INTER-AUTONOMOUS SYSTEM MPLS VPN

February 2004

Agenda

- Inter-Autonomous System (Inter-AS) Multiprotocol Label Switching (MPLS) VPN Overview
- Inter-AS Control and Forwarding Planes
- Inter-AS Connectivity Models
- Inter-AS Summary

MPLS VPNs review in one slide

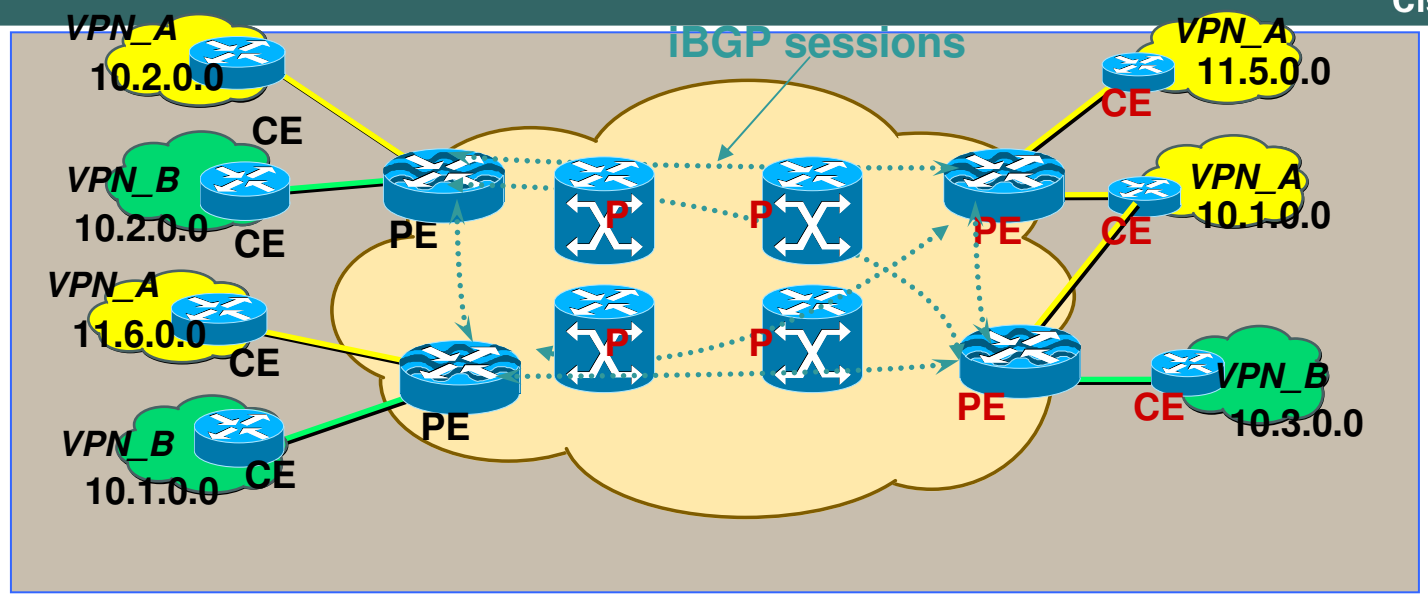

- P routers (LSRs) are in the core of the MPLS cloud
- PE routers use MPLS with the core and plain IP with CE routers
- P and PE routers share a common IGP (LS)
- PE router are MP-iBGP fully meshed or use Route-Reflectors
- Route Distinguishers (RD) are used to achieve uniqueness of IP address (RD+IPv4=VPNv4)
- Route Targets are used to export and import routes in a VPN
- Multiple routing tables (VRFs) are used on PEs
  - Each VRF contain customer routes
  - Customer addresses can overlap
  - VPNs are isolated
- MP-BGP is used to propagate these addresses between PE routers

Why Inter-AS?

- Extends MPLS VPN services across geographical boundaries, so Service Providers can support their customer base in geographical locations that do not have POPs
- Enables communication between networks under separate autonomous systems (Service Providers)
- Provides traffic separation and maintain end-to-end privacy while traffic traverses multiple MPLS VPN backbones

Inter-AS Topology Overview

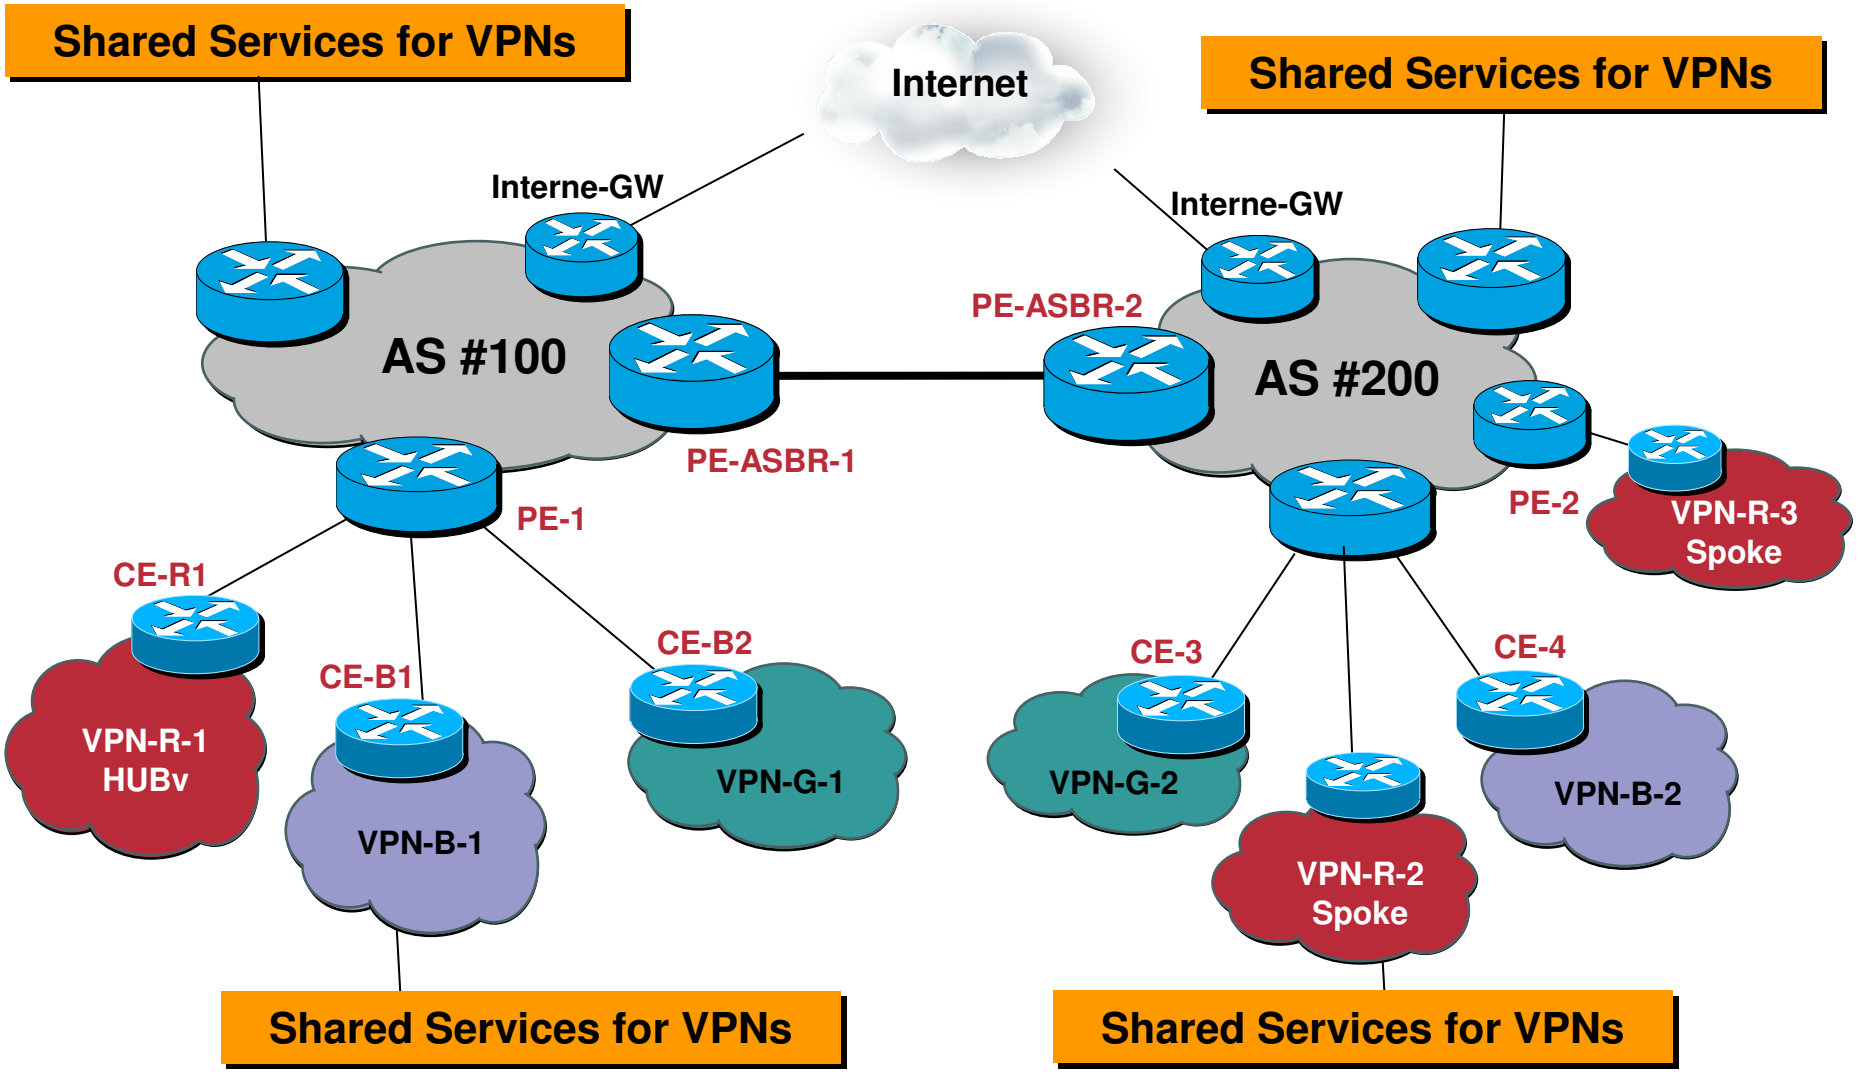

Inter-AS Functionality

- MPLS VPN providers exchange routes across VRF interfaces
- Each PE-ASBR router treats the other as a CE
- Provider edge routers are gateways used for VPNv4 route exchange
- PE-ASBR to PE-ASBR link may use any supported PE-CE routing protocol

