v4/v6 L3VPN over IP Core

- Tutorial

Madhusudan Nanjanagud

mnanjana@cisco.com

Agenda

- Introduction
- Feature Overview
- 6VPE – v6 L3VPN
- Service integration
- Configuration flow
- Summary

Introduction

- Providers have legacy infrastructure with IP core
- Provide Managed IP/ VPN services(L2 /L3),mVPN
- Converged common IP Core Backbone
- Simplify network Operations/Maintenance
- Simplified interface with other providers
- Secured infrastructure for service integrations

Feature Overview

- **RFC4364 based L3 VPN Services with IP Core**
- **L3VPN services could be for v4 or v6 over the same IP core**
- **Leverages most of the functionality from MPLS core based VPN**
- **Presence of IP Core is transparent to Customer Edge devices**
- **Two common approaches – mGRE & L2TPv3 Tunnel**

Feature Overview

- Multipoint to Point Tunnels are established between the edge routers through BGP signaling.
- Instead of manually configuring tunnels, “Tunnel Reachability Information” is signaled via BGP.
- Packets encapsulated with L2TPv3 header
- Session ID/Cookie (optional) values exchanged part of BGP updates
- No native L2TP signaling, BGP is used as the signaling protocol to convey encap header from PE to PE
- One Multipoint-to-point tunnel (configured per PE)

Tunnel Overlay Model

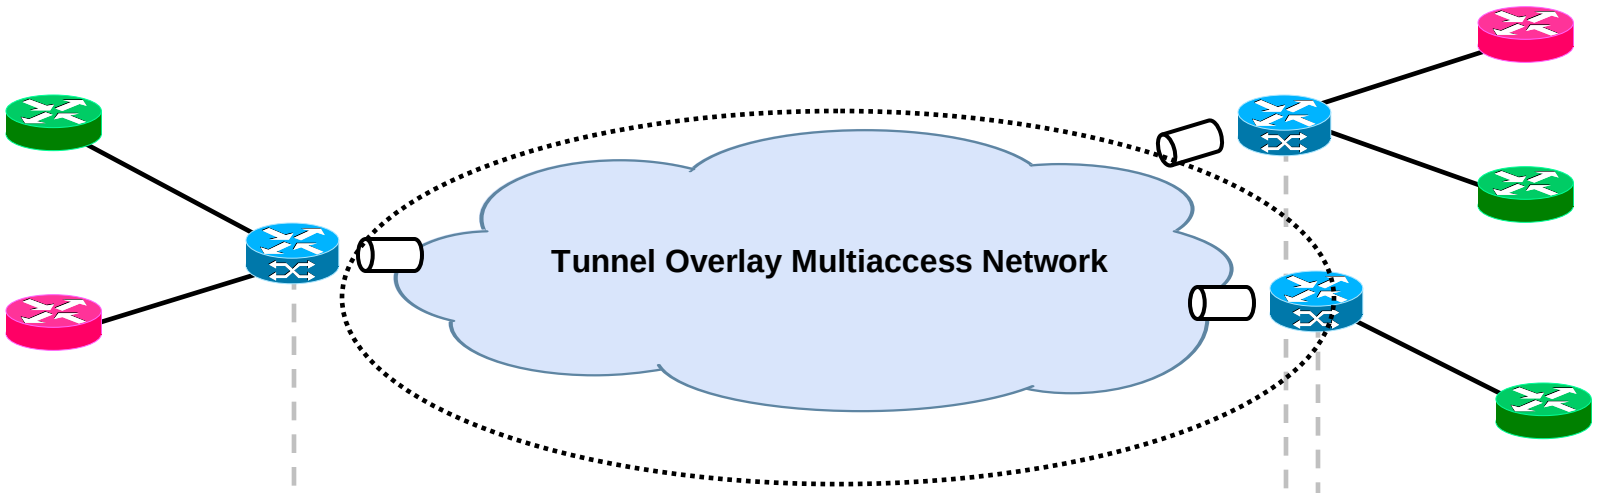

Tunnel overlay network switches VPN traffic between PEs

Transport network provides connectivity between PEs

Transport network is independent of VPN Service…

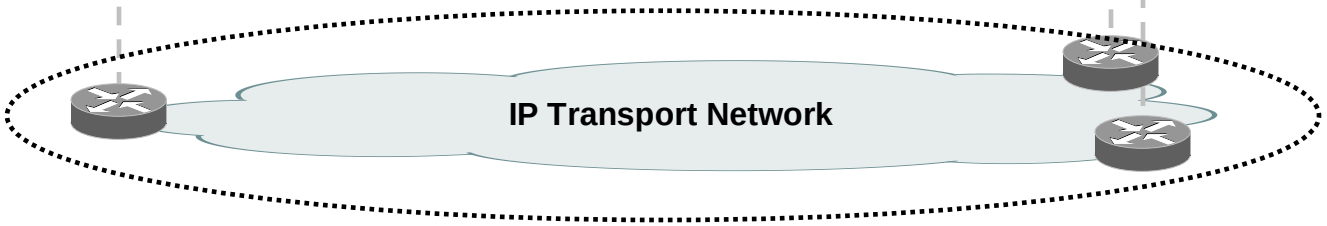

Encap PE Packet format

Tunnel IP >> To reach Egress PE with PID as 0x73 (L2TPv3)
Session ID/Cookie >> For L2TPv3 session control/Security
BGP VRF Label >> For remote End point
VPN IP >> Final customer destination prefix

