service Carving – asymmetric access BW



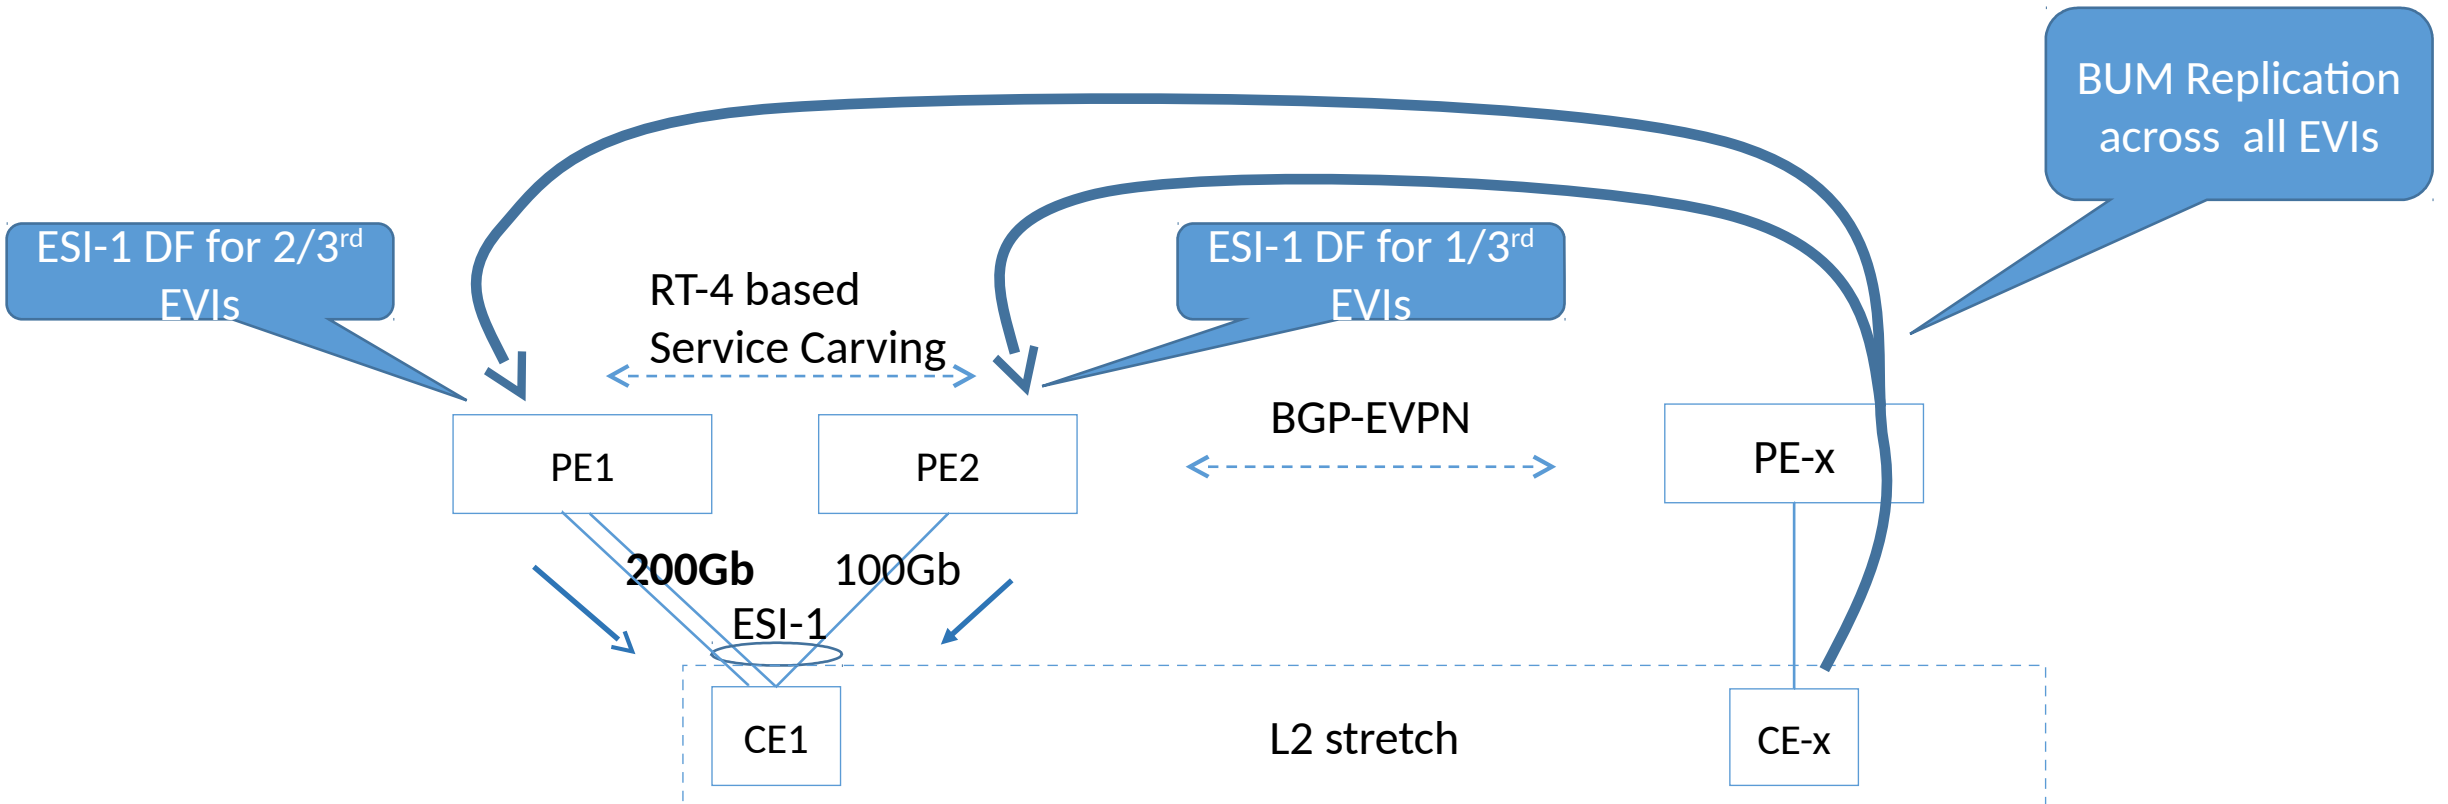
Objective

- Overlay flows load-balanced across multi-homing PEs in proportion to ESI bandwidth
- Per-service DF role (service carving) distributed across multi-homing PEs in proportion to ESI bandwidth

Objective – Unicast LB in proportion to ESI link-bandwidth



Objective - DF Service carving in proportion to ESI link-bandwidth



Solution

Unicast Traffic Load-Balancing

- Local PE
 - Advertise per-ESI link-band-width attribute as part of per-ESI EAD RT-1
 - <https://tools.ietf.org/html/draft-ietf-idr-link-bandwidth-06> (BGP Link Bandwidth Extended Community)
- Remote PE
 - ESI Path-list computed in proportion to received link-band-width attribute from each PE

DF Service Carving

- TBD – to be reconciled with draft-ietf-bess-evpn-pref-df-00 (Preference-based EVPN DF Election)

Next Steps

- DF election procedure to be reconciled with draft-ietf-bess-evpn-pref-df-00
- BGP Link Bandwidth Extended Community draft has expired (under discussion)
- Collaboration with additional co-authors in progress for rev-1

Weighted Multi-Path Procedures for EVPN All-Active Multi-Homing

(draft-malhotra-bess-evpn-unequal-lb-00)

Thank You

Neeraj Malhotra, Samir Thoria, Ali Sajassi (Cisco)
Avinash Lingala (AT&T)