Routing For Each SP AS

- Each AS runs a different IGP and no IGP routing information exchange between the domains
- All routing information exchange between the domains is via Exterior Routing Protocol
- Routing policies may differ between the exchange points
- Customer VPN routes are distributed into VRFs at the ingress PE of the ISP
- Each PE assigns labels for the routes to establish connections

INTER-AS Control and Data Planes

Inter-AS Control Plane

- Establishes EBGP session between the PE-ASBRs
- Distribute IPv4 routes for the VPNs in the form of VPNv4 addresses
- PE-ASBRs re-write Next-hop and labels when a route is distributed to a neighbor
- PE-ASBRs store ALL VPN routes that need to be exchanged
- Routes are in the MP-BGP table but not in any other routing tables
  - PE-ASBRs do not have any VRF
  - MP-eBGP labels are used in LFIB

Inter-AS Route Exchange

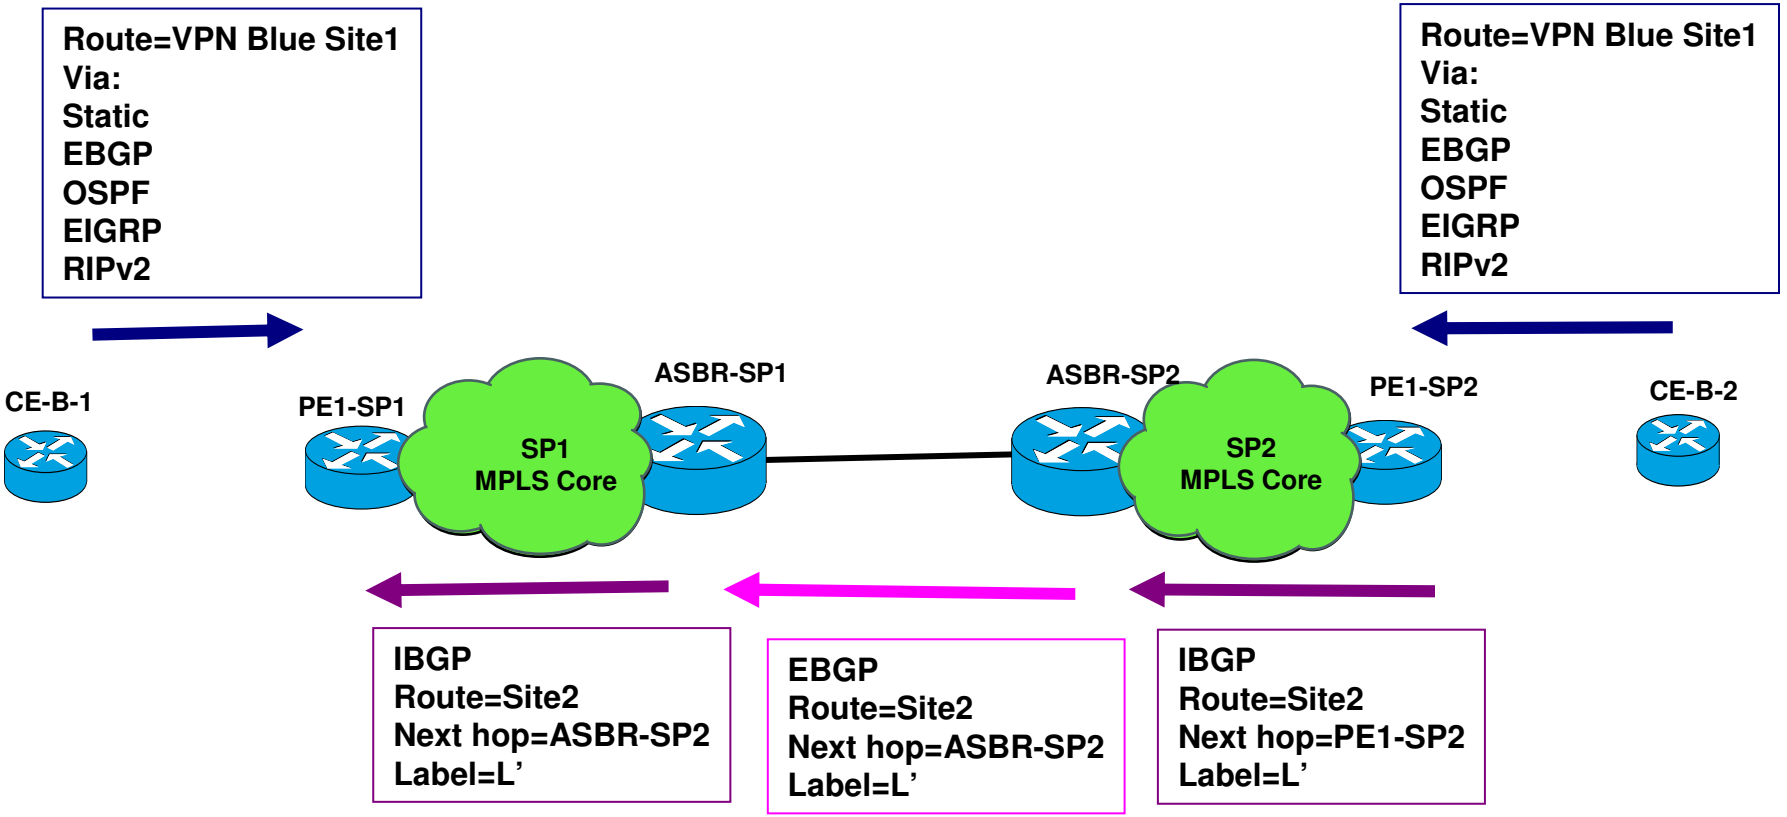