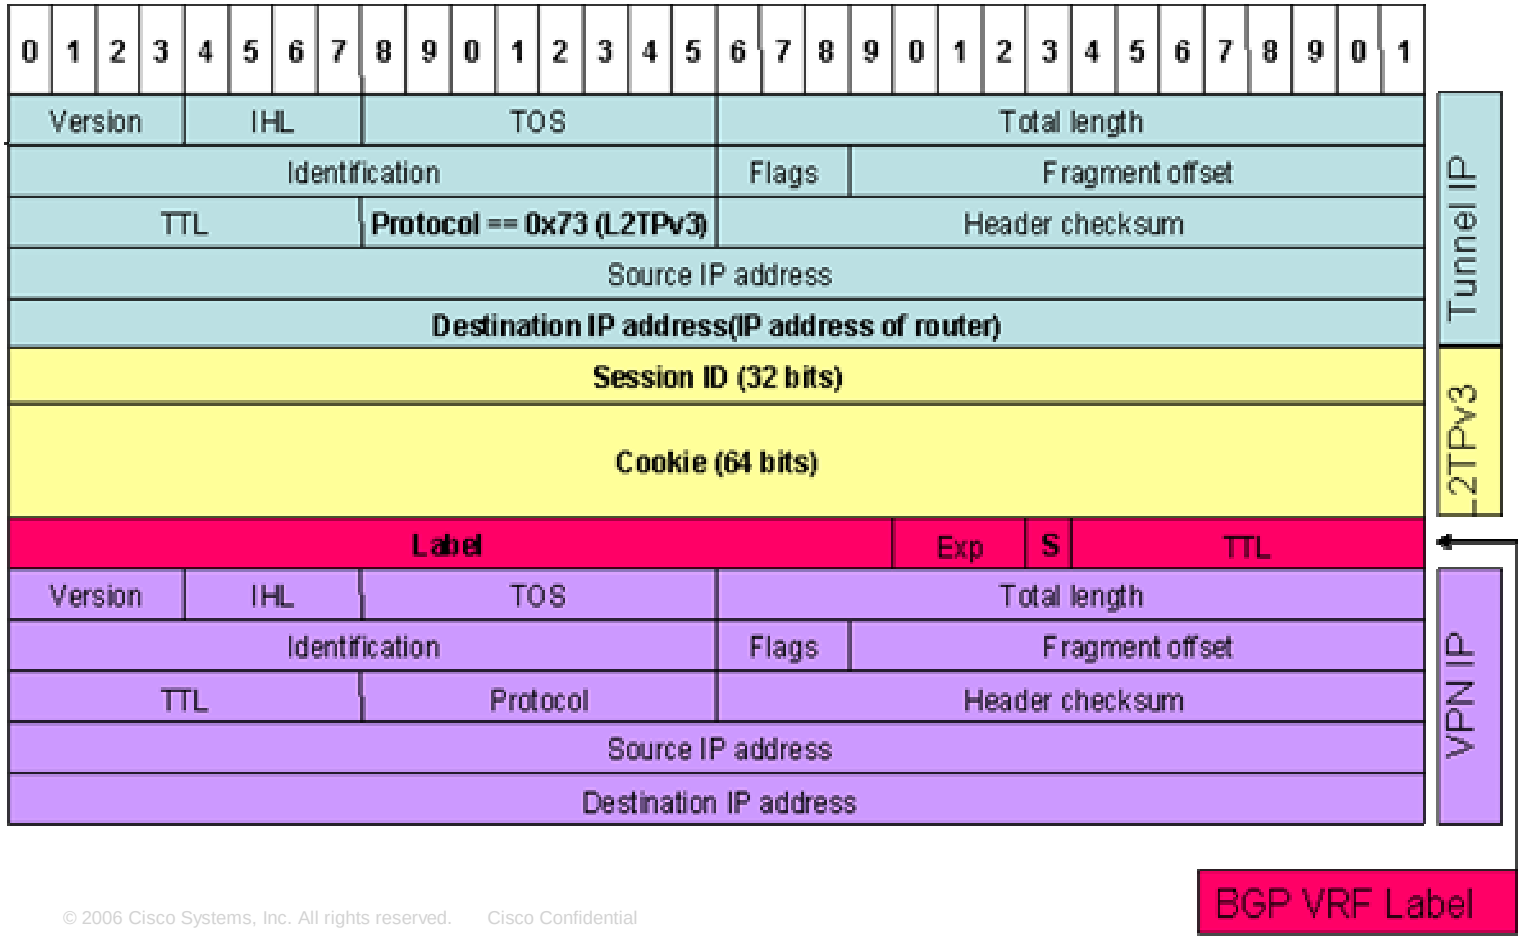

Session ID/Cookie Values

- **Part of the L2TPv3 header**
- **Optional “Cookie field” is a random 64 bit value in each data packet associated with session id**
- **To protect against a malicious blind attack, or inadvertent insertion of data into the tunnel stream.**

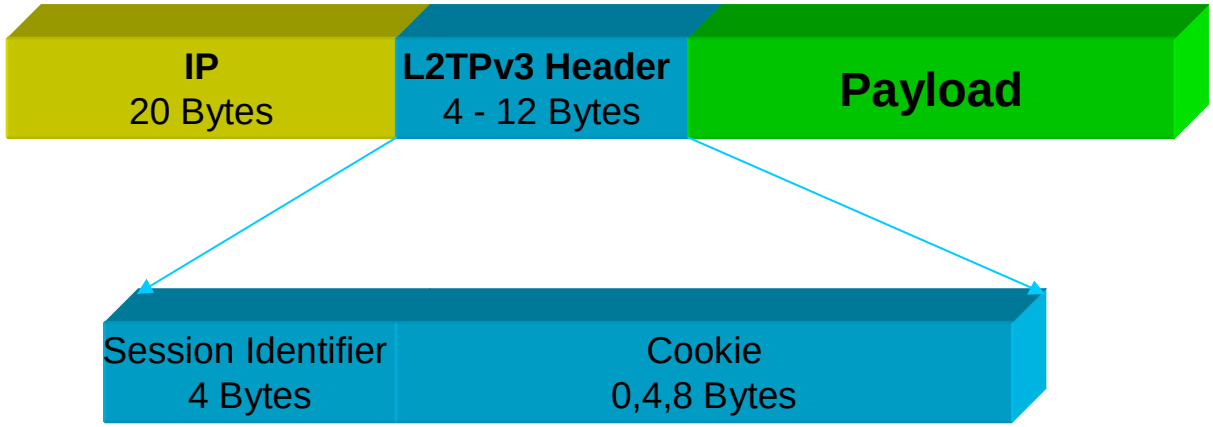

*** [REF] -** W. Mark Townsley NANOG presentations on L2TPv3

Comparison – MPLS Transport

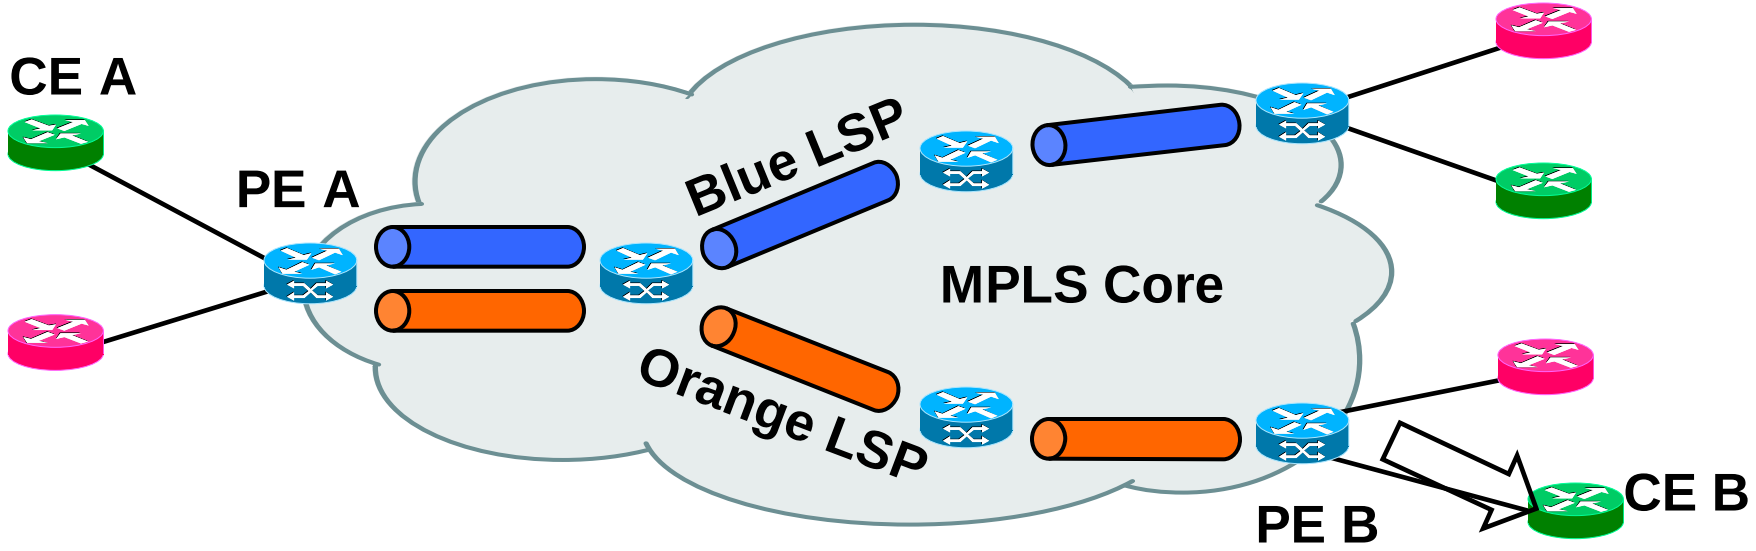

- PE-A Router imposes VPN label such that remote PE-B has enough information to determine that the payload should be delivered to CE-B in the Green VRF => UNTOUCHED
- PE-A injects VPN-labelled packet into orange LSP which will transport the payload transparently (almost!) to PE-B => LSP REPLACED BY IP TUNNEL

Comparison - MPLS Based Core – L3VPN flow

- As per RFC4364, MP-BGP would exchange customer prefixes between PE routers with Next Hop address.
- Ingress PE has LSP established with egress PE – LDP/RSVP.
- Ingress PE resolves the customer prefix within VRF table.
- Ingress PE would impose “IGP Label” + “VPN Label”+ “Customer Payload”
- Traffic to be drained over the core uplink interface – MPLS enabled links.
- Core P router does label Swap/Push/Pop
- Egress PE – Label lookup ->IP Lookup and then packet forwarded to egress interface.

