End to end BGP based VPNs for the European R&E community

Jean-Marc Uzé

Juniper Networks

juze@juniper.net

TERENA, Zagreb, May 20, 2003

What Is a VPN?

- A private network constructed over a shared infrastructure
  - Virtual: not a separate physical network
  - Private: separate addressing and routing
  - Network: a collection of devices that communicate

- Deploying VPNs in the 1990s
  - Provider-provisioned VPNs with ATM PVC
    - E.g.. JAMES, TEN-155 MBS, several NRENs…
  - CPE-based VPNs with IP tunnels (GRE, IP-IP)
    - E.g.. Mbone, 6Bone…

- Deploying VPNs in the 21st Century
  - Uses IP Infrastructure
  - Provider-provisioned VPNs and CPE-based VPNs
  - One VPN Model Cannot Fit All Requirements!

Virtual Private Network Services

- L3 IPv4/IPv6 VPNs (RFC 2547)
  - Application example: support multiple communities in MAN or Regional Network
    - Network isolation
    - Manage exterrnal access (NREN, IP commodity)
- L2 point-to-point VPN (L2 VPNs)
  - Application examples: support National/European projects that require dedicated L2 infrastructure –or- share an access loop with different services
    - Pt-to-pt Layer 2 circuits
      - FR DLCI on POS access links
      - ATM PVC on ATM access links
      - VLAN on Ethernet access links
    - IP interworking support to mix L2 access technologies
- L2 multipoint-to-multipoint VPN (VPLS)
  - Application example: Virtual Lab Service
    - Ethernet multipoint access
    - Support of broadcast and MAC learning

End to End VPNs across multiple domains (AS)

- R&E end-users of a VPN are rarely connected to the same IP domain
- IP is everywhere
  - Provides any-to-any connectivity
  - A VPN service using IP infrastructure do not require a new/separate network
    - Would be to costly and to complex
- Inter-AS VPN Service is an inevitable issue in R&E networks, but not specific to this environment
  - Also required by commercial ISP:
    - That have a big network constituted of multiple ASs
    - in a consolidation process (new AS acquisition)
    - Carrier supporting VPN across different ISPs
  - Defined in 2547bis for any BGP-based VPN (L3, L2 and VPLS)

Agenda

- The BGP/MPLS VPN Toolkit
- Inter-AS/Inter-provider operations

Network Reference Model

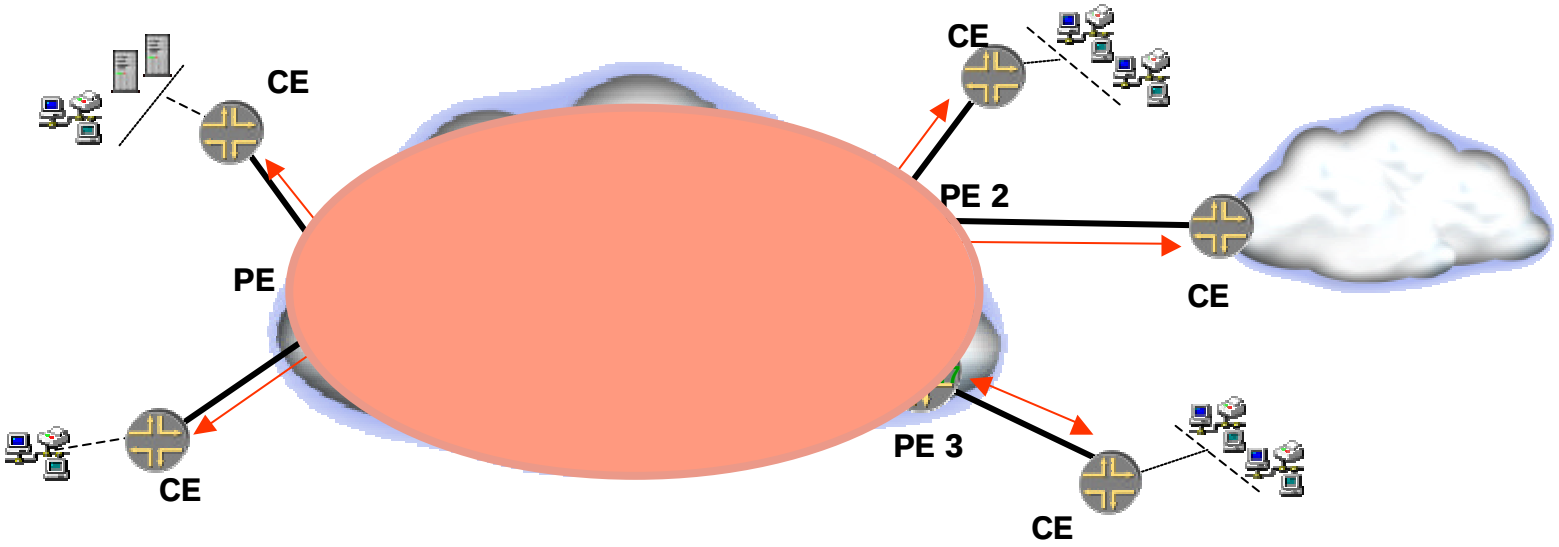

- **Addressing (loopback + interconnection)**
- **IGP (IS-IS, OSPF v2/v3)**
- **iBGP (Route Reflectors, confederation…) + EBGP**
- **Same Routing Information in all routers (P, PE)**

Requirements for scalable VPN Services

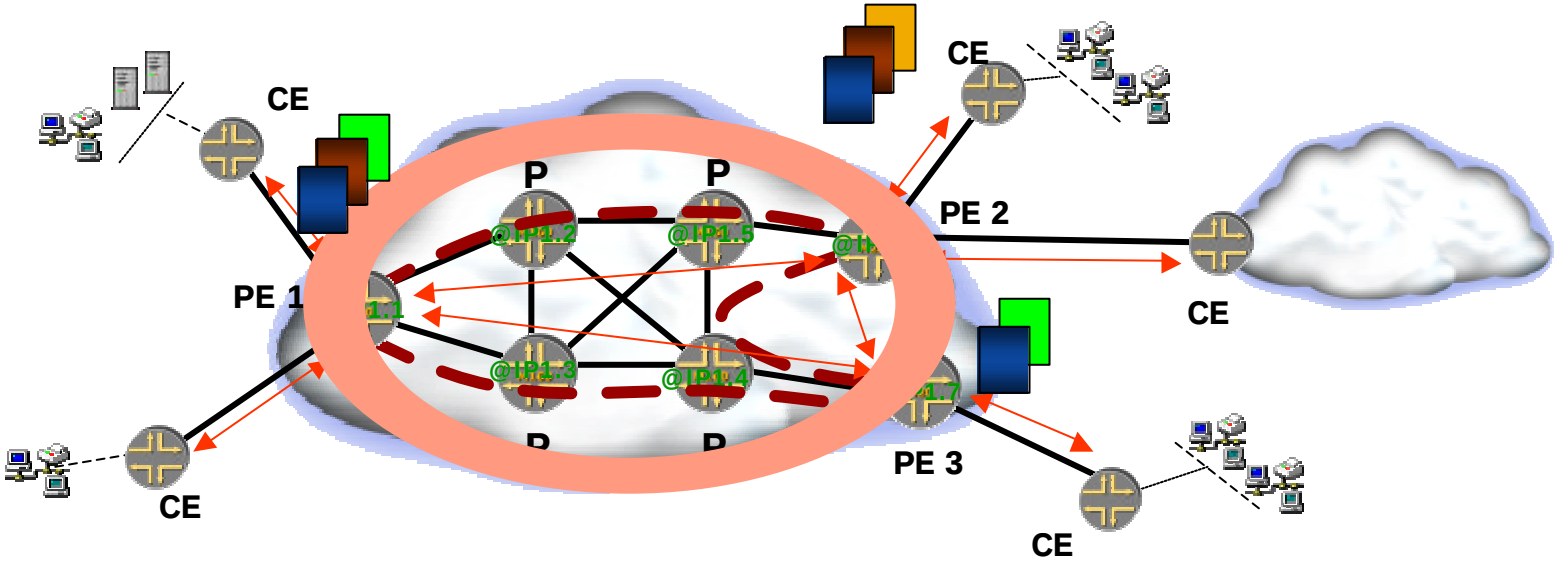

- **Distribute Routing and Forwarding information in the PEs**
  - **PE router has to maintain VPN information only for VPNs whose sites are directly connected to the PE router**
  - **P routers must be free of all the VPN routing information (v4, v6, L2 VPNs & VPLS)**
    - **Tunnels required between PEs**

VPN Service Components

- CE-PE : routing protocol or Layer 2 protocol
- Tunnel setup
  - Outer tunnel – PE to PE
    - MPLS tunnels: RSVP-TE, LDP (P are MPLS nodes)
    - IP tunnels: GRE, IPSec, L2TPv3 (P are IP nodes)
- PE-PE Auto-Discovery
  - which PEs are members of a given VPN
- PE-PE Signaling a demultiplexor
  - to which VPN (and, for Layer 2 VPNs, which source site) does a given packet belong
- PE: VPN Connection/Routing/Forwarding Tables