Inter-AS Control Plane

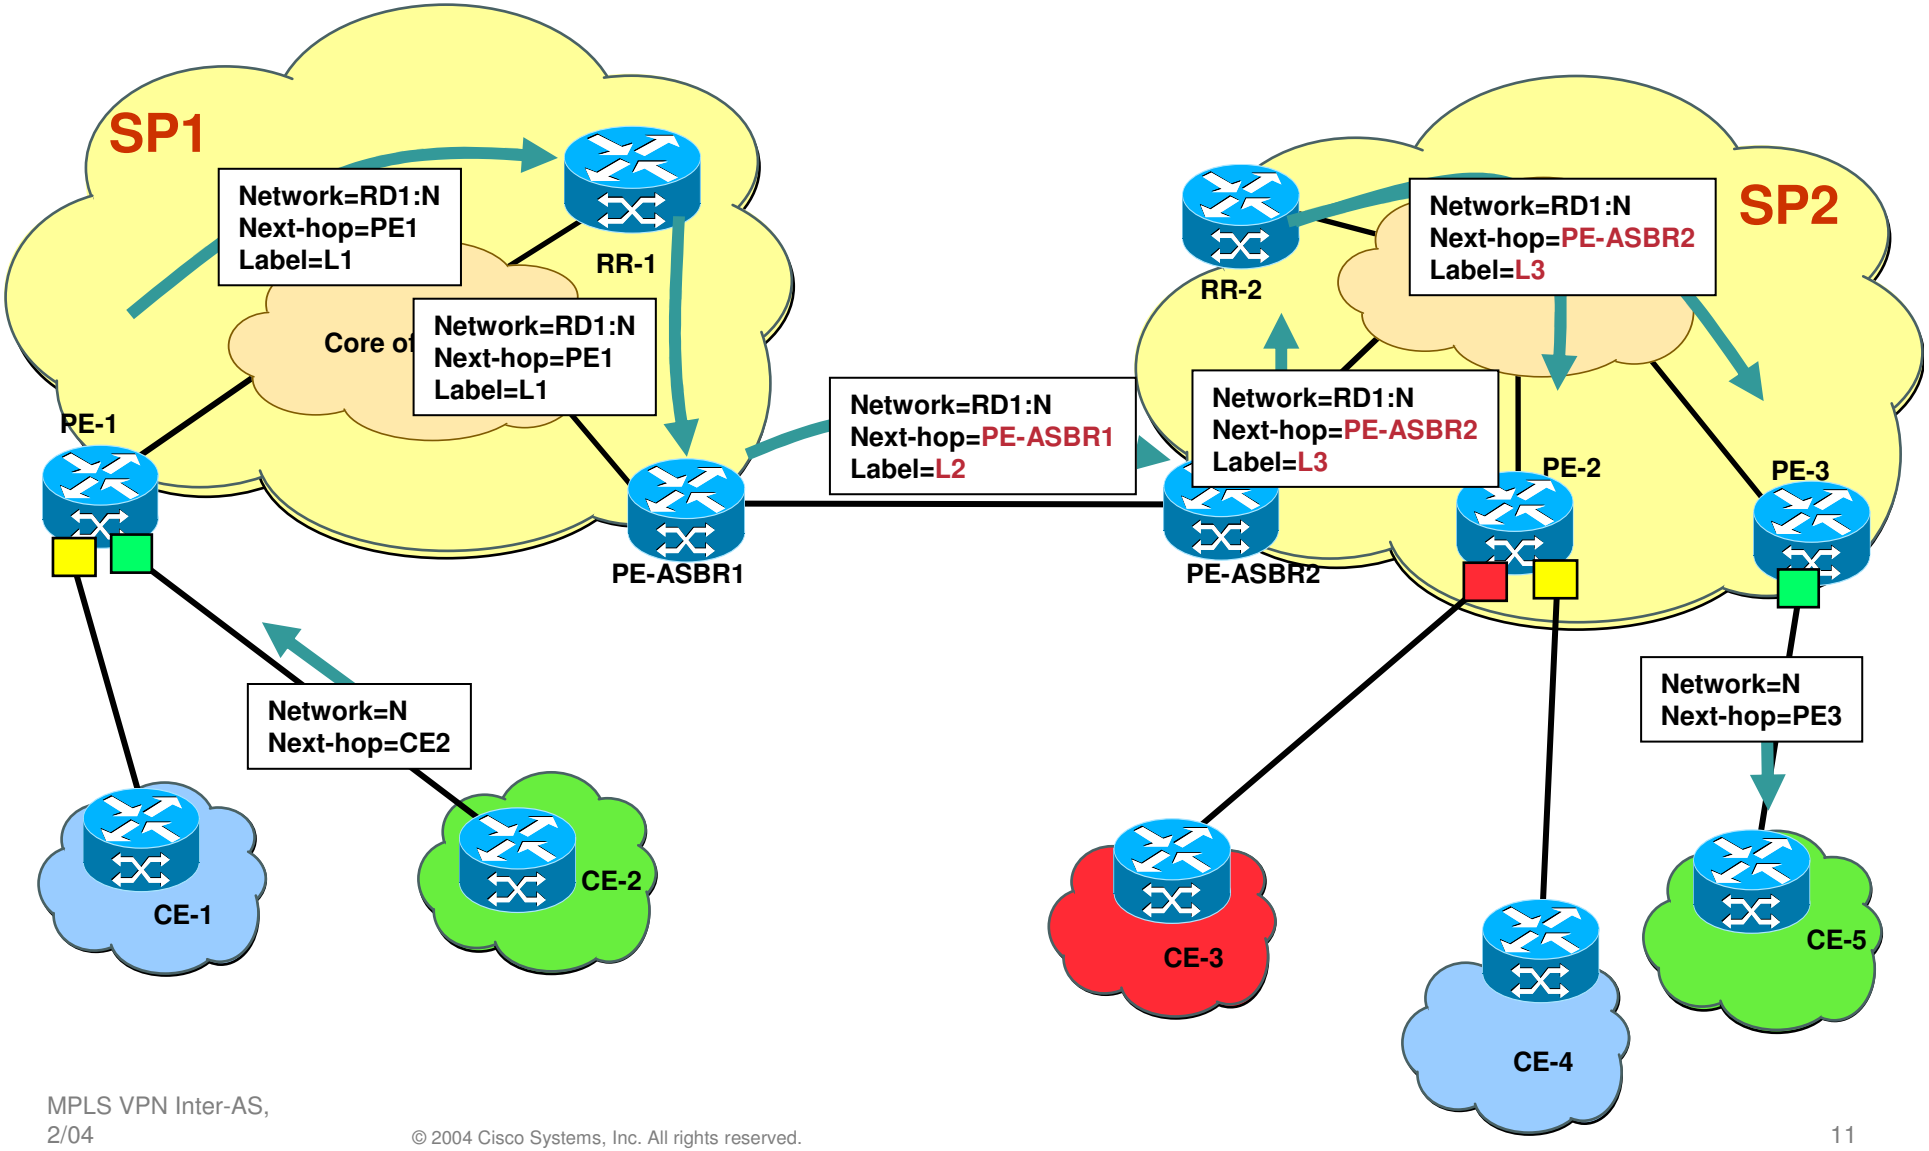

External MP-BGP for VPNv4 Forwarding Plane

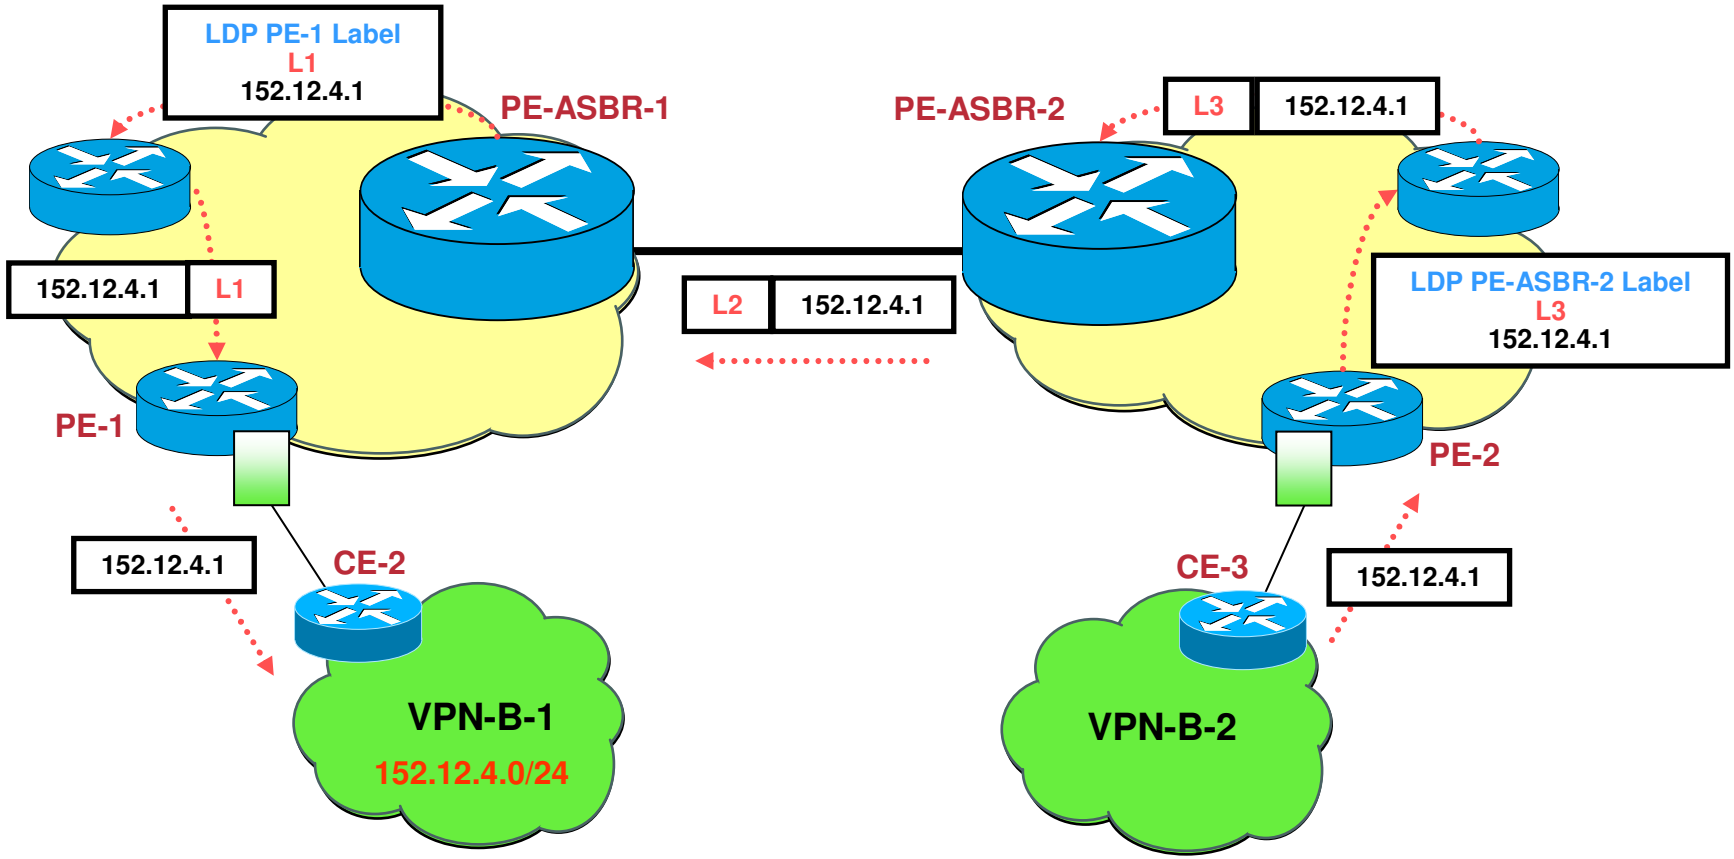

Inter-As Forwarding Plane

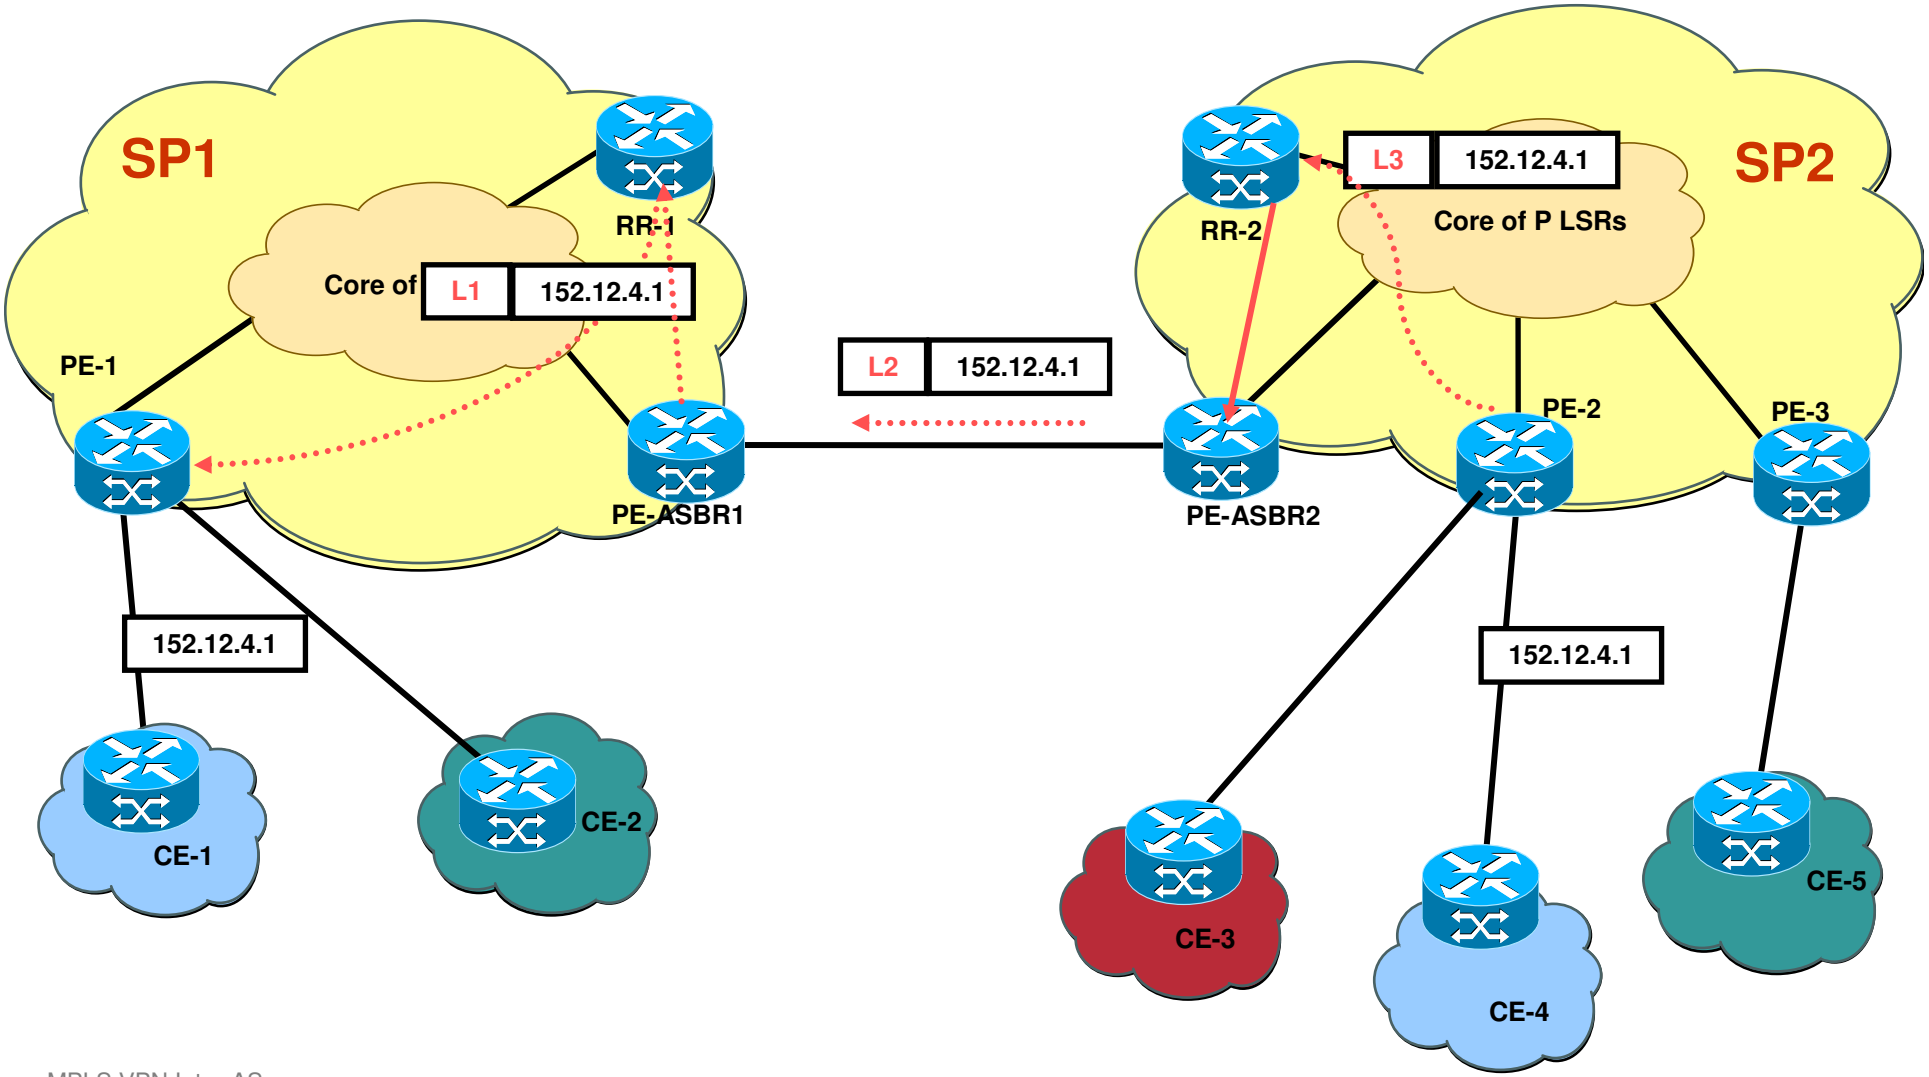

Inter-AS Basic Configuration

- Create a loopback address on participating ASBRs
- Setup ASBRs for VPNv4 route distribution
- Setup ASBRs for IPV4 route distribution
- Disable automatic route filtering feature
- Set ASBR as Next-Hop-Self