IP Core - Ingress PE

- As per RFC4364, MP-BGP would exchange customer prefixes between PE routers with Next Hop address.
- Along with this, VRF labels are exchanged as well – vpnv4 AFI
- Ingress PE's routing table is populated with the prefix info and having Egress PE as Next Hop.
- Ingress can reach Egress PE as per IGP update – IP Core
- Ingress PE resolves the customer prefix with a multipoint L2TPv3 tunnel data structures.
- Ingress PE would impose "Tunnel IP/L2TPv3 header"+ "VPN Label"+ "Customer Payload"
- Traffic to be drained over the tunnel

IP Core - Core P router

- **P router's run standard IGP.**
- **Incoming traffic from Ingress PE would result in standard IP look over the Carrier IP.**
- **The Destination IP being Egress PE.**
- **Lookup resolves into packet getting forwarded as is to the NH.**
- **NH could be another P router or the Egress PE itself.**

- **Core P router – Not Aware of VPN prefix – Standard IP Lookup**

IP Core – Implementation – Egress PE

- **Egress PE would do IP Lookup on the Tunnel IP and then process the L2TPv3 header.**
- **L2TPv3 Session ID matched with Local Tunnel records and Cookie value compared with incoming packet.**
- **After admittance check, Session Type indicates L3 which results in BGP VPN Deagg thread.**
- **BGP Label lookup would result in VRF and its associated outgoing interface after VPN IP lookup.**
- **The Tunnel IP, L2TPv3 Header, BGP VRF label are all removed and the IP packet is forwarded.**

Forwarding Into and Out Of Tunnel

Encapsulation Behaviour

- Forwards IP packet by looking up IP destination in 'customer' VRF. Add VPN label and tunnel encap via standard adjacency rewrite
- Forward packet using outer header IP destination via global routing table

Decapsulation Behaviour

- Forwarding matches a receive entry that matches the local tunnel address and finds L2TPv3 protocol type in IP header
- Lookup in VPN forwarding table based on the MPLS label in the packet
- Forward the inner IP packet using the interface derived from MPLS label lookup