IPv4 VPNs

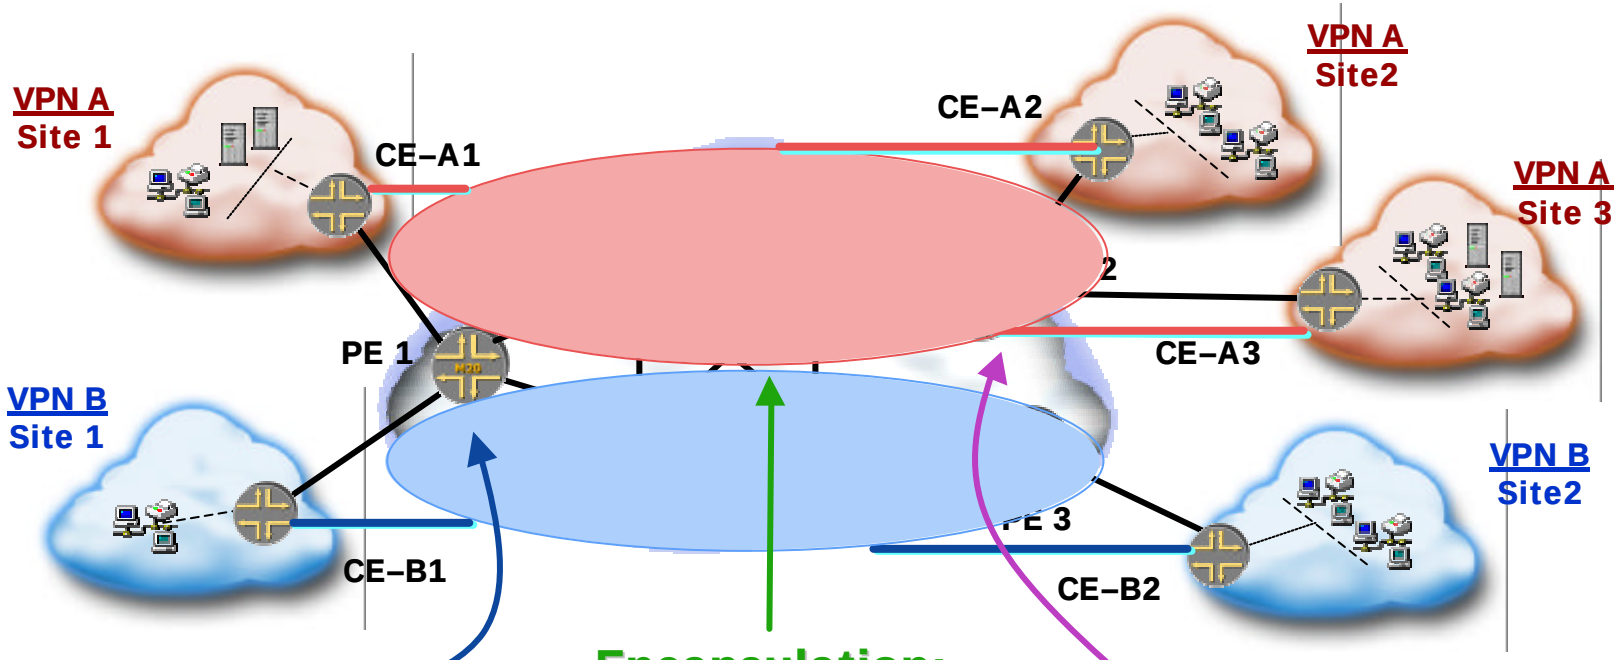

Encapsulation:
IPv4 over IP/MPLS

At ingress, IPv4 packet forwarded based on its IPv4 dest address and its associated VPN Routing and Forwarding Table

At egress, label is used to determine where to send the IPv4 packet

SP network acts as a dedicated IPv4 network

IPv6 VPNs

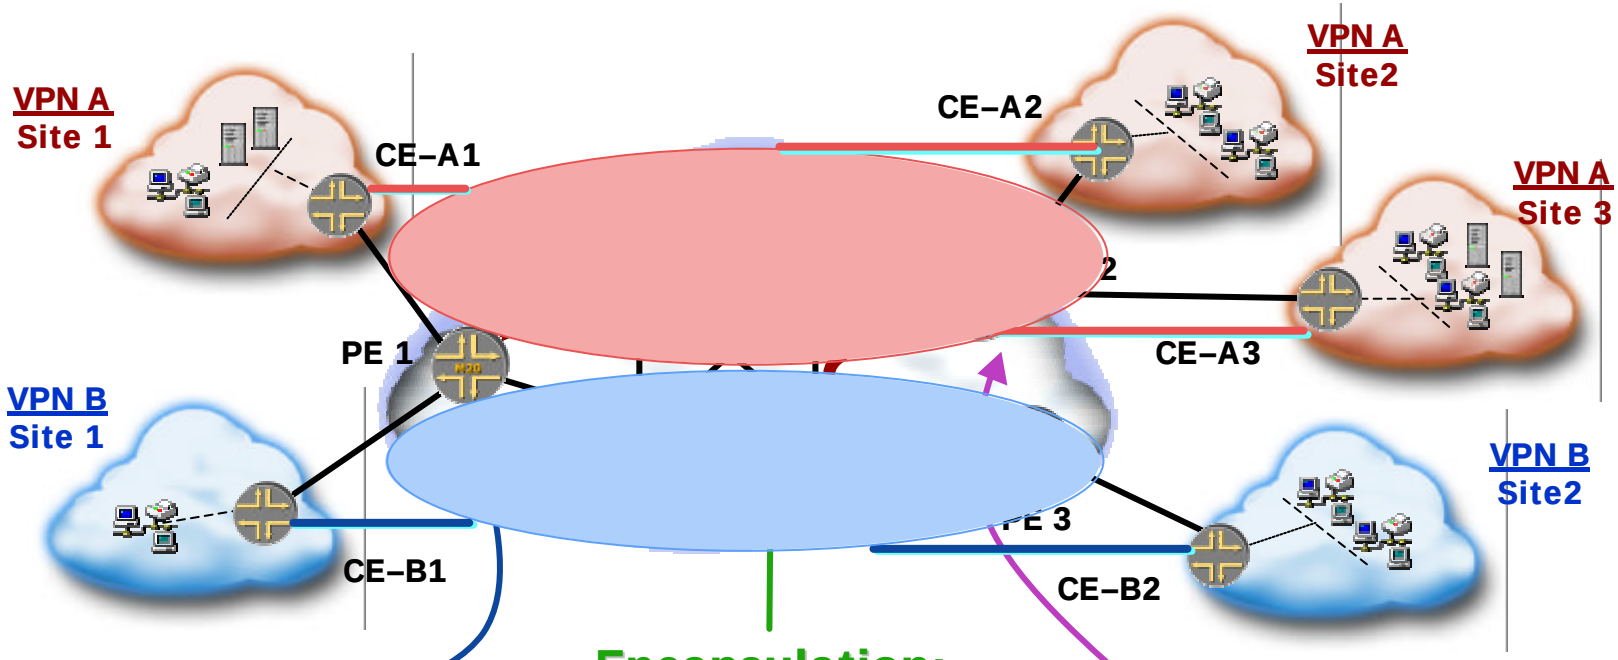

At ingress, IPv6 packet forwarded based on its IPv6 dest address and its associated VPN Routing and Forwarding Table

At egress, label is used to determine where to send the IPv6 packet

SP network acts as a dedicated IPv6 network

Point-to-point Layer 2 VPNs

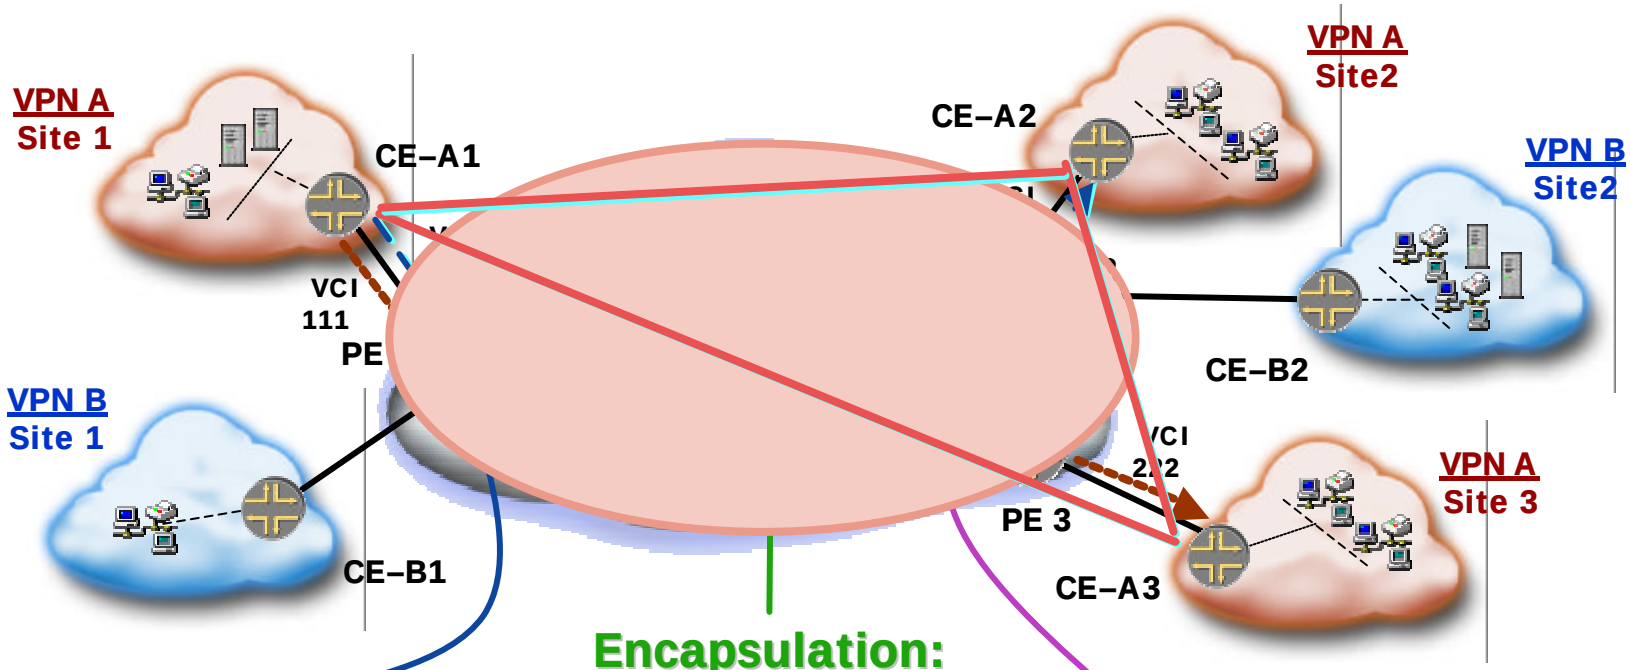

At ingress, L2 frame switched based on its layer 2 address

Encapsulation: Layer 2 over IP/MPLS

At egress, label is used to determine where to send frame

SP network acts as giant Layer 2 switch

IP Interworking (TCC)