The Problem : VPN Client Connectivity

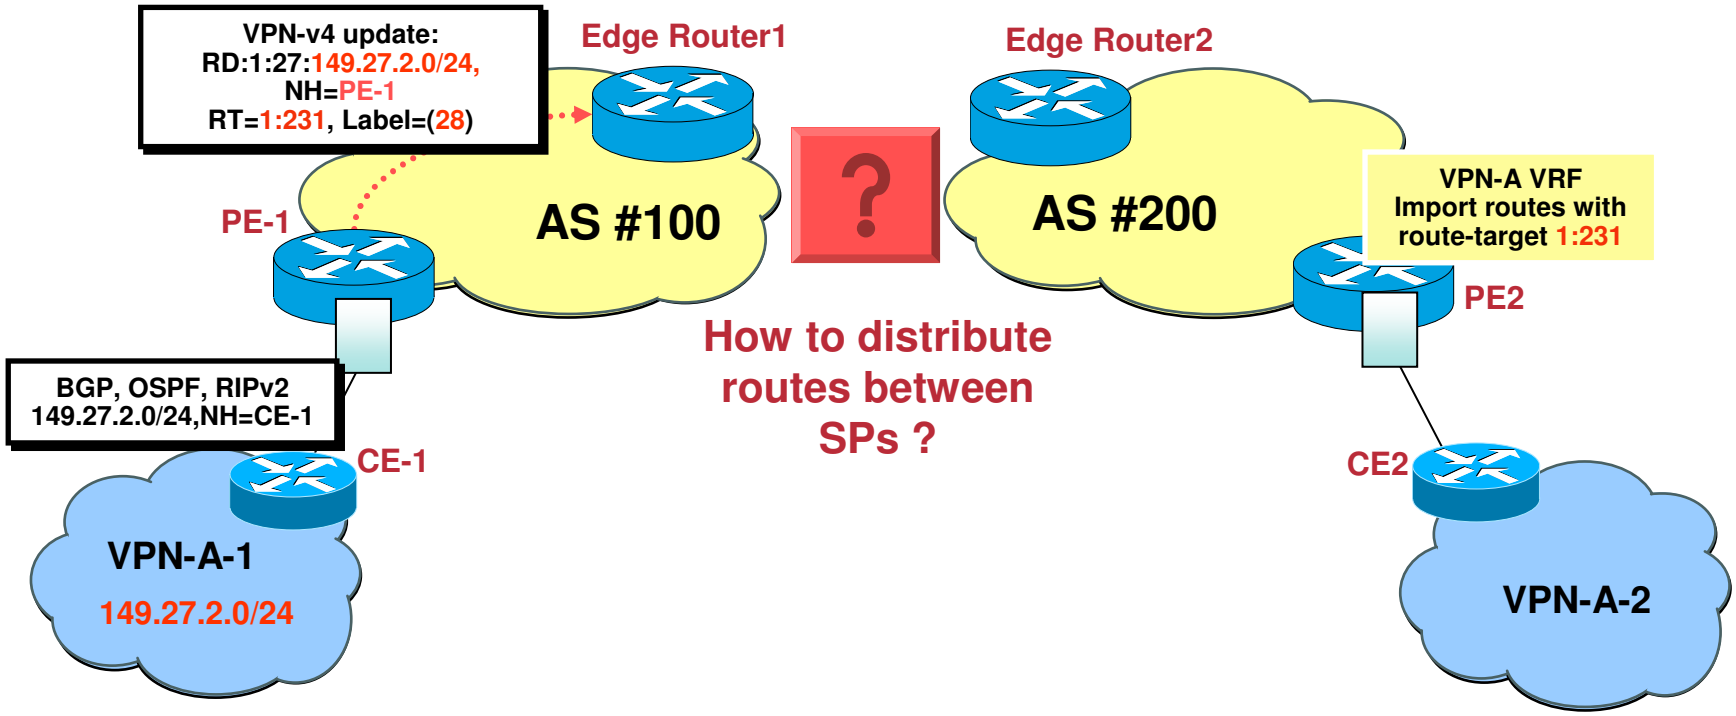

VPN Sites attached to different MPLS VPN Service Providers

The Solution : VPNv4 Distribution in Inter AS

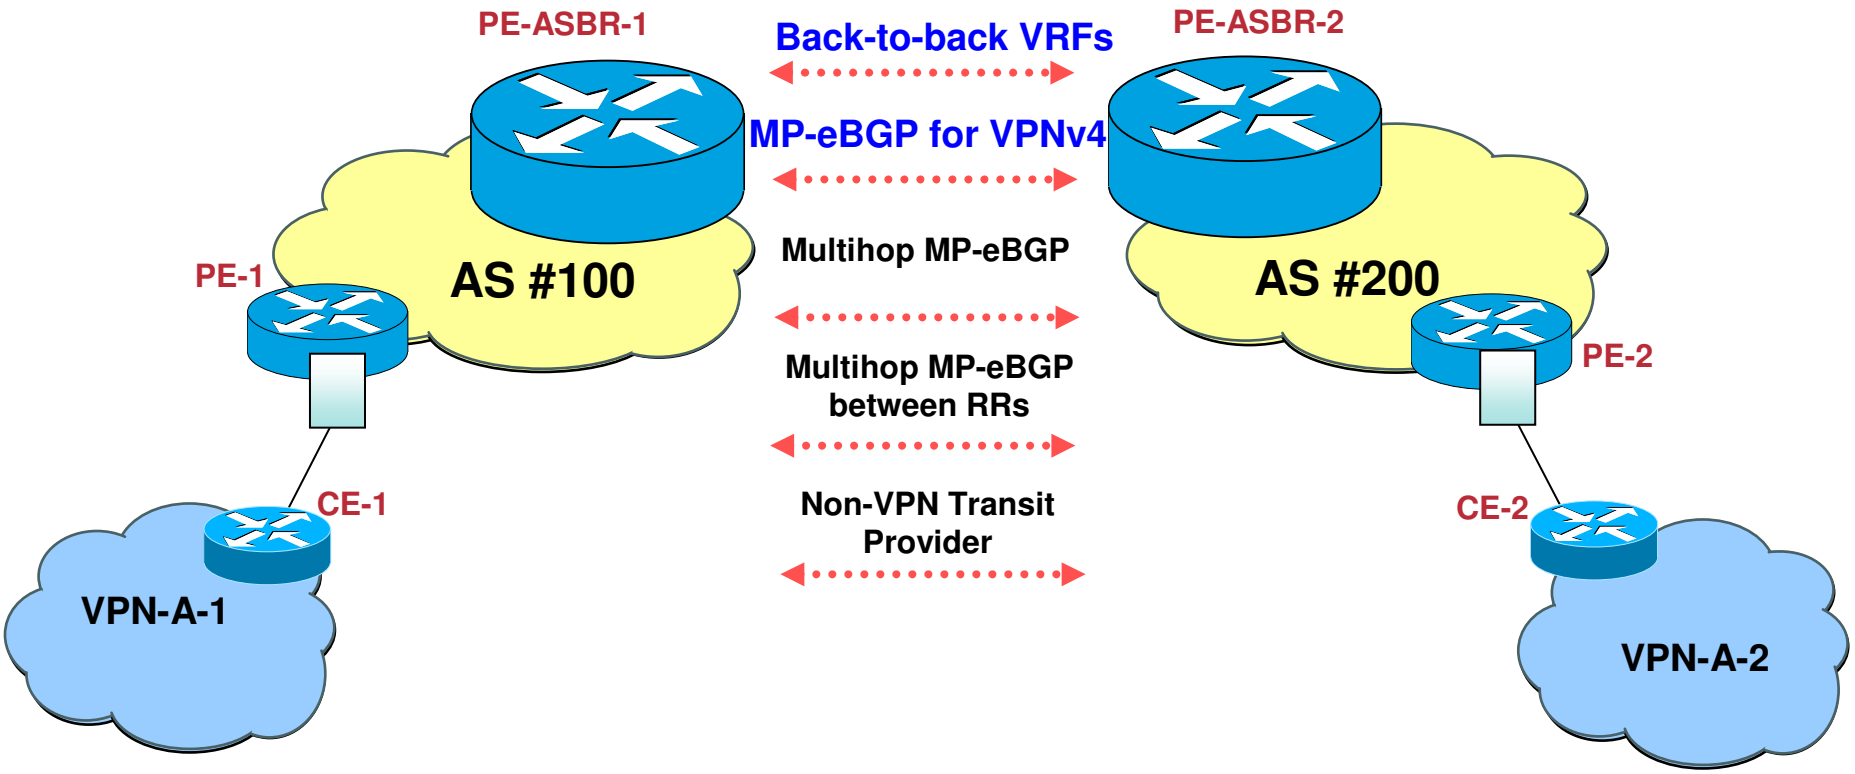

Several options available for distribution of VPNv4 prefix information (we will discuss only 2)

INTER-AS CONNECTIVITY MODELS

Option 1: Back-to-Back VRF Connectivity

Cisco.com

- Recommended for fewer VRFs requiring simpler connectivity when ASBRs are directly connected over a physical interface
- ASBRs are directly connected over a physical interface
- Sub-interface per VRF is created and mapped
- Packet is forwarded as an IP packet between the ASBRs
- Each PE-ASBR router treats the other as a CE
- PE-ASBR to PE-ASBR link may use any supported PE-CE routing protocol
- Scalability issues if need to support large numbers of VRFs