Control Plane Events

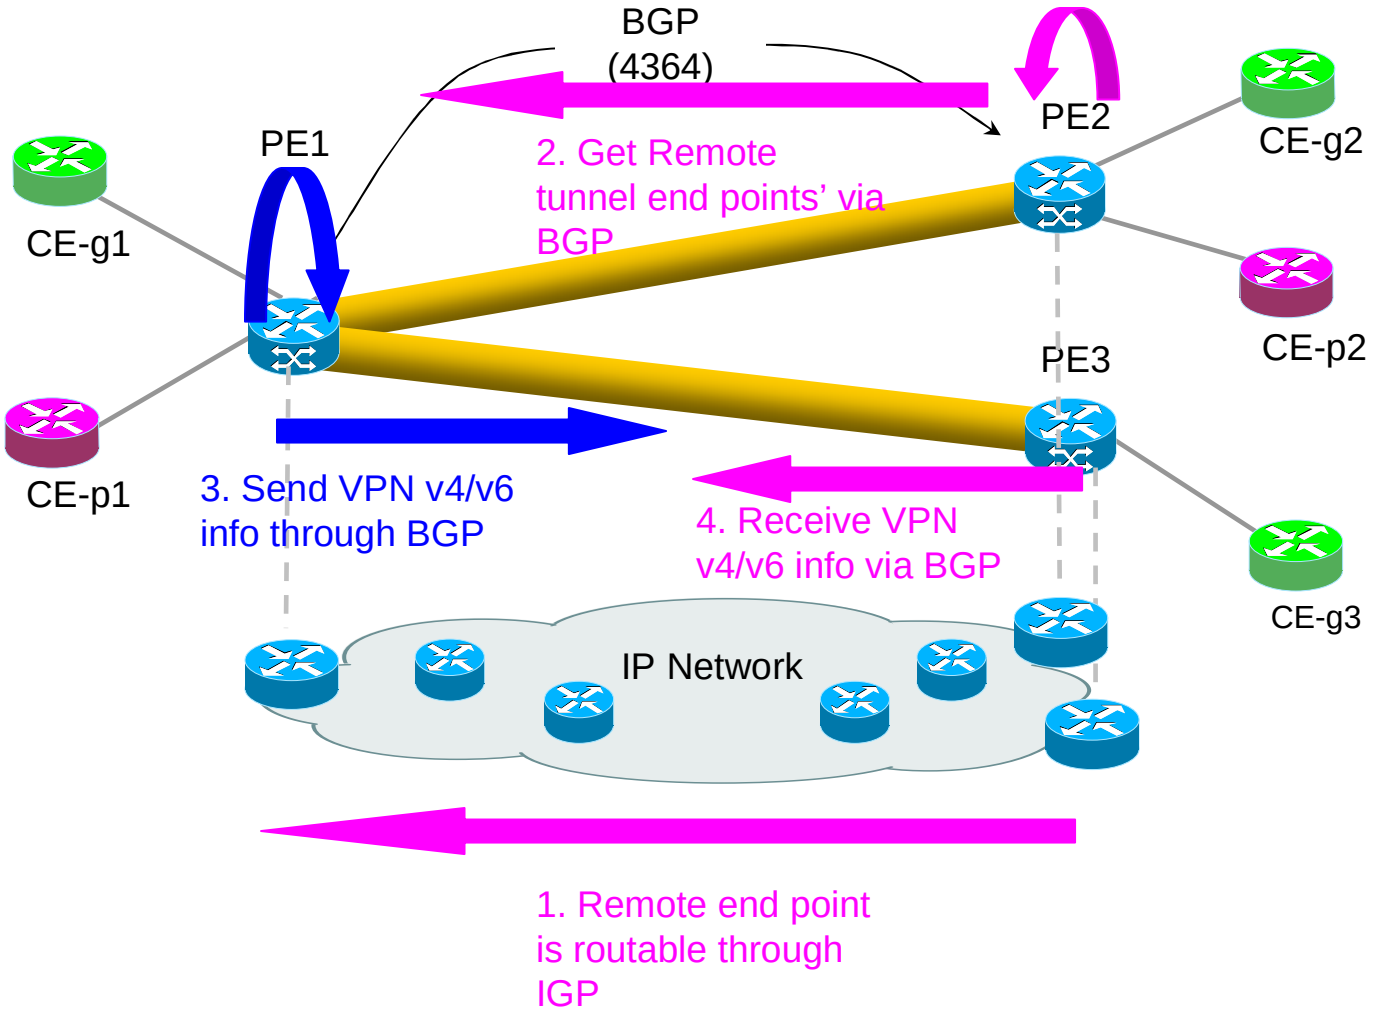

Control Plane Operation

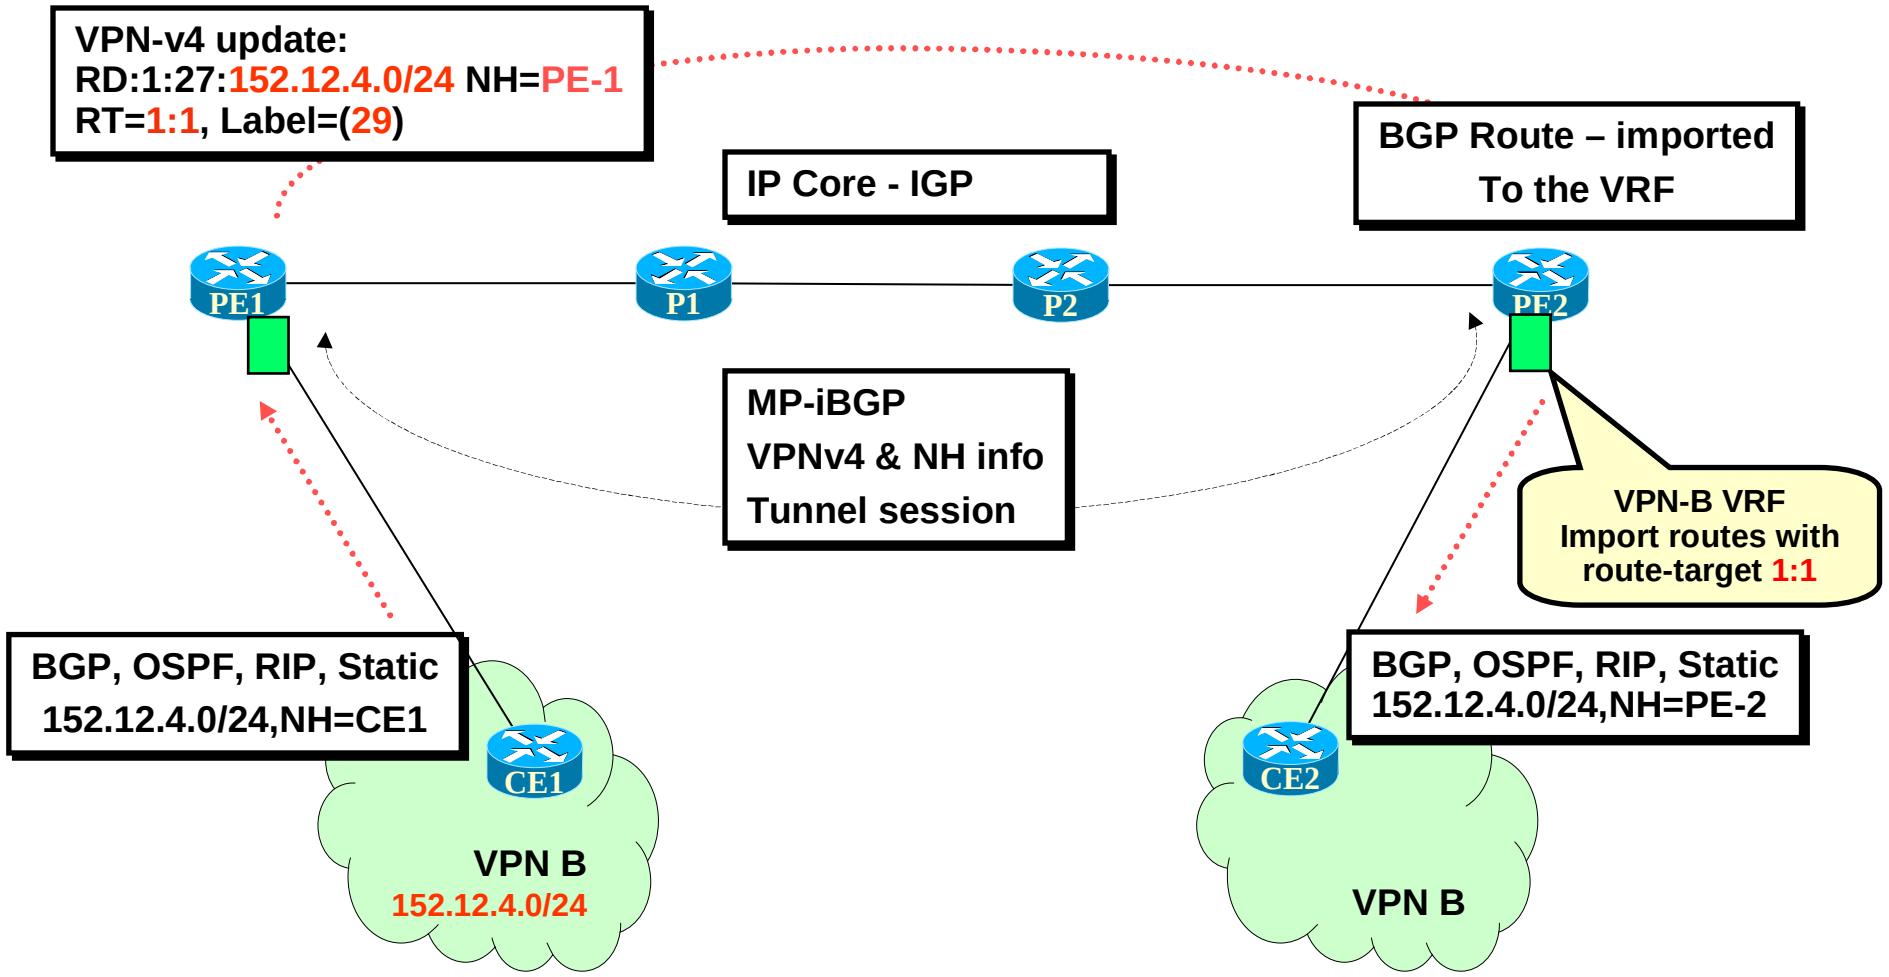

Forwarding Plane Operation

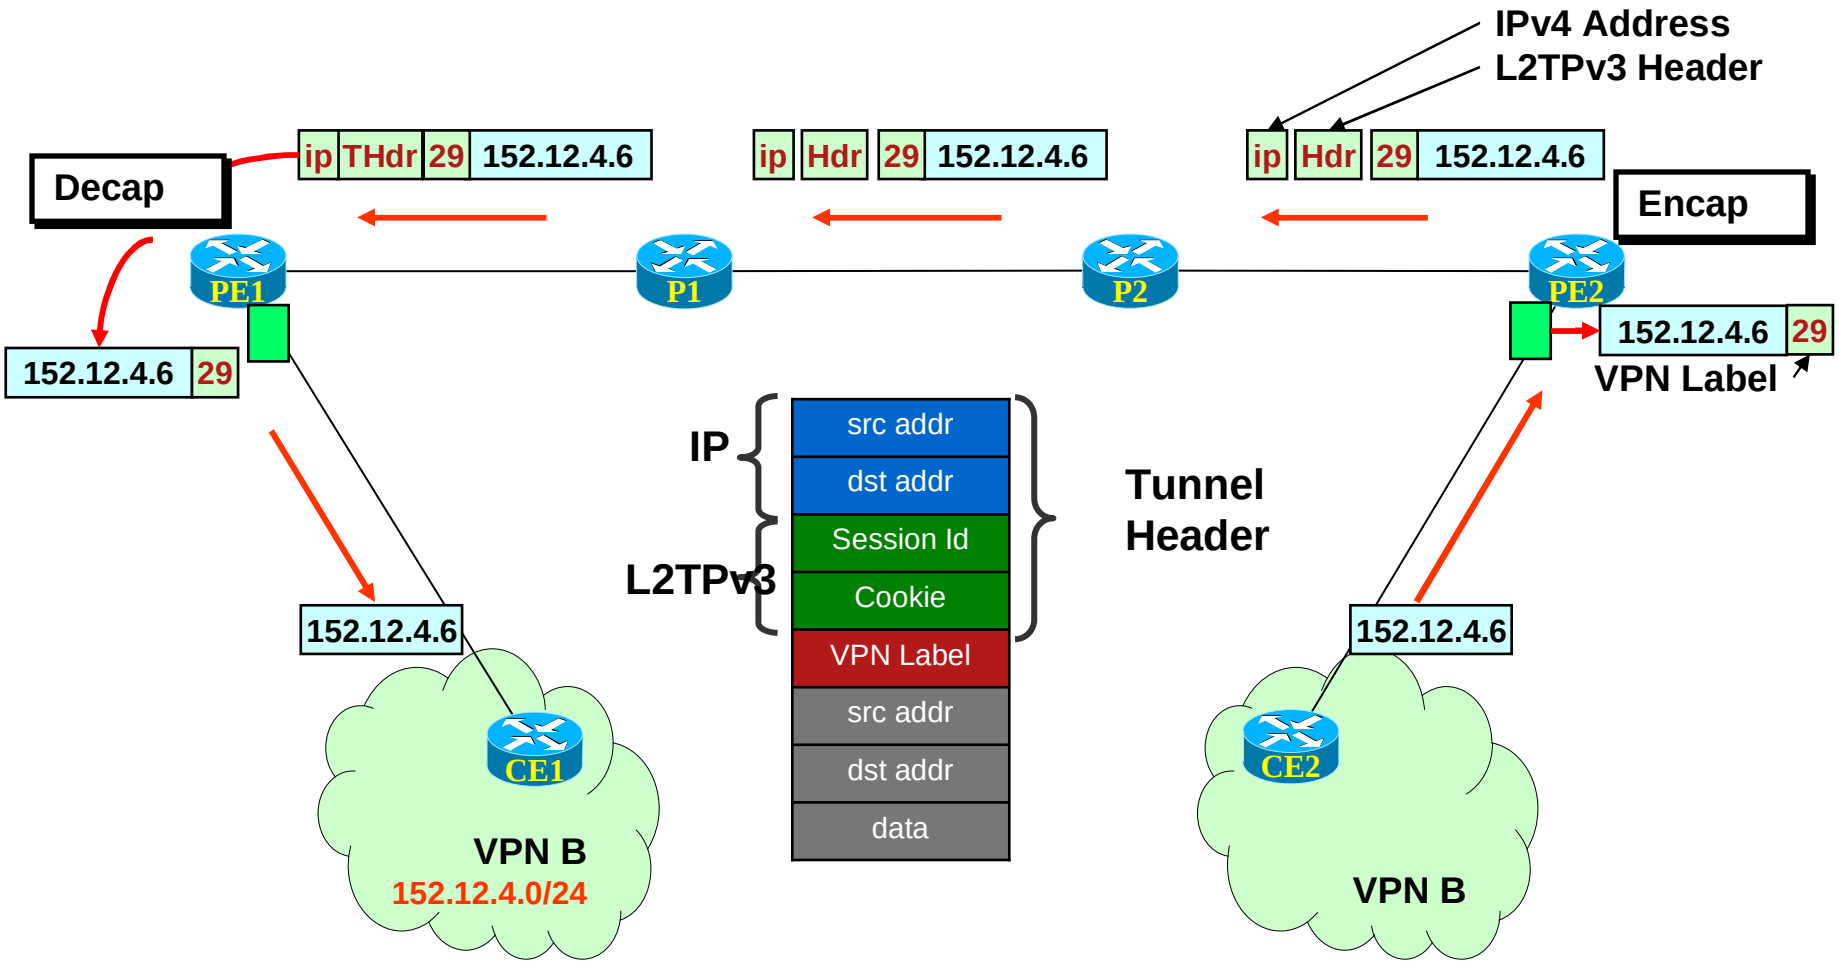

6VPE Analysis – v6 L3VPN

- **Multiple approaches for IPv6 over v4 core**
  - **IPv6 over EoMPLS/AToM**
  - **IPv6 over L2TPv3**
  - **IPv6 CE-to-CE IPv6 over IPv4 tunnels**
  - **IPv6 provider edge router (6PE) over MPLS/IP**
  - **IPv6 VPN provider edge (6VPE) over MPLS/IP**

6VPE Feature overview

- 6VPE is L3VPN services for v6 VPN customers
- RFC4659 – L3VPN Extension for v6 VPN
- Could be MPLS based core or IP core