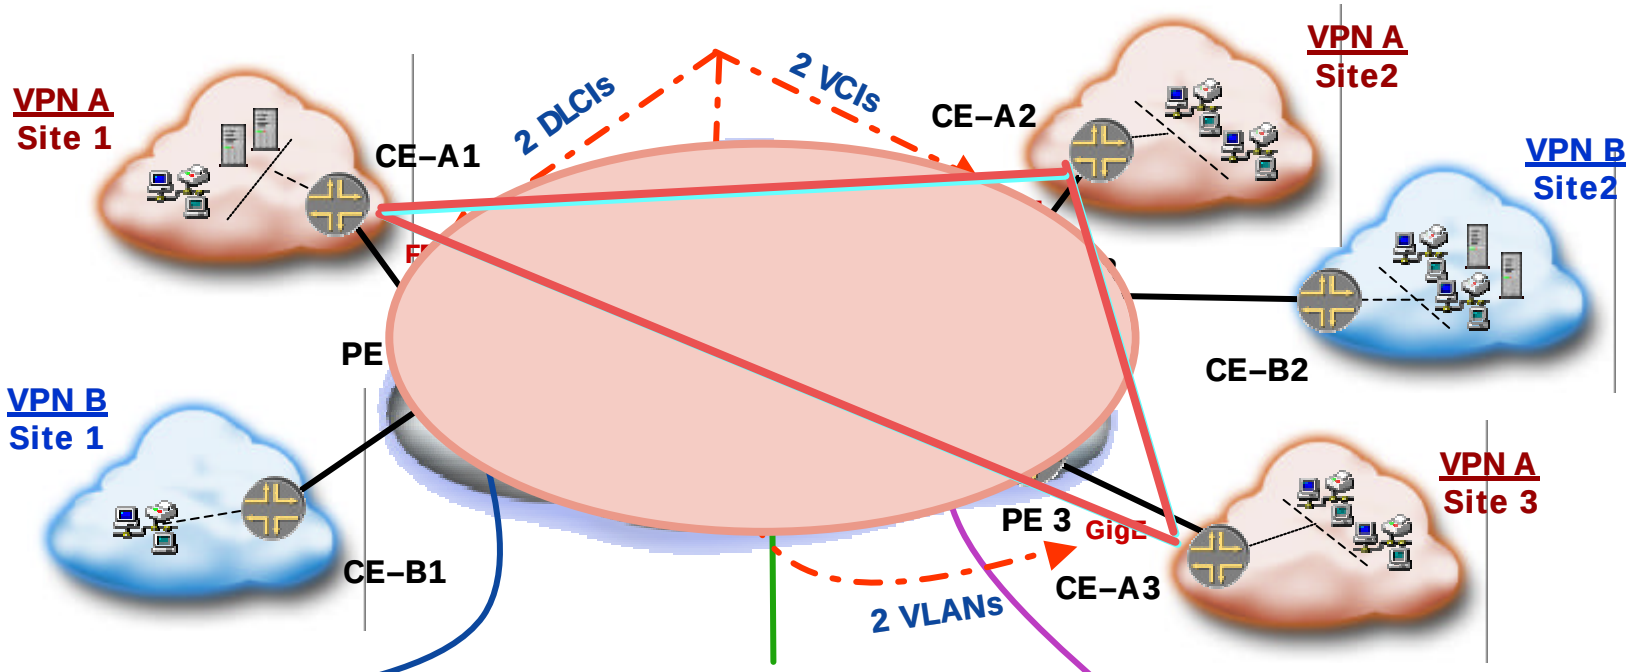

Ingress switches frame based on L2 address, strips L2 header, adds label stack to IP packet

Encapsulation: IP over IP/MPLS

SP network acts as a giant l/w switch

At egress, label stack is stripped, new L2 header added and packet is sent to CE

Virtual Private LAN Service

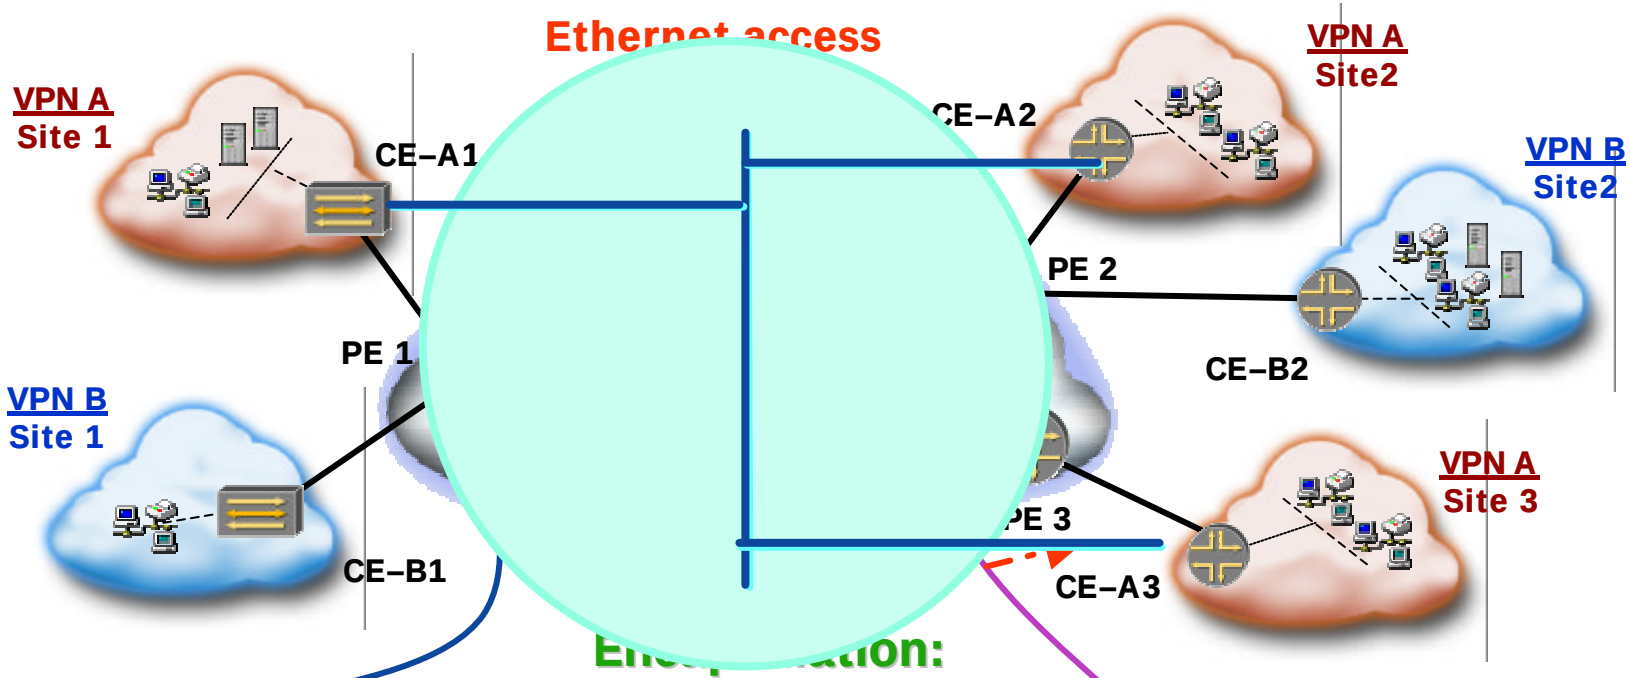

At ingress, Ethernet frame is switched (or flooded) based on dest MAC address, label stack added

SP network acts as a giant Ethernet switch

At egress, label stack is stripped, source MAC address learnt, and Ethernet frame is sent to CE

Basic RFC2547 operation

- Labeled Path between PEs -> outer-label distributed by LDP in the AS
- Outgoing Interface -> inner-label distributed by MP-iBGP

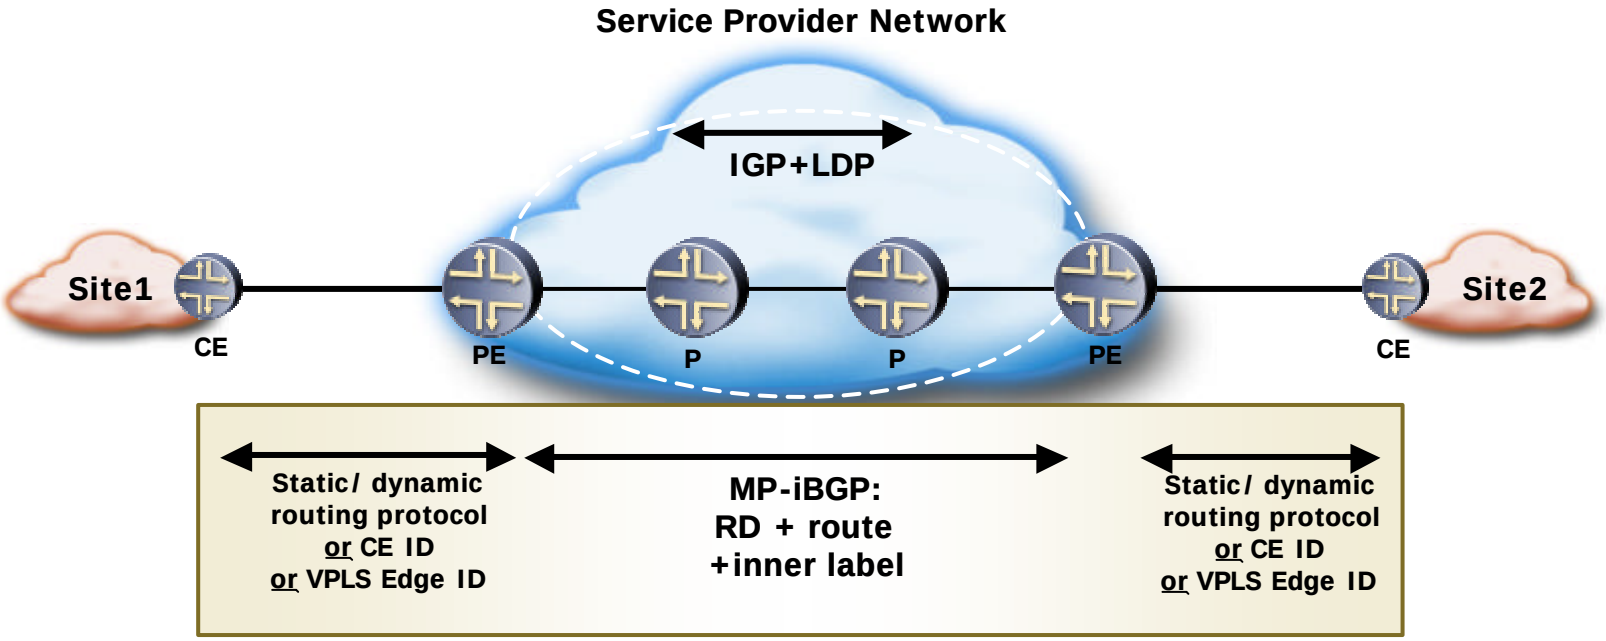

Forwarding state: basic RFC2547 VPNs

- Labeled Path between PEs -> outer-label distributed by LDP in the AS
- Outgoing Interface -> inner-label distributed by MP-iBGP