Back-to-Back VRF Connectivity

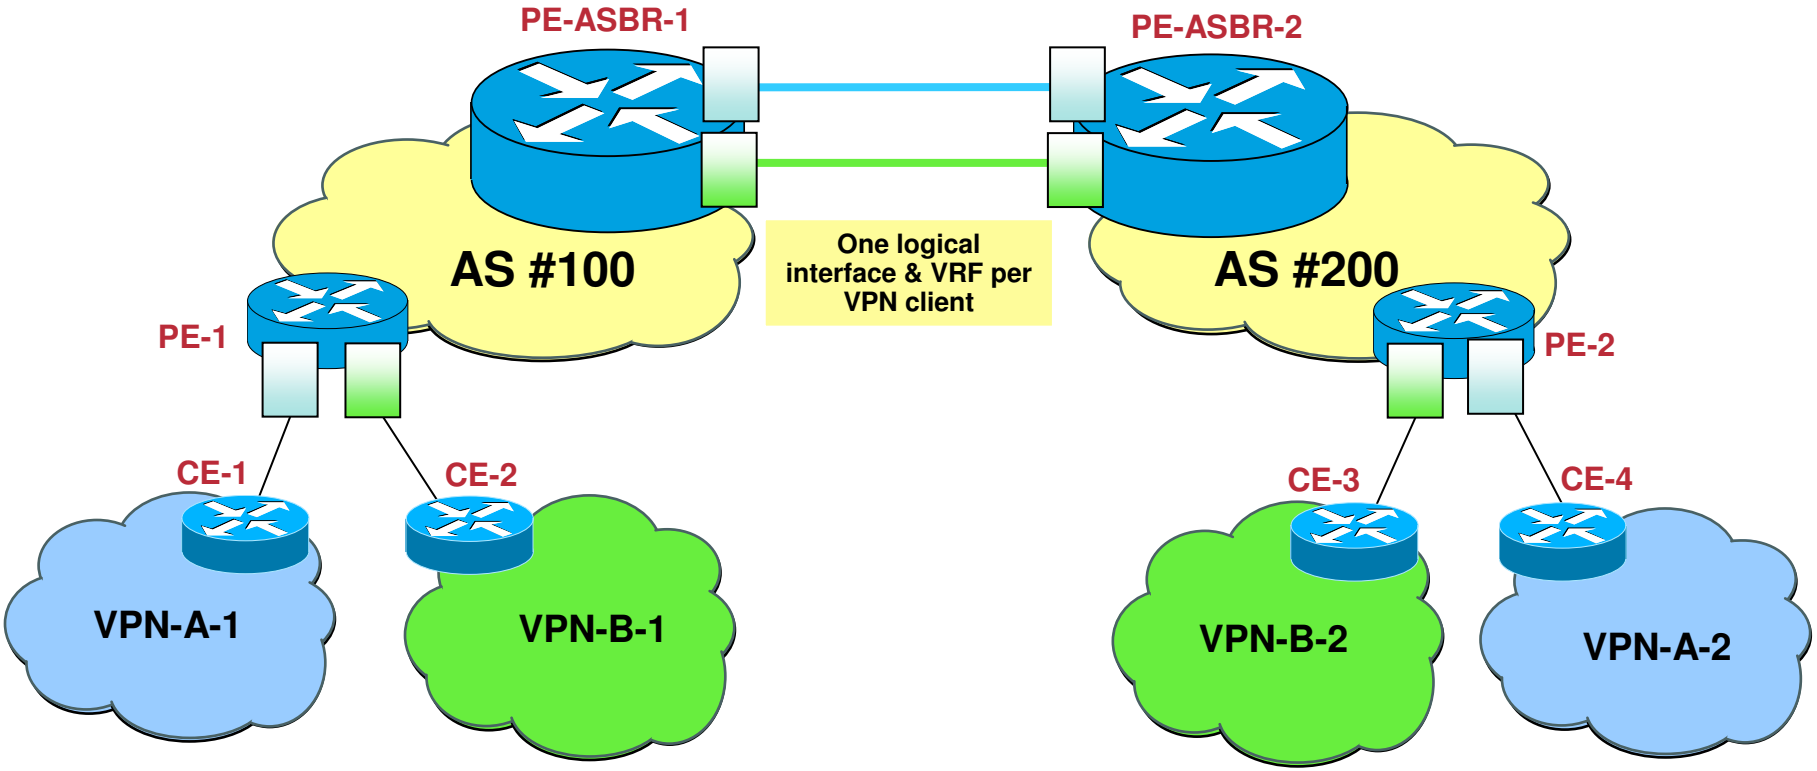

VRF to VRF Connectivity between PE-ASBRs

Back-to-Back VRF Connectivity Control Plane

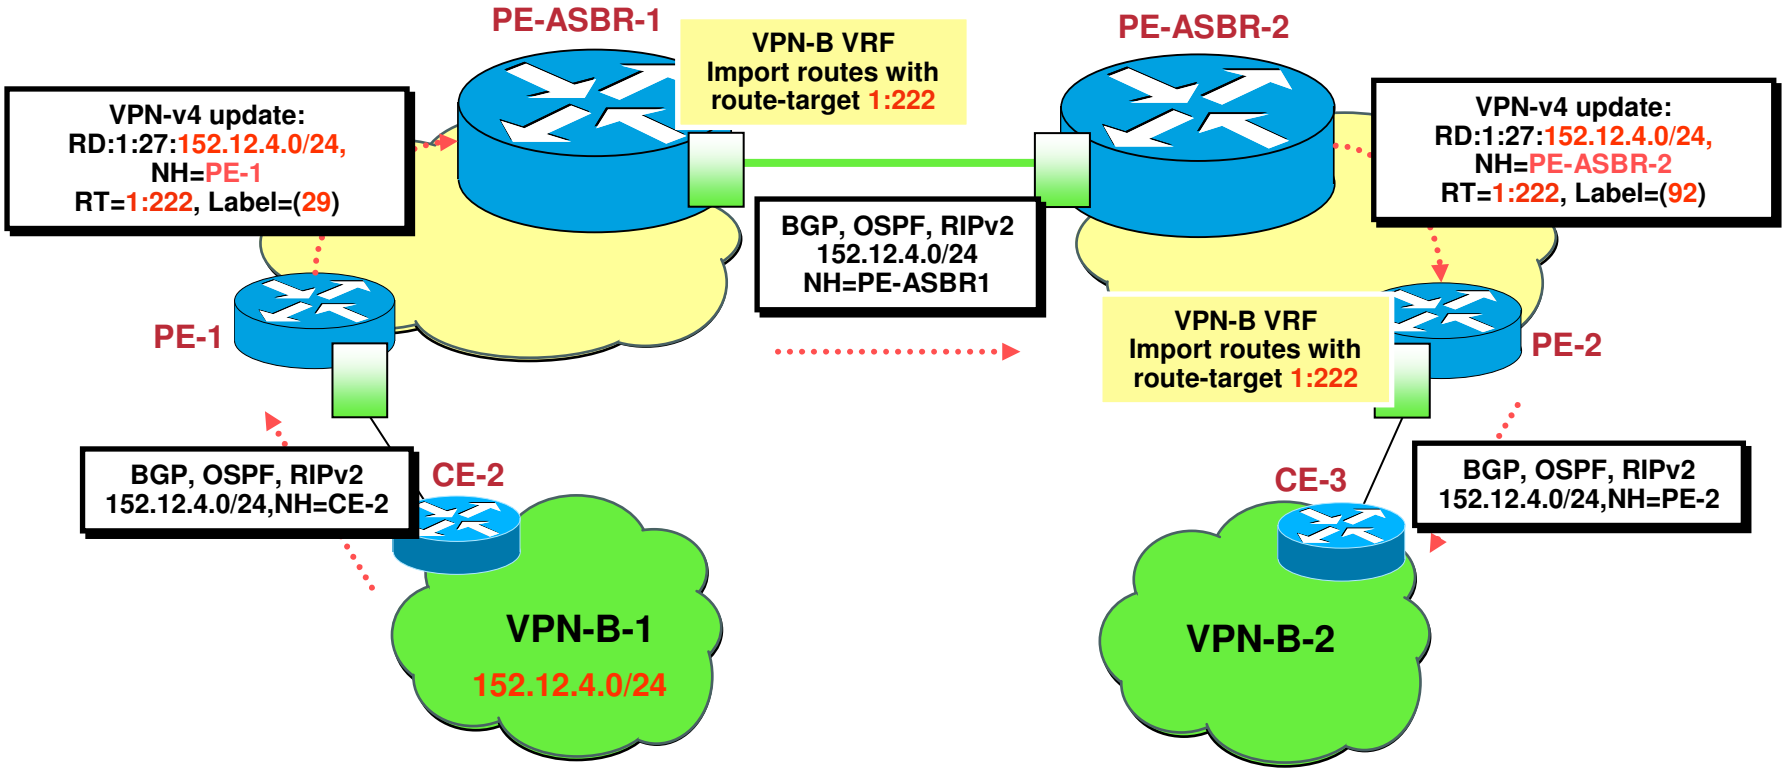

VRF to VRF Connectivity between PE-ASBRs

Back-to-Back VRF Connectivity Forwarding Plane

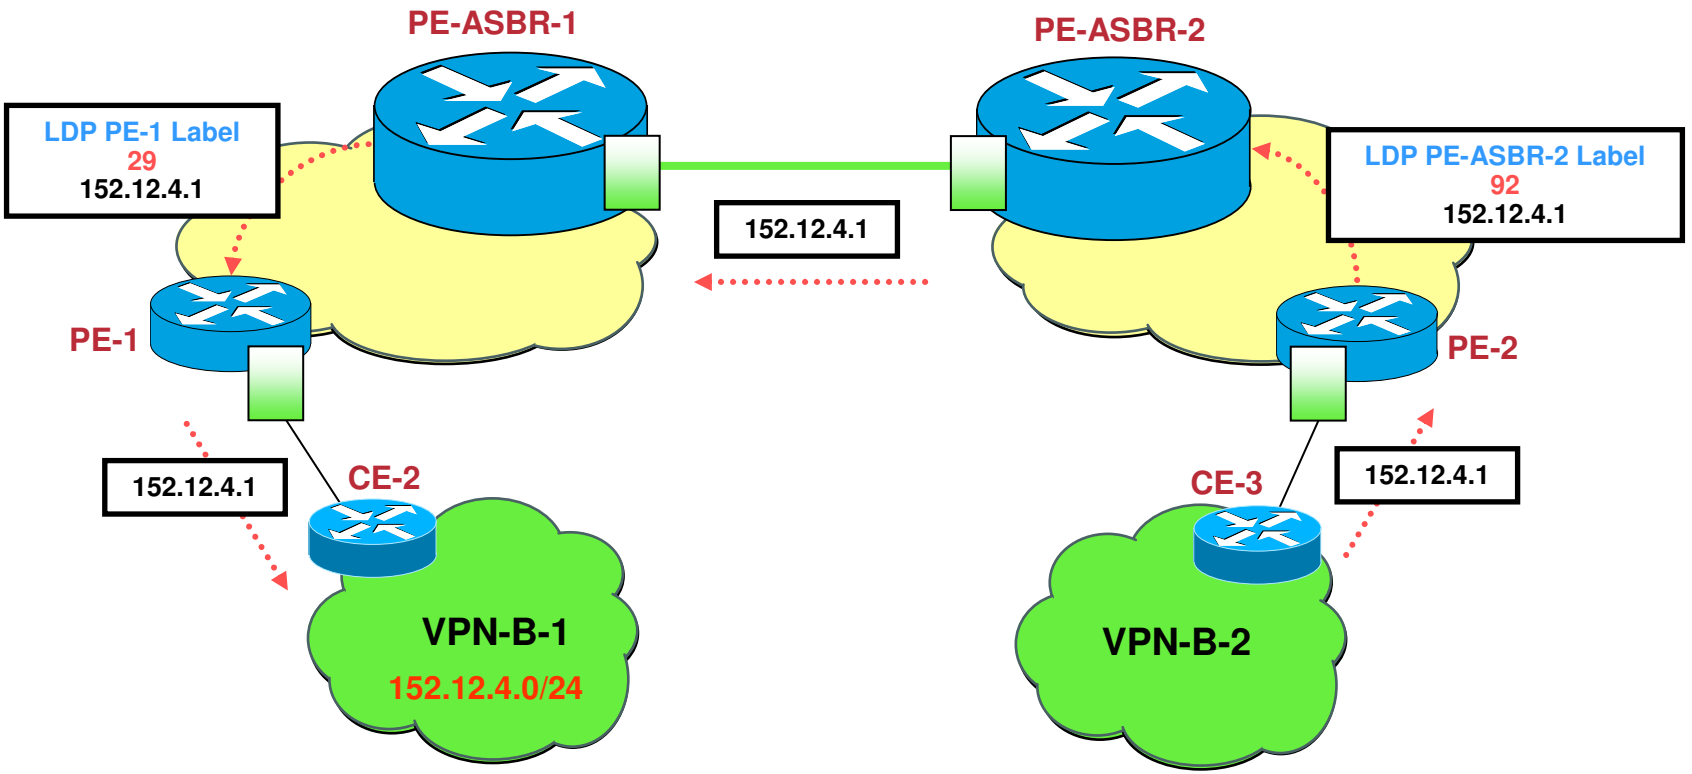

VRF to VRF Connectivity between PE-ASBRs

Option 2: External MP-BGP for VPNv4 Prefix Exchange

- Recommended when a larger number of VRFs need to be supported
- ASBRs are directly connected and belong to only couple service providers
- Traffic will be crossing only single hop network

External MP-BGP for VPNv4 Prefix Exchange (Cont.)

- **Gateway PE-ASBRs exchange routes directly using BGP**
  - **External MP-BGP for VPNv4 prefix exchange**
  - **No LDP or IGP**
- **MP-BGP session with next-hop set to advertising PE-ASBR**
  - **Next-hop and labels are rewritten when advertised across the Inter-Provider MP-BGP session**
- **PE-ASBR stores all VPN routes that need to be exchanged**
  - **Only within the BGP table (no VRFs)**
  - **Labels are populated into the LFIB of the PE-ASBR**

External MP-BGP for VPNv4

- **Receiving Gateway PE-ASBRs may allocate new label if desired**
  - **Controlled by configuration of next-hop-self (default is off)**
- **Receiving PE-ASBR will automatically create a /32 host route for its PE-ASBR neighbor**
  - **Which must be advertised into receiving IGP if next-hop-self is NOT in operation to maintain the LSP;**
- **PE-ASBRs need to hold all Inter-AS VPN routes**