- No dual stacking in core! No v6 in Core
- Feature and functionality parity with v4 L3VPN
- v6 L3VPN services offered over the same v4 L3VPN infrastructure

6VPE : MP-BGP role

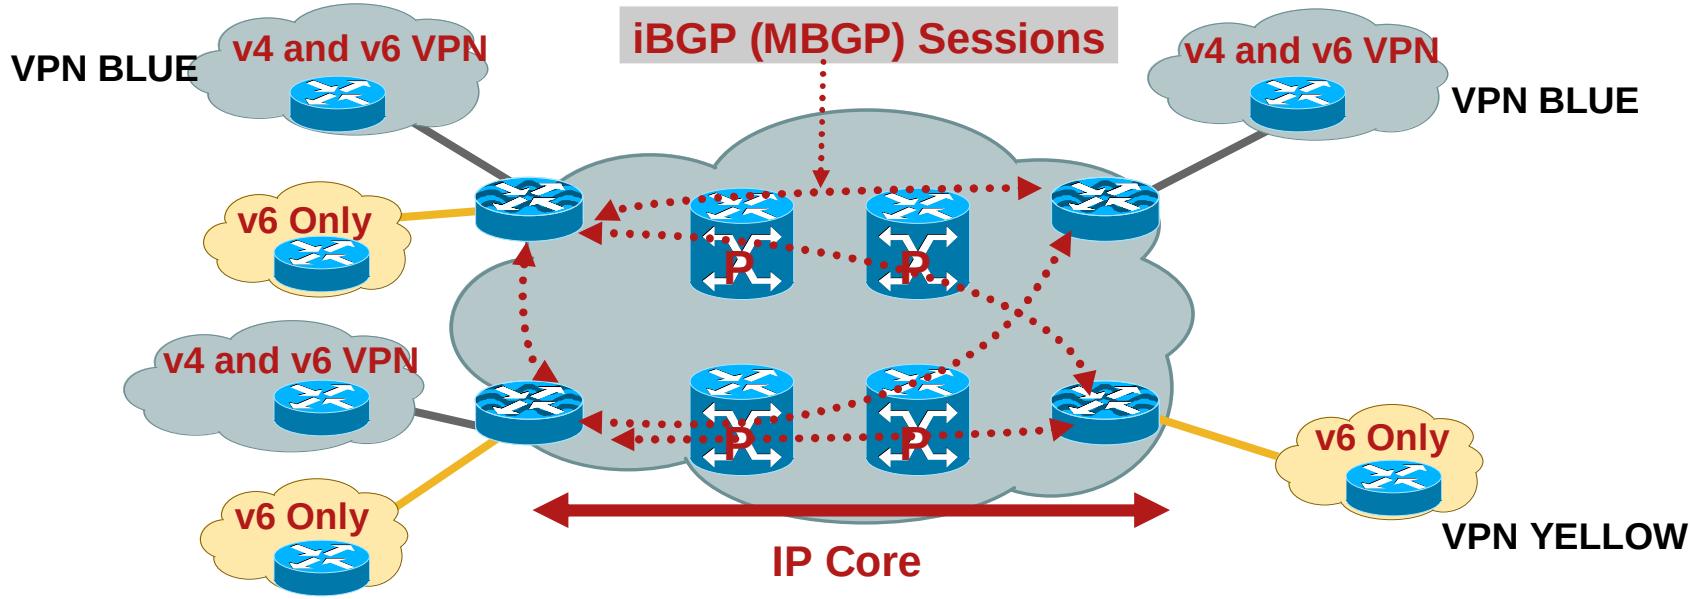

- 6VPE ~ IPv6 customer payload + BGP VPN label + IP transport
- 6VPE is an implementation of RFC4659
- VPNv6 address:

  Address including the 64 bits route distinguisher and the 128 bits IPv6 address
- Next Hop is carried as RD:v4-mapped-v6-address

- MP-BGP VPNv6
  AFI "IPv6" (2), SAFI "VPN" (128)
- VPN IPv6 MP_REACH_NLRI

  With VPNv6 next-hop (192bits) and NLRI in the form of <length, IPv6-prefix, label>
- Encoding of the BGP next-hop