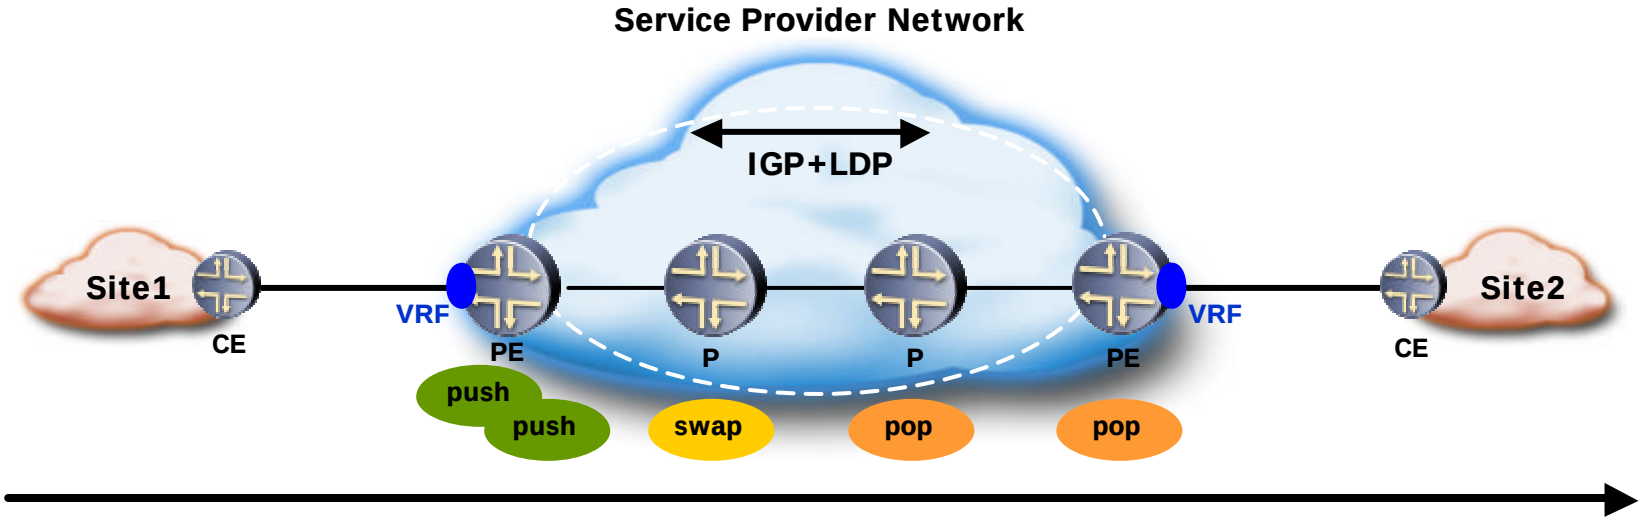

Agenda

- The BGP/MPLS VPN Toolkit
- Inter-AS/Inter-provider operations

VRF-to-VRF Connections at AS Border Routers

Inter-Provider Backbones Option A in 2547bis

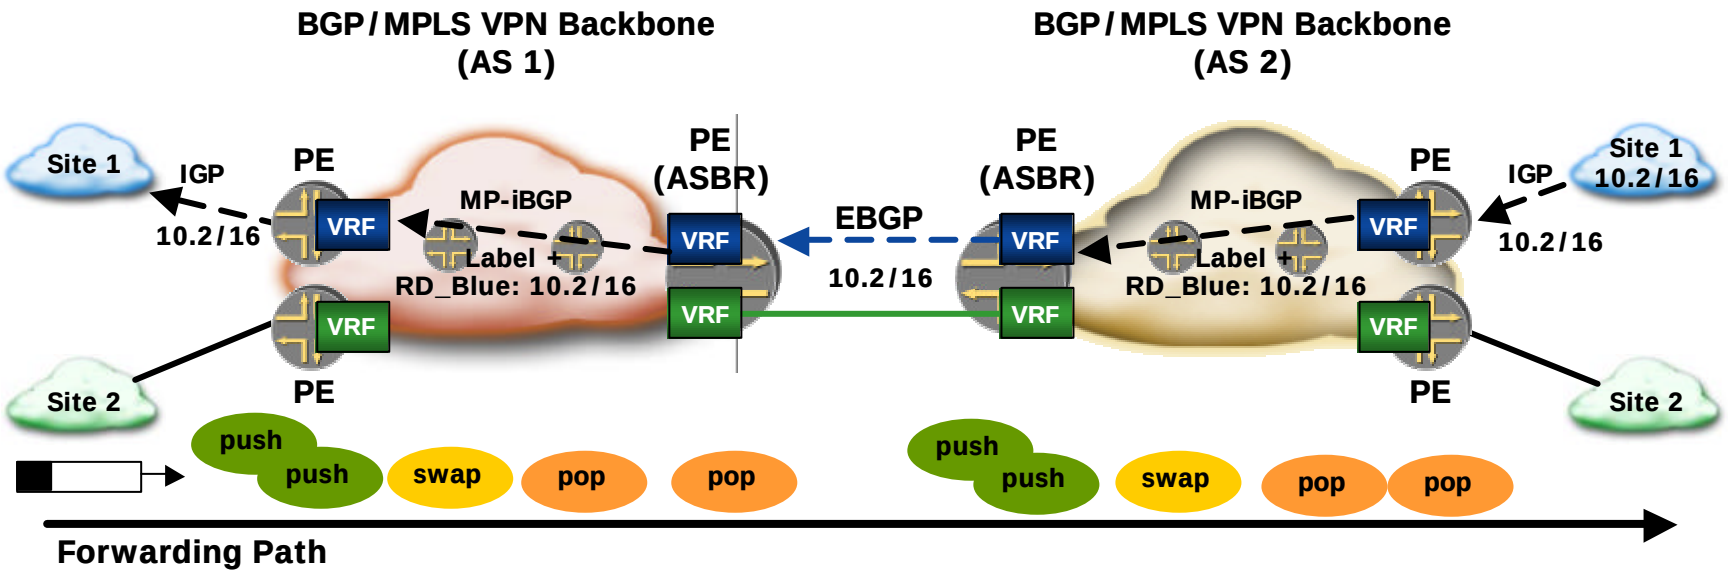

- MPLS not required at the boundary between ASs
- scalability limitations:
  - requires per-VPN configuration on the PE (ASBR) routers
  - requires ASBRs to maintain an extremely large number of VPN-IPv4 routes

MP-eBGP Distribution of Labeled VPN-iPv4 Routes between ASBRs

Inter-Provider Backbones Option B in 2547bis

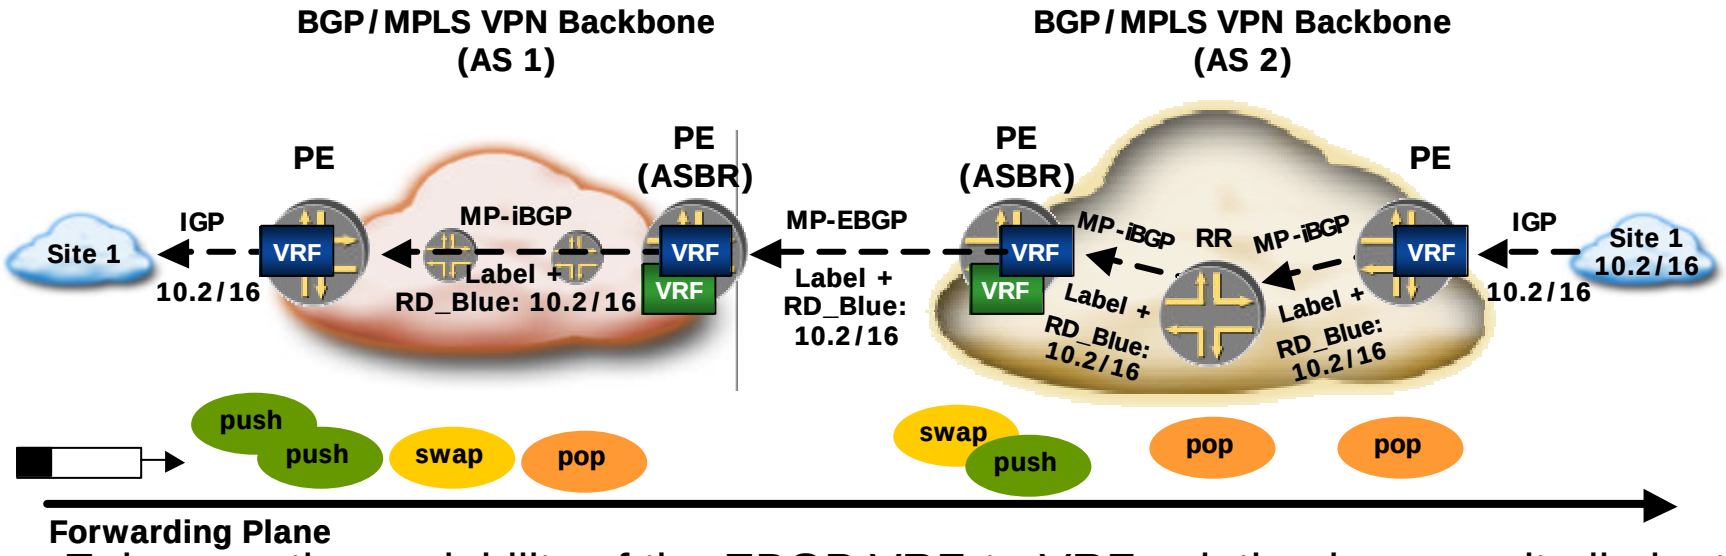

- Enhances the scalability of the EBGP VRF-to-VRF solution because it eliminates the need for per-VPN configuration on the PE (ASBR)s
- Requires an LSP be established from the ingress PE router to the egress PE router
- Requires trust relationships between and among the set of autonomous systems along the path from the ingress PE router to the egress PE router
- Requires understandings between and among the ASs concerning which ASBRs receive routes with specific Route Target attributes.