External MP-BGP for VPNv4

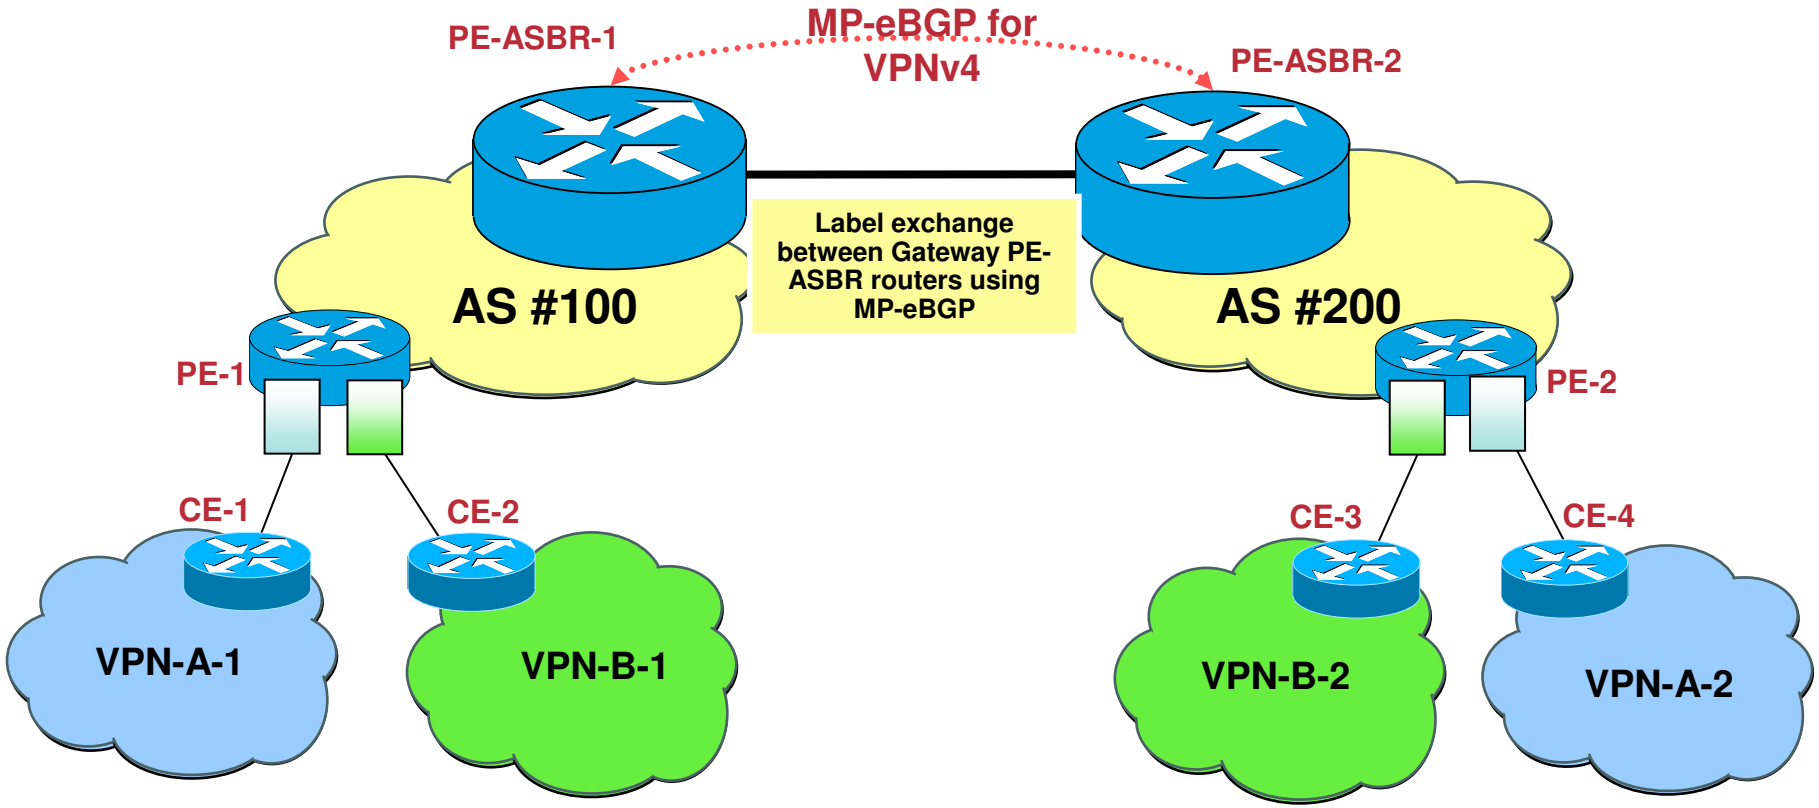

MP-BGP VPNv4 prefix exchange between Gateway PE-ASBRs

External MP-BGP for VPNv4 Control Plane

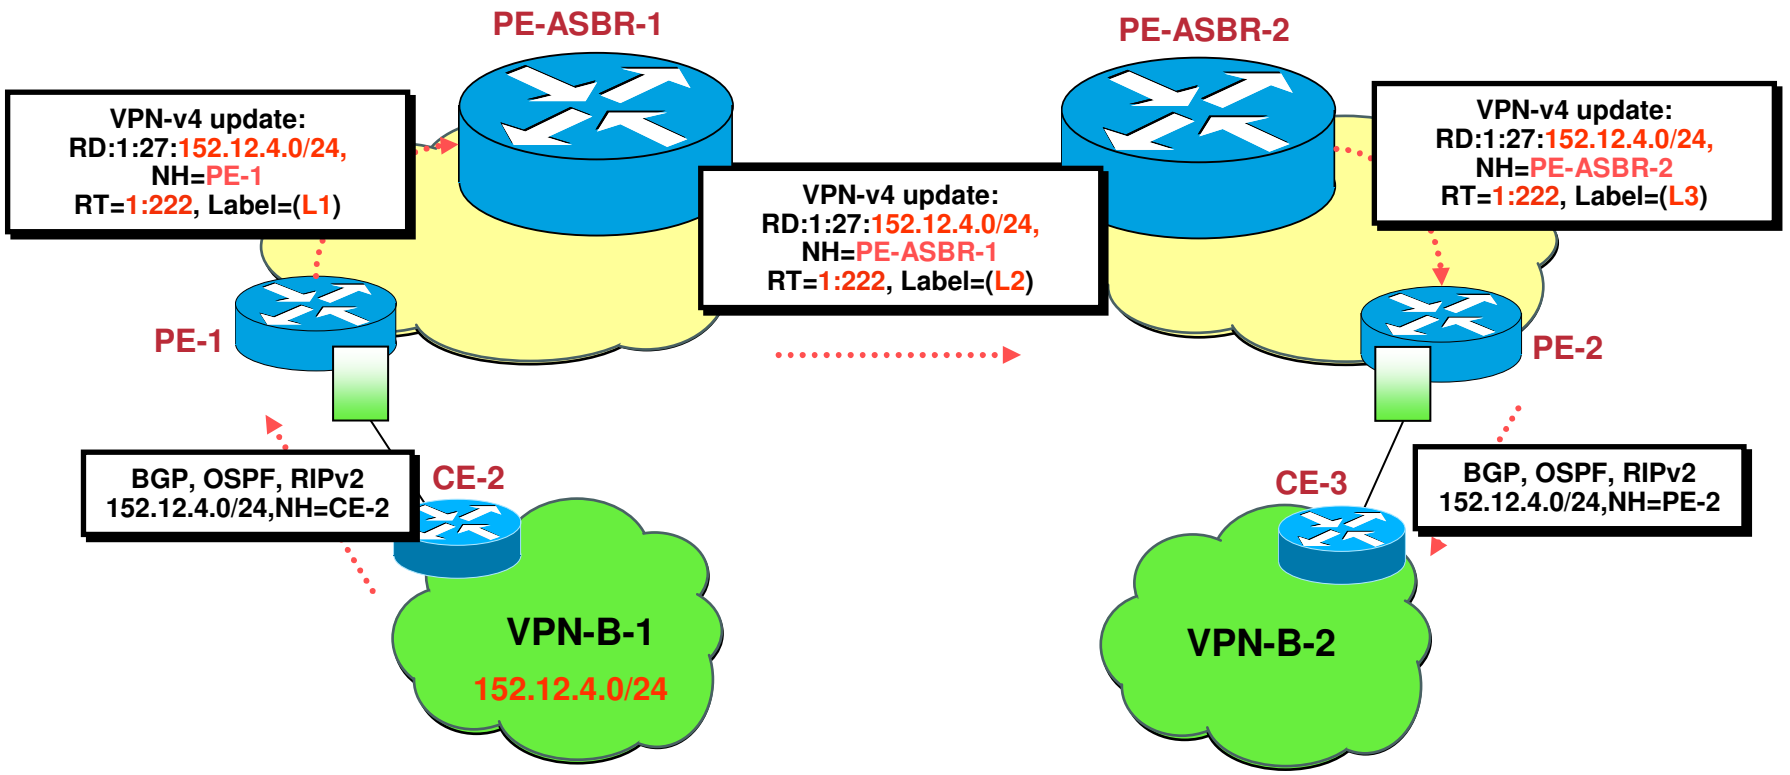

External MP-BGP for VPNv4 Forwarding Plane

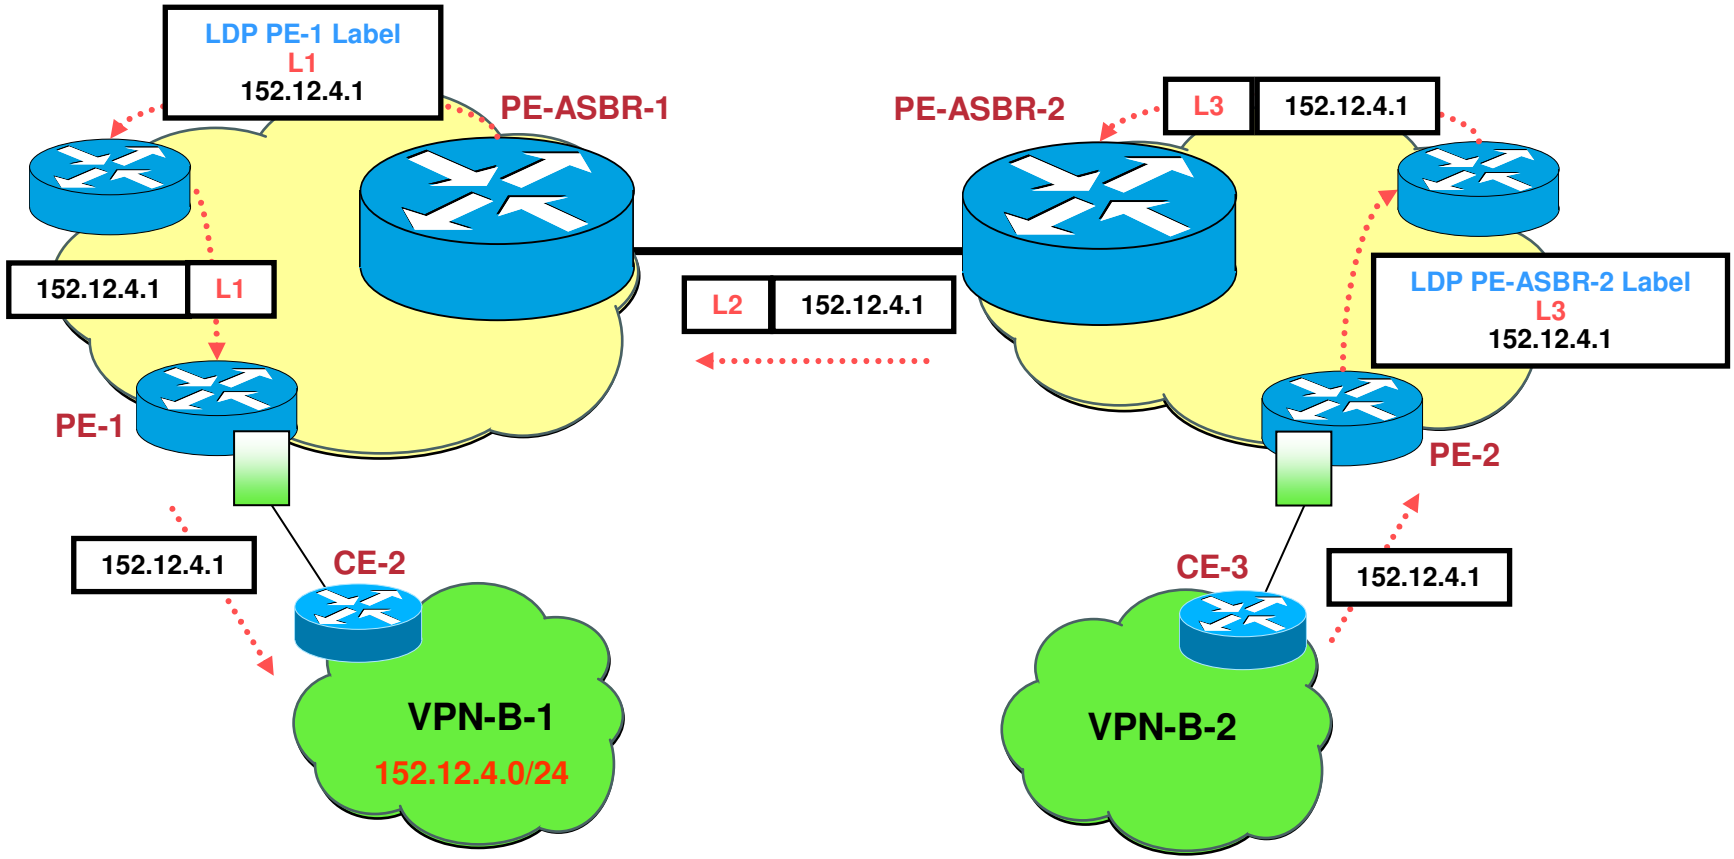

Why IPV4 BGP Label Distribution?

- Allows a VPN service provider network to exchange IPv4 routes with MPLS labels
- Use BGP to distribute labels associated with the routes at the same time it distributes the routes

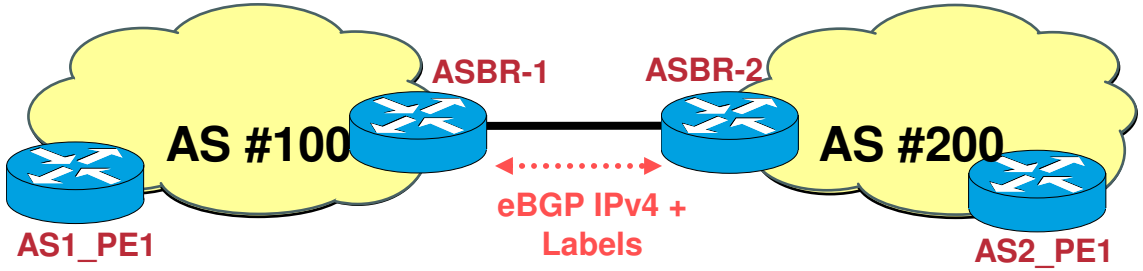

Benefits:

- Eliminate the need for any other Label distribution protocol between the two ASBRs
- Allow a non-VPN core network to act as a transit network for VPN traffic

Summary

Back to Back VRFs

- Scalability is an issue with many VPNs
  - One VRF & logical interface required per VPN client;
  - Gateway PE-ASBR must hold ALL routing information
- PE-ASBR must filter & store VPNv4 prefixes
  - Plus import into VRFs thus increasing MPLS & routing table memory
- No MPLS label switching required between providers
  - Standard IP between gateway PE-ASBRs;
  - No exchange of routes using MP-eBGP;
  - Simple solution, works today but limited in deployment scope

MP-eBGP for Prefix exchange

- Scalability less of an issue when compared to back-to-back VRF connectivity
  - Only one interface required between PE-ASBR routers;
  - No VRF requirement on any PE-ASBR router interfaces
- Automatic Route Filtering must be disabled
  - Hence filtering on RT values essential, and good filtering policy must be applied on EVERY PE-ASBR;
  - Import of routes into VRFs is not required which reduces the memory impact on PE-ASBR routers
- MPLS label switching required between providers
  - Routes exchanged using MP-eBGP;
  - Still simple, more scalable & works today