6VPE : MP-iBGP role

- **MP-BGP is modular to facilitate distinct peering relationships**
- **Using the *Address Family Identifier* (AFI) – VPNv4 & VPNv6 AFI**
- **Two provider edge routers to exchange labeled IPv6 VPN prefixes, they must use BGP capabilities negotiation for vpnv6 AFI.**
- **MP-iBGP peering would include typically vpnv4 AFI and vpnv6 AFI capability exchange along with other associated AFI/ SAFI's like Tunnel SAFI for L3VPN over IP Core**
- **6VPE could have iBGP peering with RR which serve both vpnv4 and vpnv6 reflection.**

6VPE : Flow

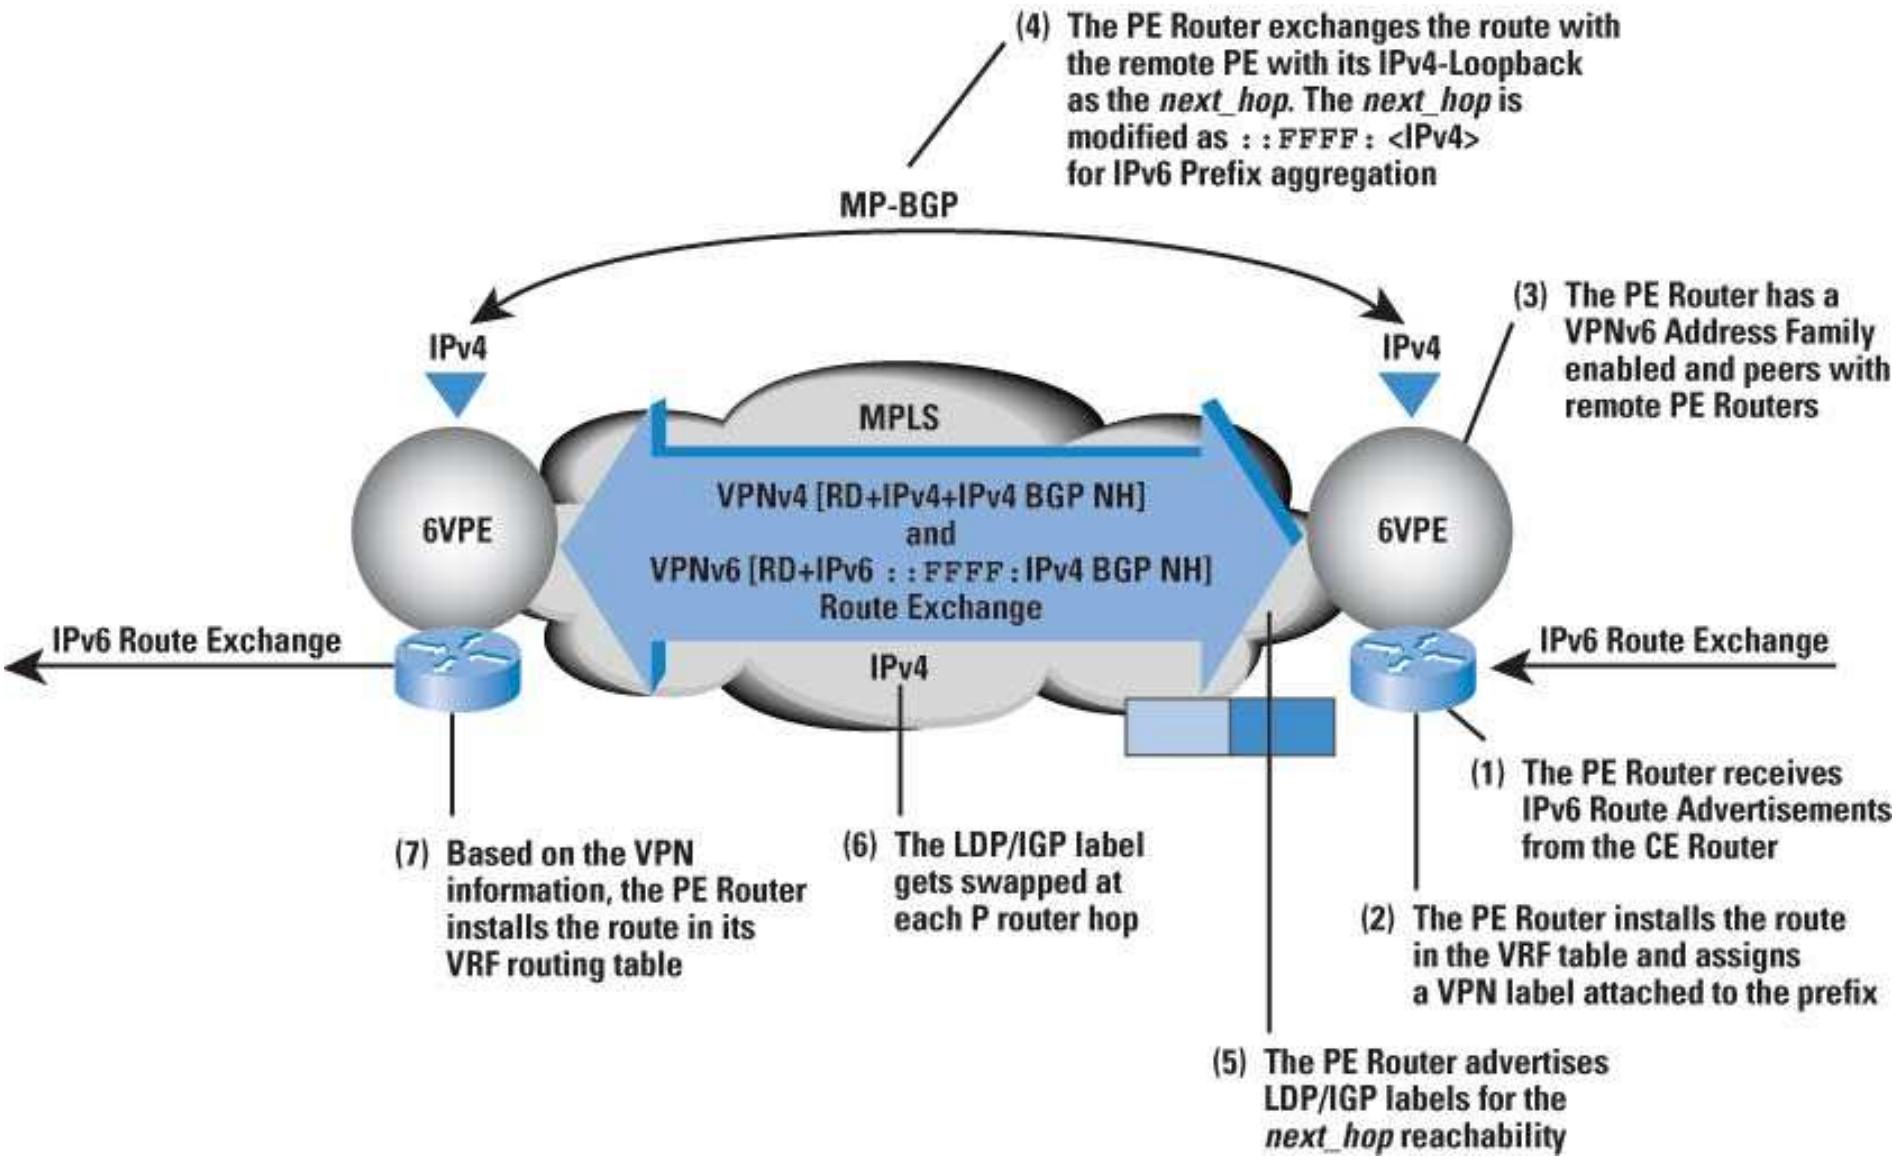

6VPE Forwarding Operation

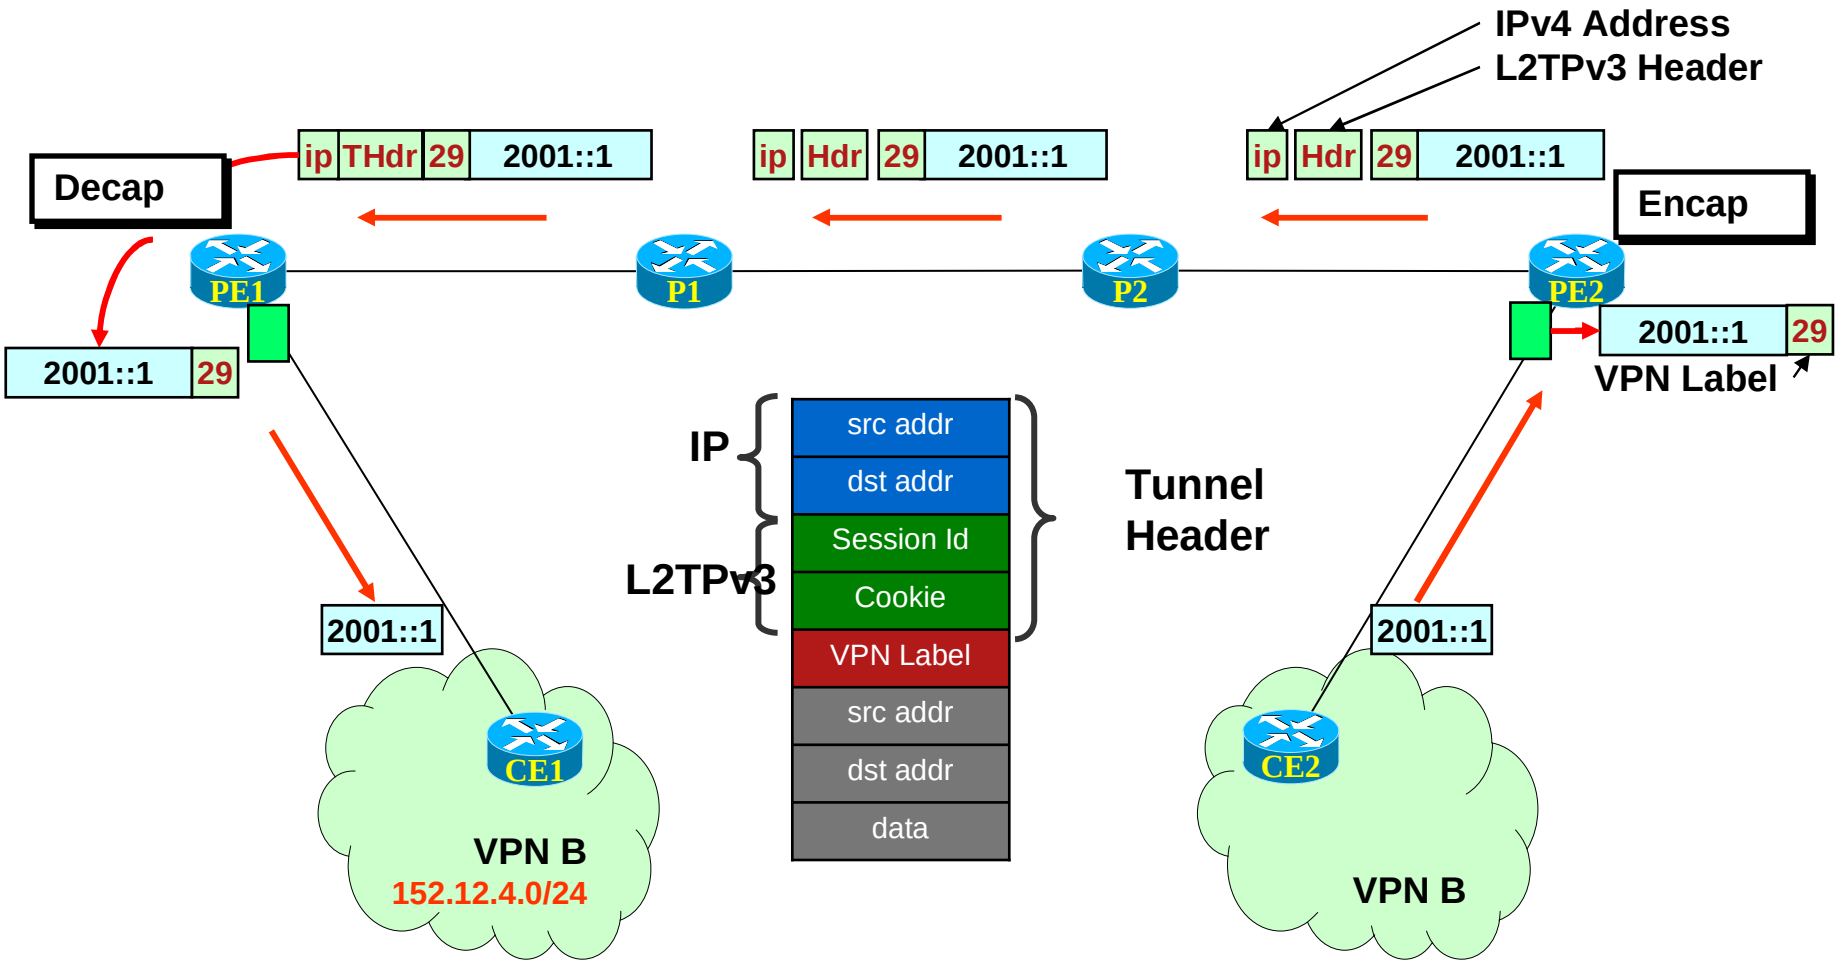

6PE : Global v6 services

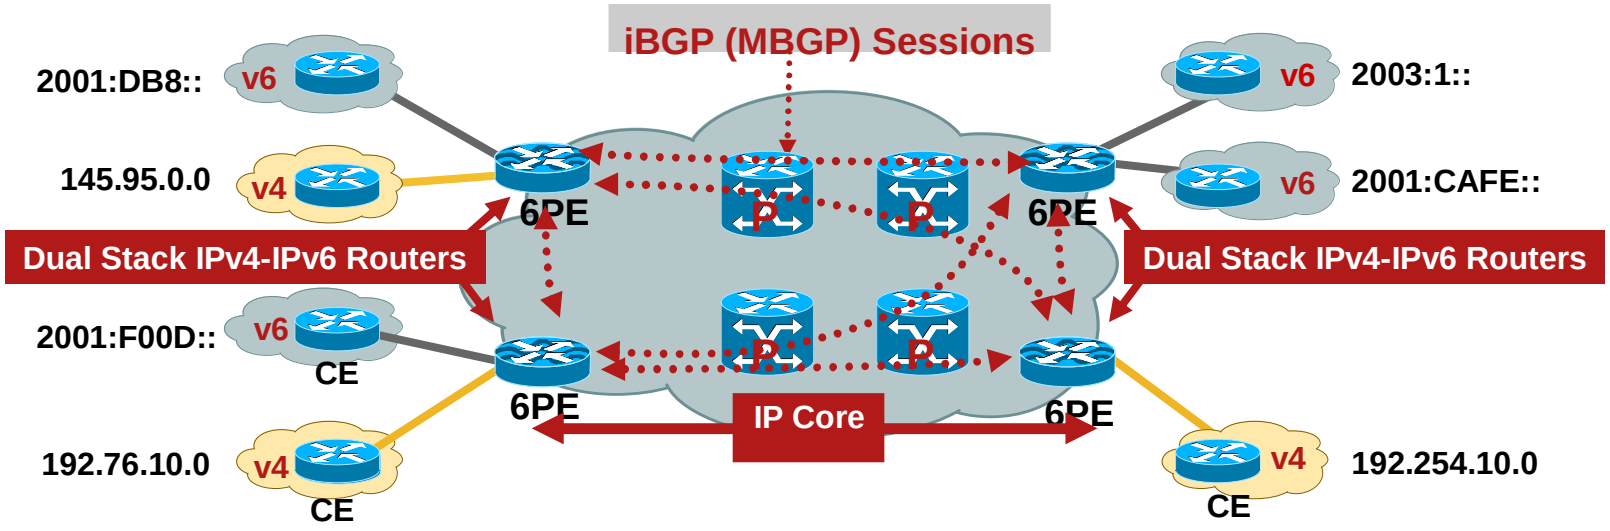

- **IPv6 global connectivity over and IPv4 core**
- **PE to support dual stack/6PE**
- **IPv6 reachability exchanged among 6PEs via iBGP**
- **IPv6 packets transported from 6PE to 6PE**