Multi-hop MP-EBGP Distribution of Labeled VPN-IPv4 Routes Between PE Routers (1)

Inter-Provider Backbones Option C in 2547bis

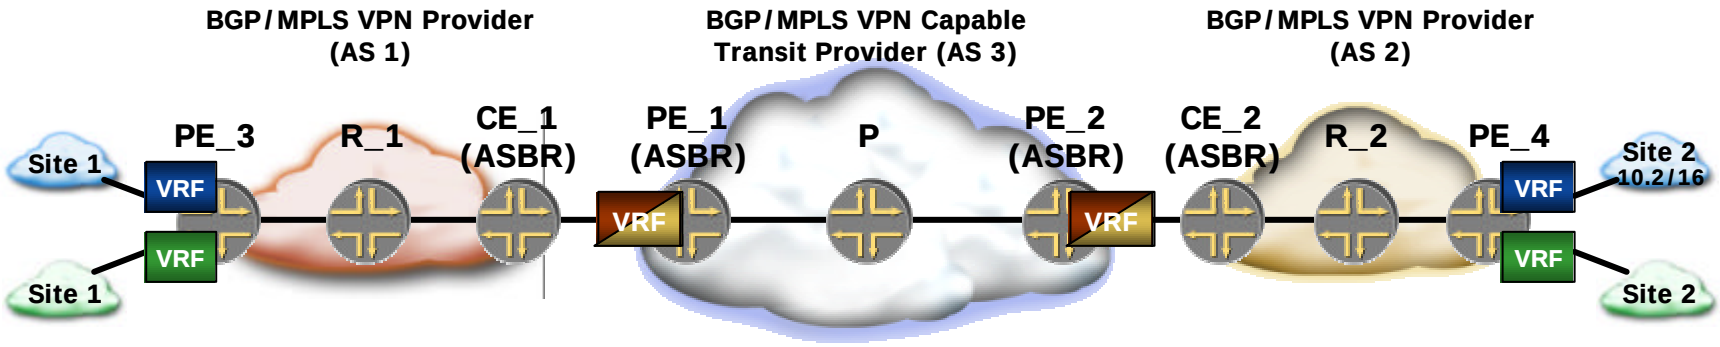

- Advertise labeled Internal Routes (/32) routes into other AS
- Establish LSP between ingress and egress PE
- Use multihop EBGP over established LSP
- If /32 PE addresses not advertised to P router can use 3-level label-stack
- ASBR is not aware of VPN information (scalable !)

Multi-hop MP-EBGP Distribution of Labeled VPN-IPv4 Routes Between PE Routers (2)

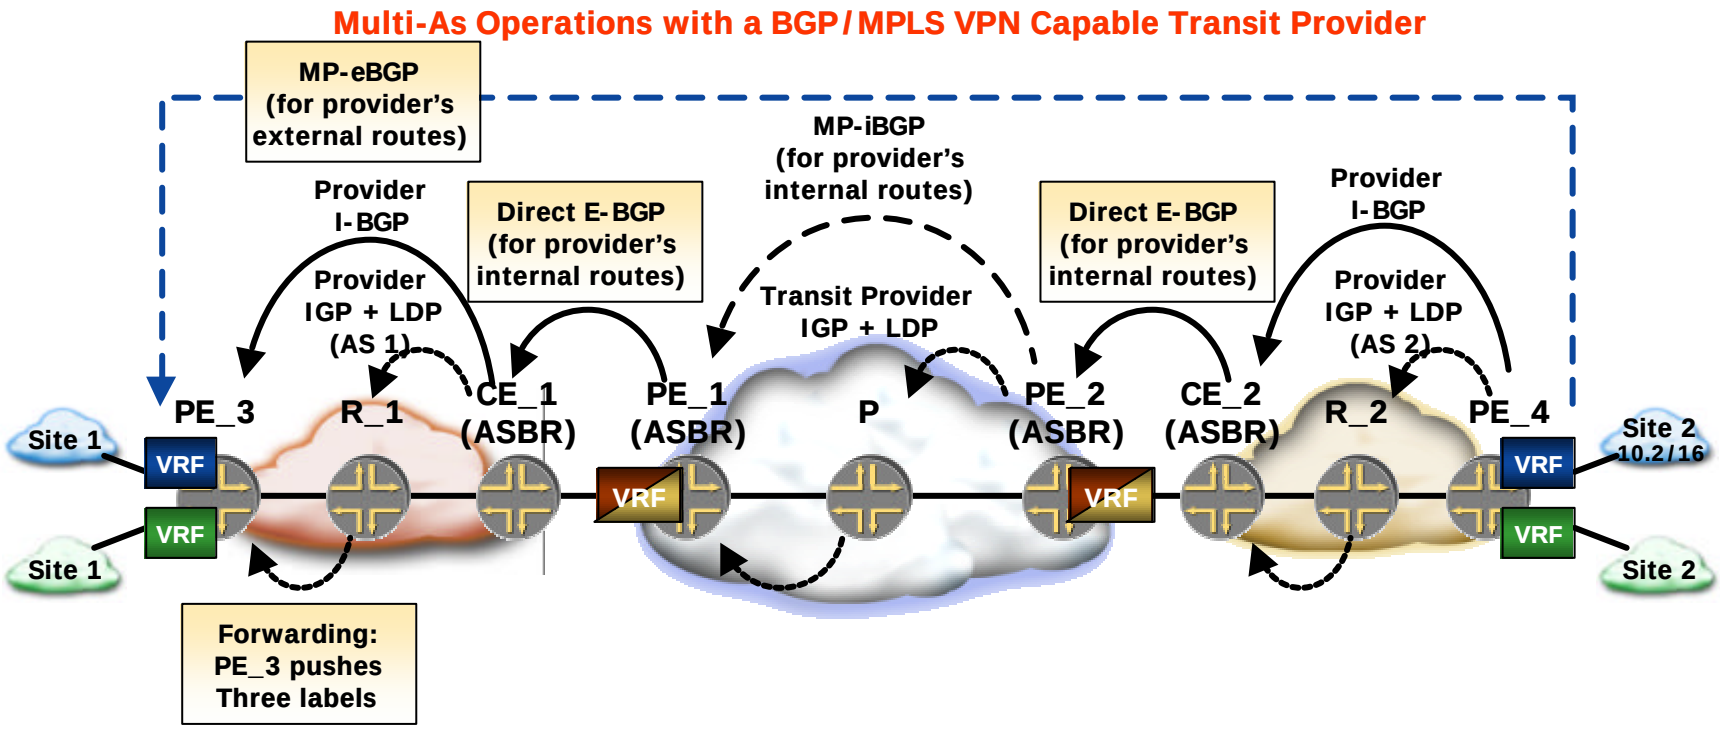

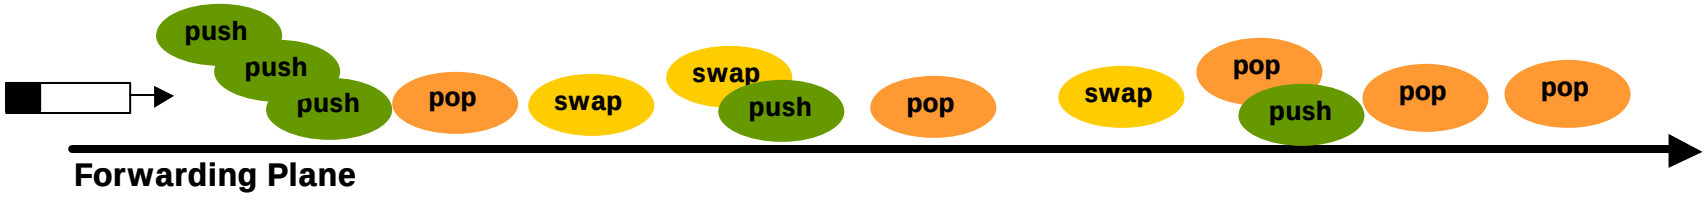

Multi-hop MP-EBGP Distribution of Labeled VPN-IPv4 Routes Between PE Routers (3)

