

MPLS VPN (RFC2547bis) Seminar

Umberto Poschi

umberto@poschi.it

uposchi@cisco.com

MPLS

What it Is Not and What It Is

- ***MPLS IS NOT***

- a mechanism that allows router to forward packets in a faster way.....router are already faster.....

- only the integration of IP and ATM (this is one of the MPLS applications)

- ***MPLS IS***

- based on a forwarding algorithm simpler than traditional IP forwarding algorithm

- the enabler of much more new services and functionalities than the traditional IP paradigm

MPLS is a Component

- MPLS is part of the overall solution:

- MPLS VPN

MPLS+BGP+Address extensions+multiple forwarding tables

- Traffic Engineering

MPLS+OSPF/IS-IS extensions+Constrained SPF

- DiffServ aware TE

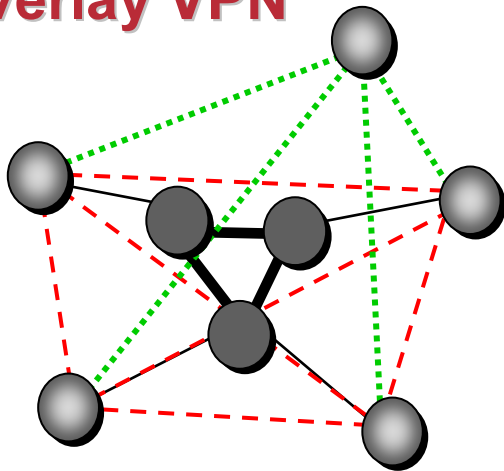
MPLS+OSPF/IS-IS extensions+Constrained SPF+DiffServ

MPLS VPN Benefit

- **MPLS VPN combines the best features of overlay VPN and peer-to-peer VPN:**
 - Provider Edge routers participate in customer routing, guaranteeing optimum routing between sites and easy provisioning.
 - Provider Edge routers carry a separate set of routes for each customer (similar to the dedicated Provider Edge router approach).
 - Customers can use overlapping addresses.


MPLS VPN Benefit: Lower Operational Cost

Overlay VPN



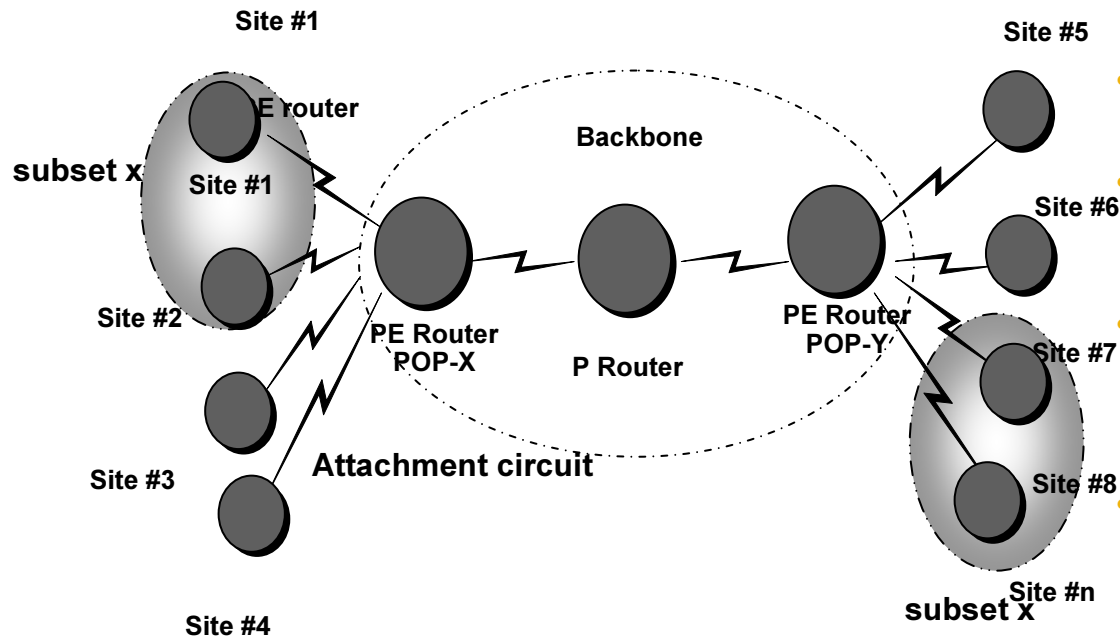
- Update traffic matrix
- Add (n-1) PVCs to connect new CPE
- Resize full PVC mesh
- Update OSPF design
- Reconfigure each CPE for new L3 topology

MPLS-based VPN



- Configure new CPE
- Update edge LSR

RFC2745bis Network Components and Definition of VPN



- Ip1: A set of sites are attached to a common network: the Backbone
- Ip2: Defined policies allows to create subsets of this set (i.e. subset x)
- St1: Two sites may have IP connectivity over the Backbone ONLY IF at least one subset contains them both
- St2: The subsets are VPNs -> Two sites have IP connectivity over the common backbone only if there is some VPN which contains them both
- St3: If all the sites in a VPN are owned by the same enterprise the VPN maybe thought as an INTRANET. If various sites in a VPN are owned by different enterprises the VPN maybe thought as an EXTRANET