Sample Configuration Flow

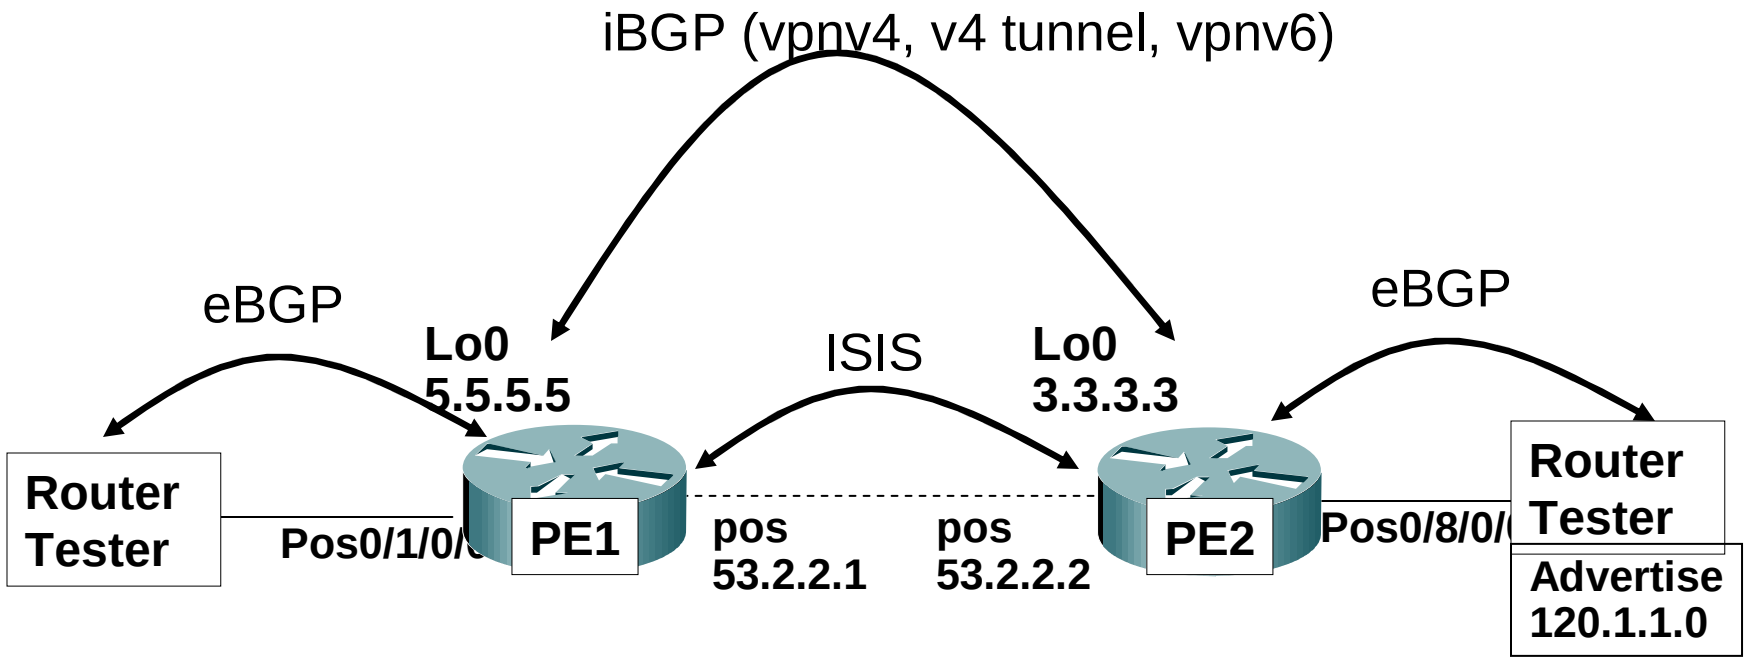

BGP Config – vpnv4 & vpnv6 AFI

(Cisco config guide reference)

PE1 Config – BGP Related

```
router bgp 1803
 bgp router-id 5.5.5.5
 address-family ipv4 unicast
 !
 address-family ipv4 tunnel
 !
 address-family vpnv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family vpnv6 unicast
 !
 neighbor 3.3.3.3
  remote-as 1803
  update-source Loopback0
  address-family ipv4 unicast
  !
  address-family ipv4 tunnel
  !
  address-family vpnv4 unicast
  !
  address-family vpnv6 unicast
  !
```

PE2 Config – BGP Related

```
router bgp 1803
 bgp router-id 3.3.3.3
 address-family ipv4 unicast
 !
 address-family ipv4 tunnel
 !
 address-family vpnv4 unicast
 !
 address-family ipv6 unicast
 !
 address-family vpnv6 unicast
!
 neighbor 5.5.5.5
  remote-as 1803
  update-source Loopback0
  address-family ipv4 unicast
  !
  address-family ipv4 tunnel
  !
  address-family vpnv4 unicast
  !
  address-family ipv6 unicast
  !
  address-family vpnv6 unicast
  !
```

Local VRF (PE-CE) config

(Cisco config guide reference)

```
RP/0/8/CPU0:PE1#sh run int poS 0/1/0/0
interface POS0/1/0/0
 vrf vpn55
 ipv4 address 15.1.1.1 255.255.255.0
 ipv6 address 1511::1/64
 ipv6 enable
 encapsulation hdlc
 keepalive disable
!
RP/0/8/CPU0:PE1#sh run vrf vpn55
vrf vpn55
 address-family ipv4 unicast
  import route-target
   1803:55
  !
  export route-target
   1803:55
  !
 !
 address-family ipv6 unicast
  import route-target
   1803:55
  !
  export route-target
   1803:55
```

```
RP/0/8/CPU0:PE1#sh run router bgp vrf vpn55
 vrf vpn55
  rd 1803:55
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
  !
  neighbor 15.1.1.2
   remote-as 7000
   address-family ipv4 unicast
    route-policy allpass in
    route-policy allpass out
   !
  !
  neighbor 1511::2
   remote-as 7001
   address-family ipv6 unicast
    route-policy allpass in
    route-policy allpass out
   !
```

Service Integration

- **With this L3VPN v4/v6 over IP infrastructure, new services can be offered**
- **Layer 2 VPN – L2TP Signaling**
- **Native Multicast or Multicast VPN – BGP MDT SAFI**
- **Quality of Service – Edge & Core**
- **Service Provider Edge feature set can be integrated over the same infrastructure – Transparent to customers**
- **CsC and Inter-AS – Parity with MPLS based core!**
- **Transit carriers can be IP Core based while Baby Carriers could be MPLS core based.**

Summary

- **Investment Protection for Providers – IP Core**
- **Enabling L3VPN –v4 & v6 services over the legacy infrastructure**
- **Newer services can be easily integrated – L2VPN multicast VPN / Native**
- **Network Troubleshooting / Maintenance simplified**
- **Parity with MPLS based services**
- **Implementation on networking gear is simplified**

Thank You!