Multi-As Operations with a Direct Connection Between BGP/MPLS VPN Providers

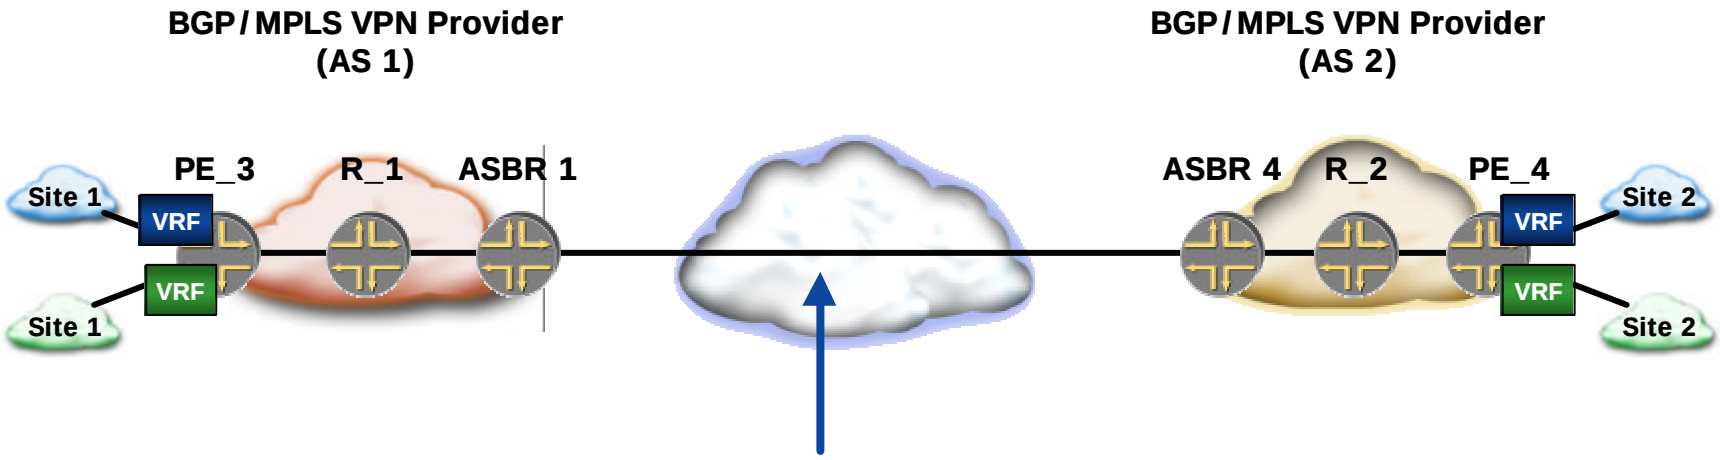

Can be:
- **a direct L2 link**
- **a L2 VPN pt-to-pt connection**
- **a GRE/IPSec tunnel**

Multi-hop MP-EBGP Distribution of Labeled VPN-IPv4 Routes Between PE Routers (4)

Multi-As Operations with a Direct Connection Between BGP/MPLS VPN Providers

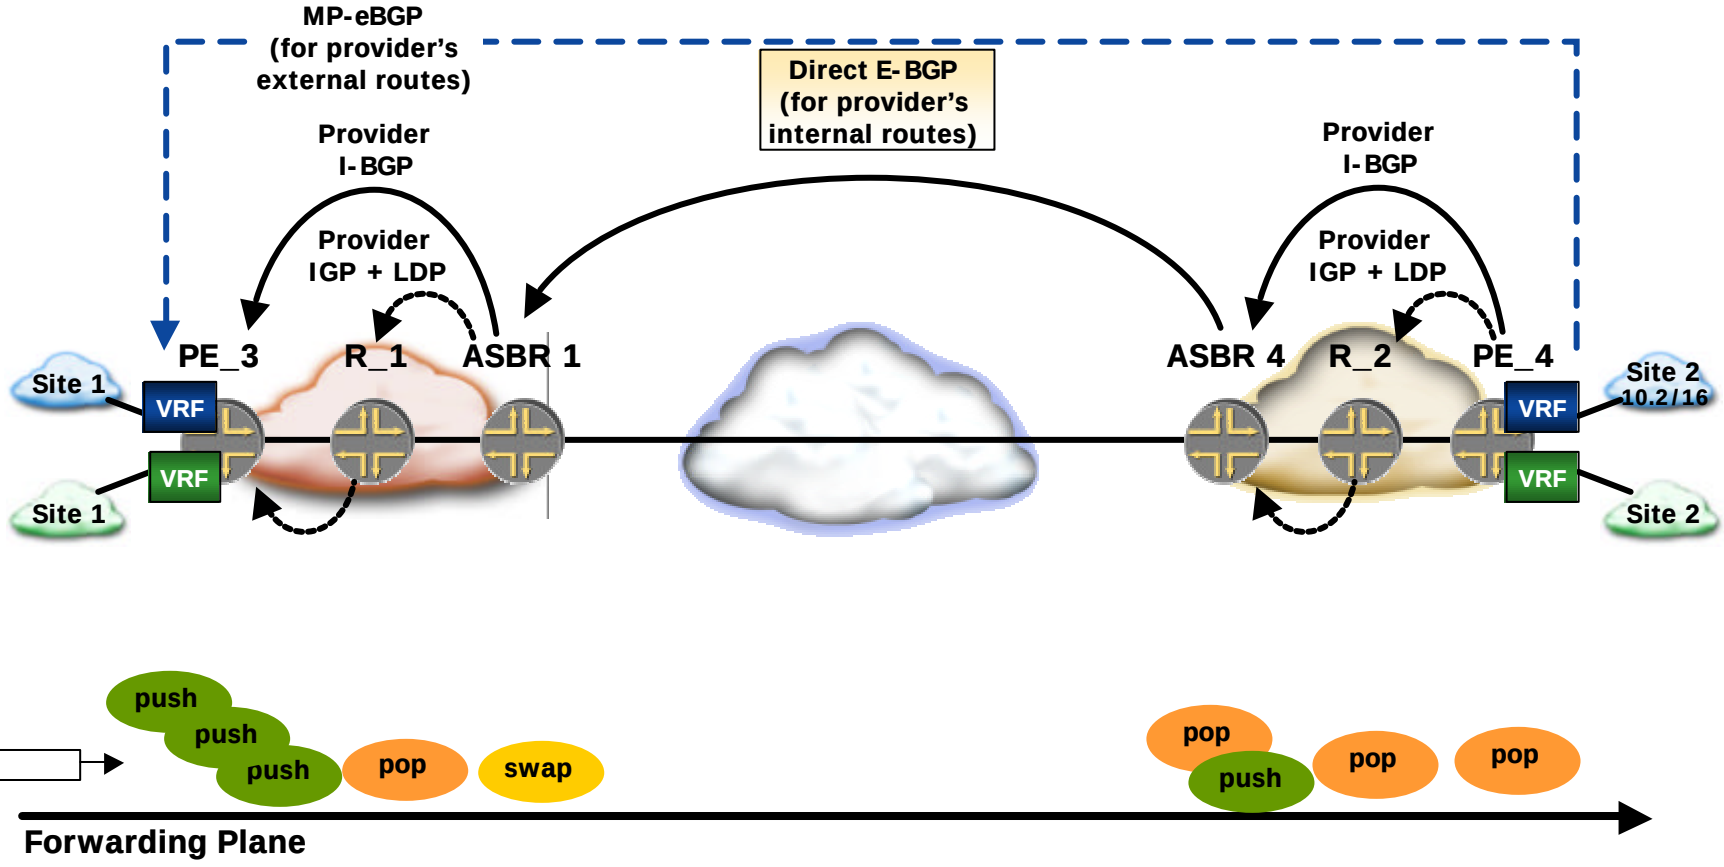

Recursive Multi-AS Operations

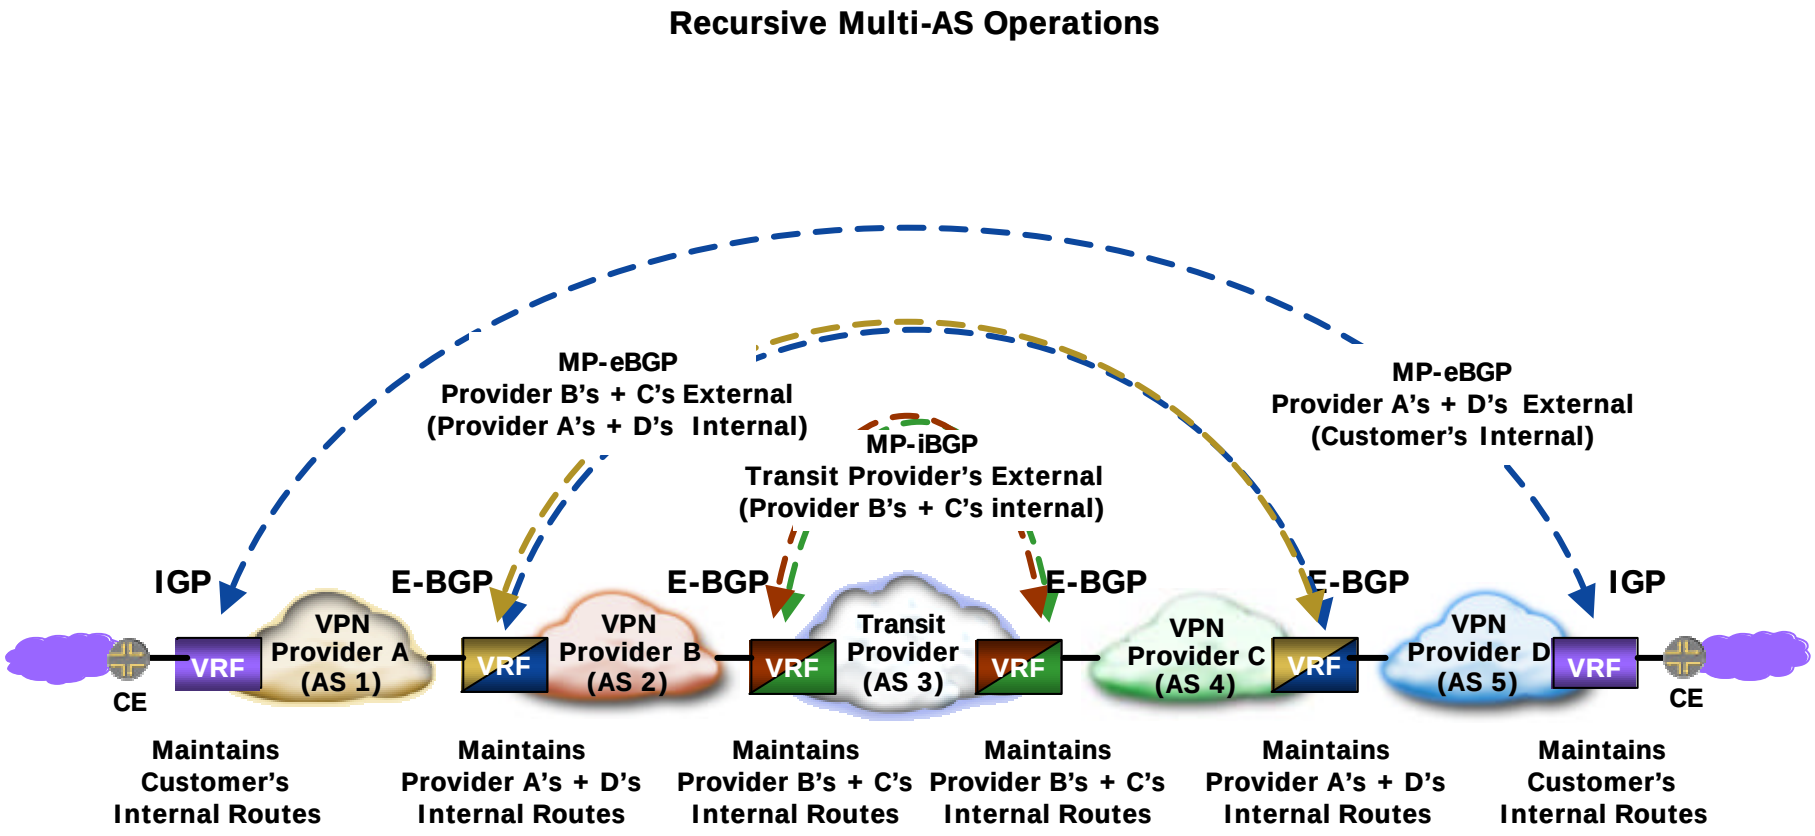

Recursive Multi-AS Operations

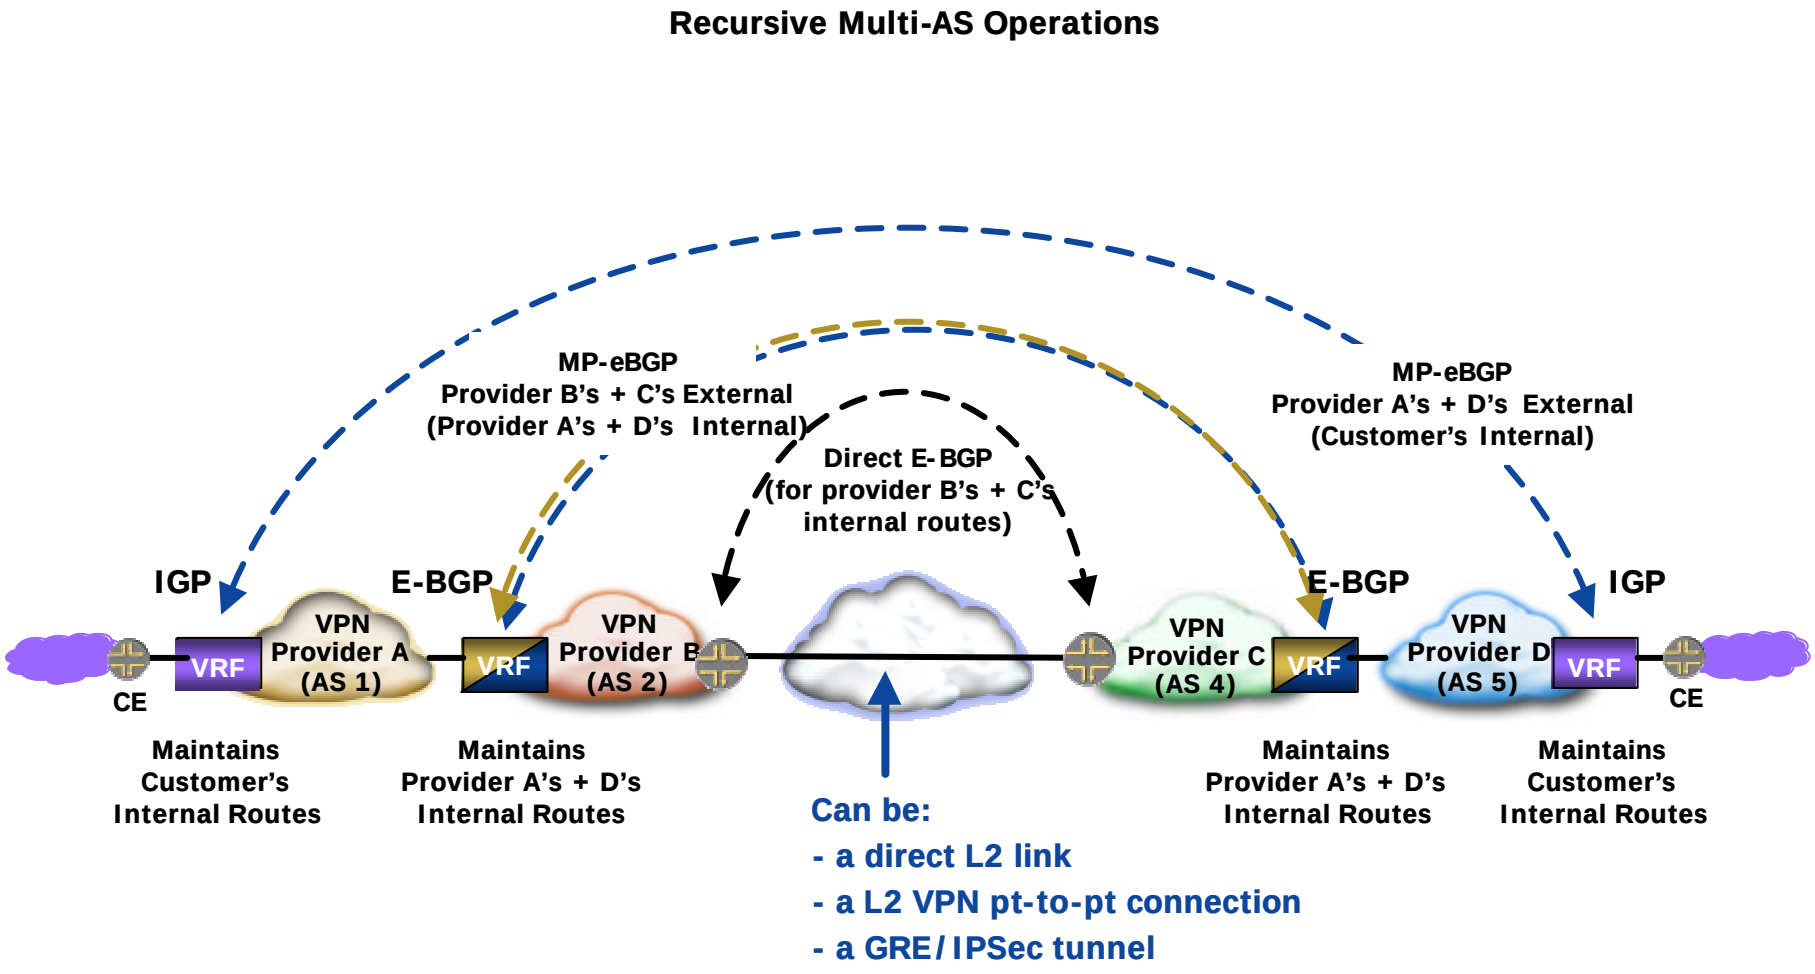

Recursive Multi-AS Operations

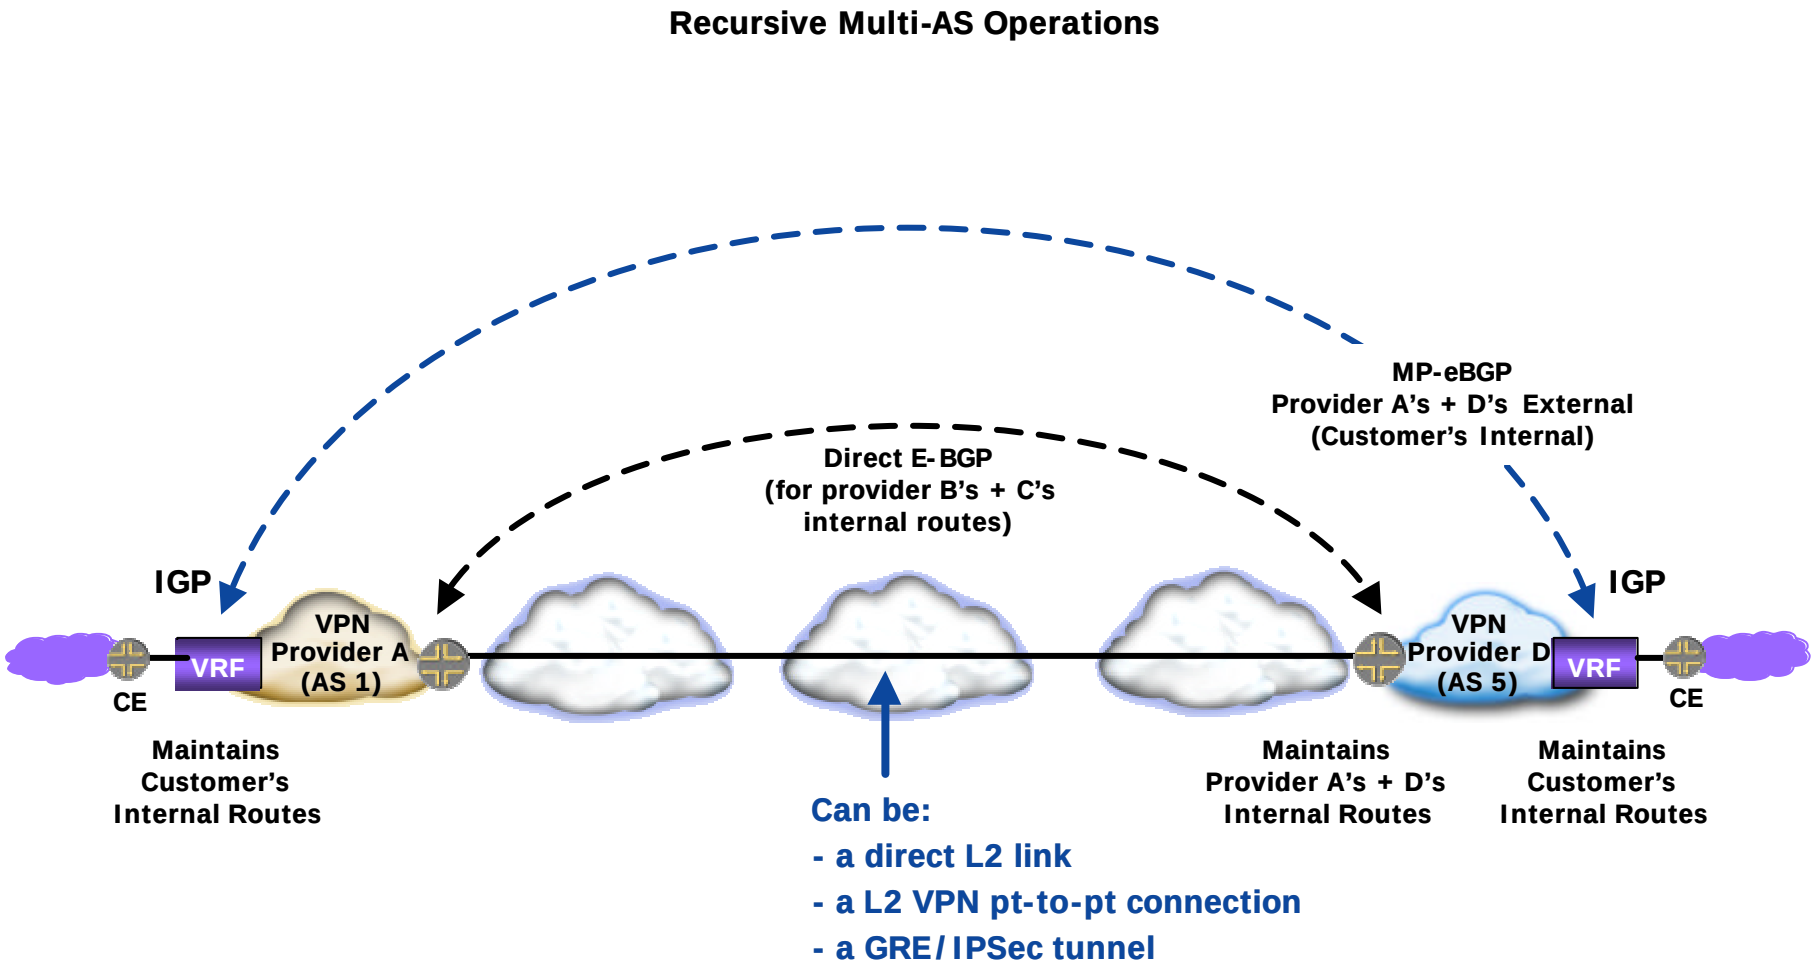

Recursive Multi-AS Operations

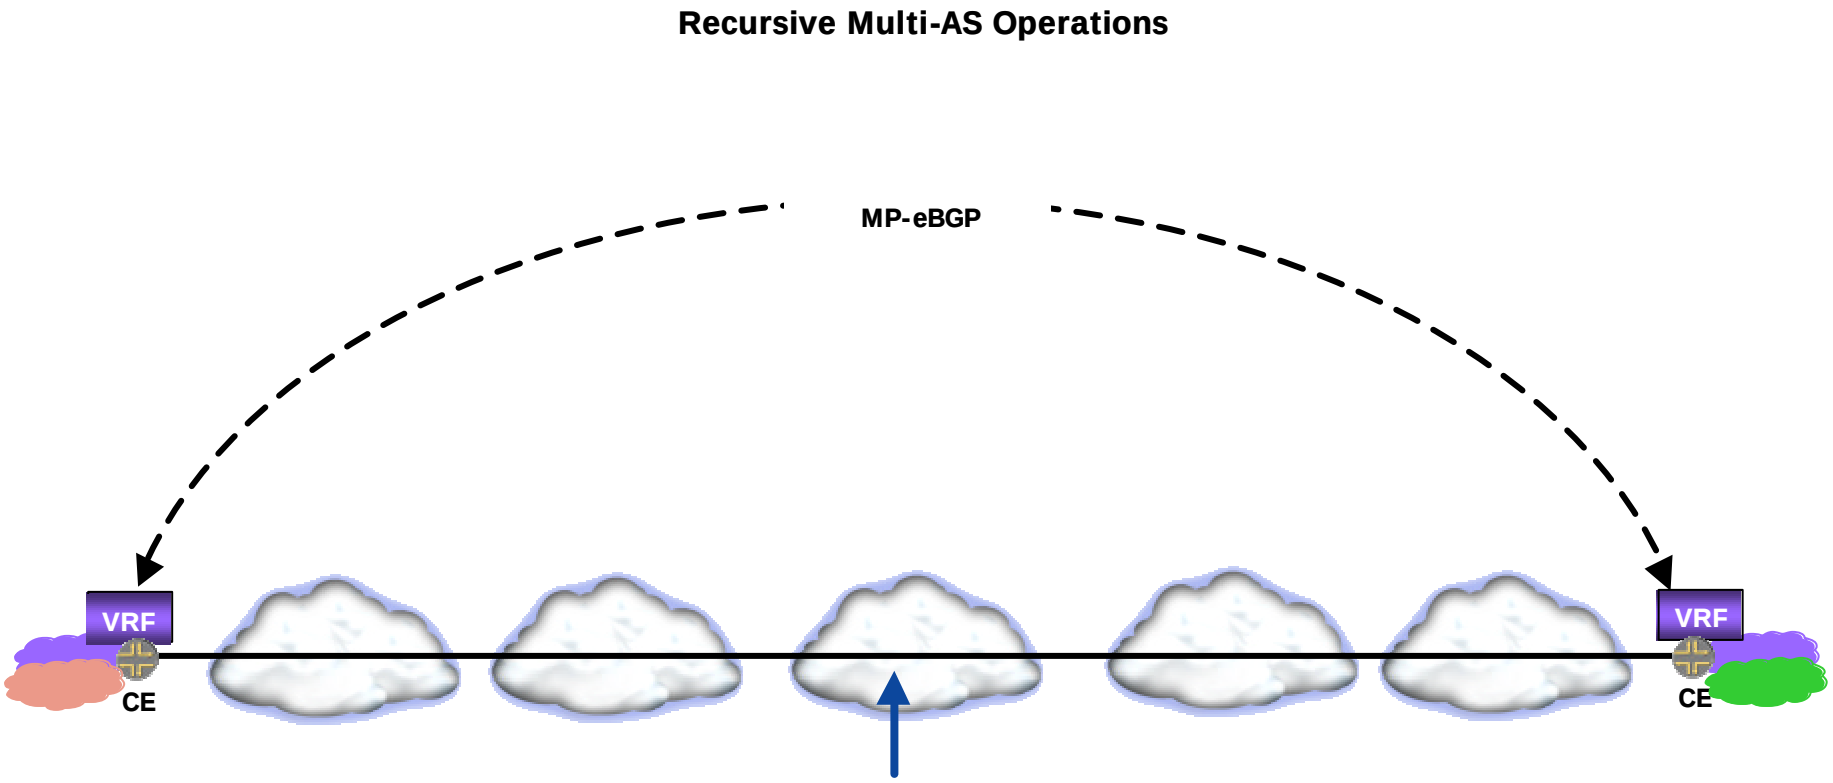

This is actually a CPE based VPN:

- Complexity managed by end-users

- Scalability issue

- Do NOT require any VPN service from transit provider (if GRE Tunnel)

Inter-AS/Inter-provider operations

- Exchange VPN information + VPN labels across AS/provider boundary by using BGP between BGP Route Reflectors in each AS/provider
  - Route Reflectors preserve the next hop information and the VPN label across the AS/provider
- PEs learn routes and label information of the PEs in the neighboring ASes through ASBRs
  - Using labeled IPv4 routes
- No VPN information (e.g., VRF, VFT) on ASBRs

Applies to RFC2547 VPN, L2 VPN, and VPLS !!!

Scalability - “divide and conquer”

(1) Two levels of labels to keep P routers free of all the VPN routing information

(2) PE router has to maintain VPN information only for VPNs whose sites are directly connected to the PE router

(3) Partition BGP Route Reflectors within the VPN Service Provider among VPNs served by the Provider

⇒ No single component within the system is required to maintain information for all the VPNs

⇒ Routing capacity of the system isn’t bounded by the capacity of an individual component

Applies to RFC2547 VPN, L2 VPN, and VPLS !!!

Summary

End Users want:

- Point-to-point Layer 2 VPNs
- Virtual Private LAN Service (VPLS)
- IPv4 and IPv6 VPNs (RFC 2547 VPN)

Research & Education Networks can offer all of the above:

- over a common infrastructure (MPLS)
- with a common framework (Multi-Protocols BGP/MPLS)
  - Taking advantage of BGP scalability and multi-AS/multi-provider support
- with common concepts (Route Distinguisher, Route Target, VRF/VFTs, …)

Can be supported over any forwarding infrastructure (MPLS, IP Tunnels…)

References

- RFC 2547 “BGP/MPLS VPNs”
- draft-ietf-ppvpn-rfc2547bis
- draft-ietf-ppvpn-bgpvpn-auto
- draft-ietf-ppvpn-bgp-ipv6-vpn
- draft-kompella-ppvpn-l2vpn
- draft-kompella-ppvpn-vpls

Thank you!

http://www.juniper.net

juze@juniper.net