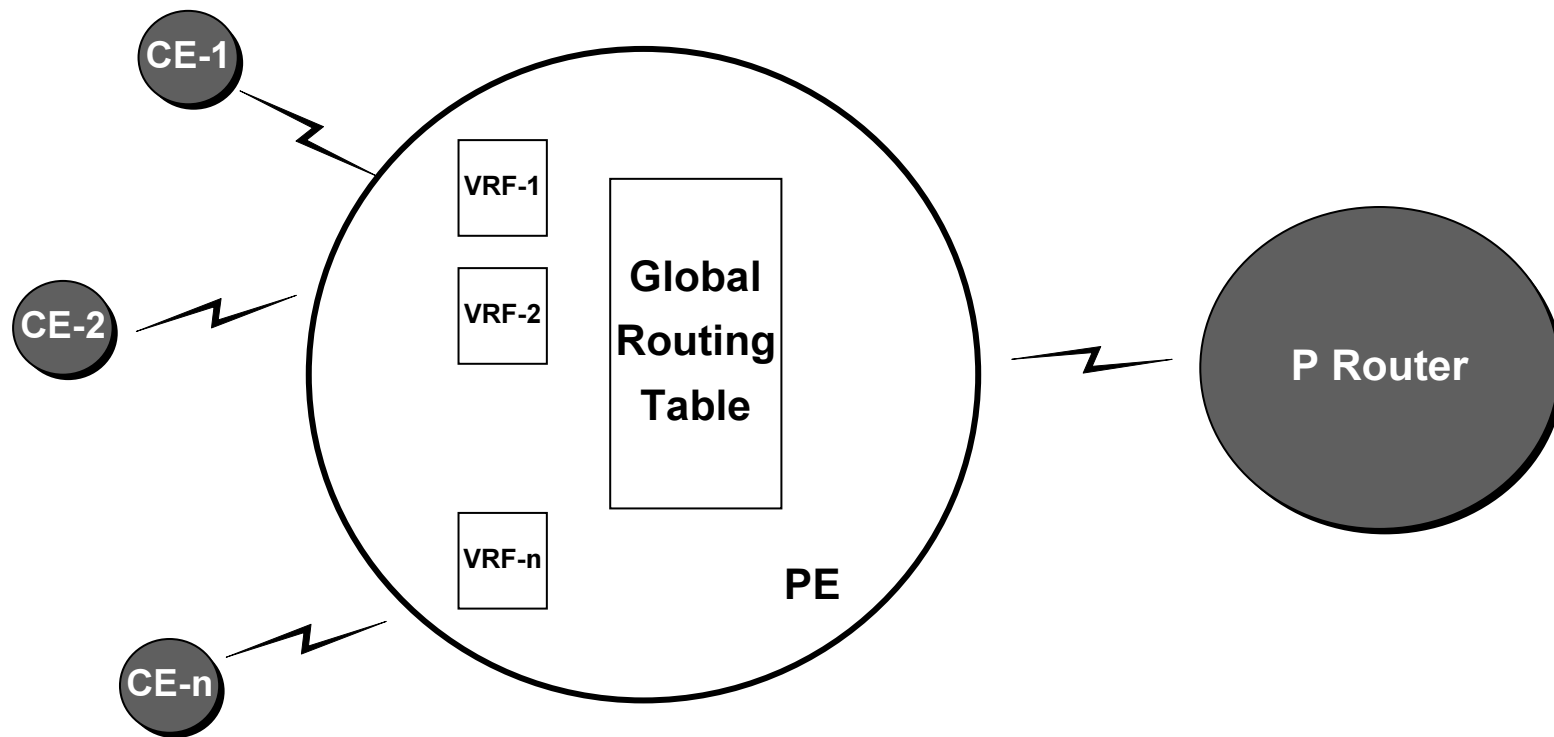
Sites and CEs

- **Site is consider to be a set of IP systems that have mutual IP connectivity that does not requires the Backbone (as previously defined).**
- **CE is always regarded as being in a single site.**
- **A site may belong to multiple VPNs**

PE Architecture

How to constrain distribution of routing information at PE that has sites from multiple (disjoint) VPNs attached to it.

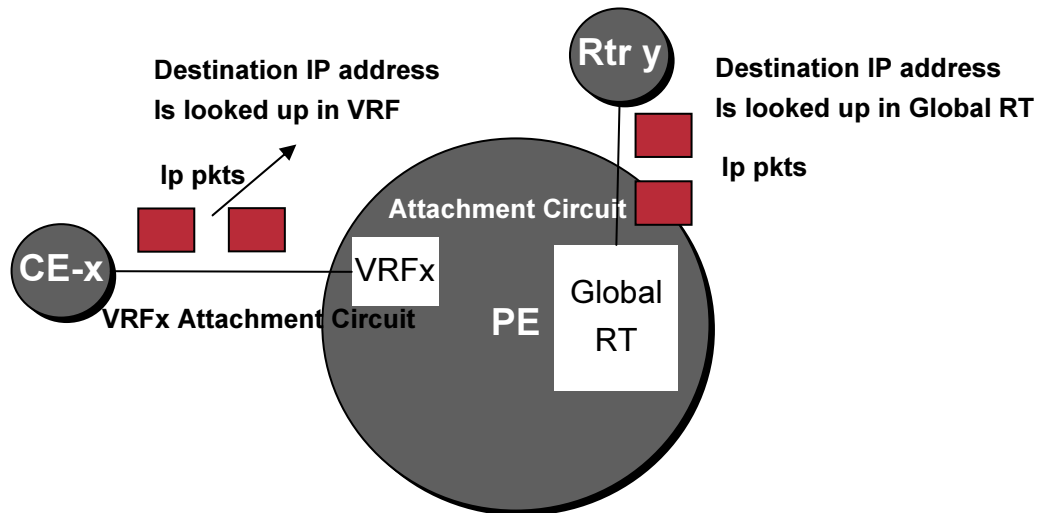
Single Forwarding Table on PE doesn't allow per VPN segregation of routing information



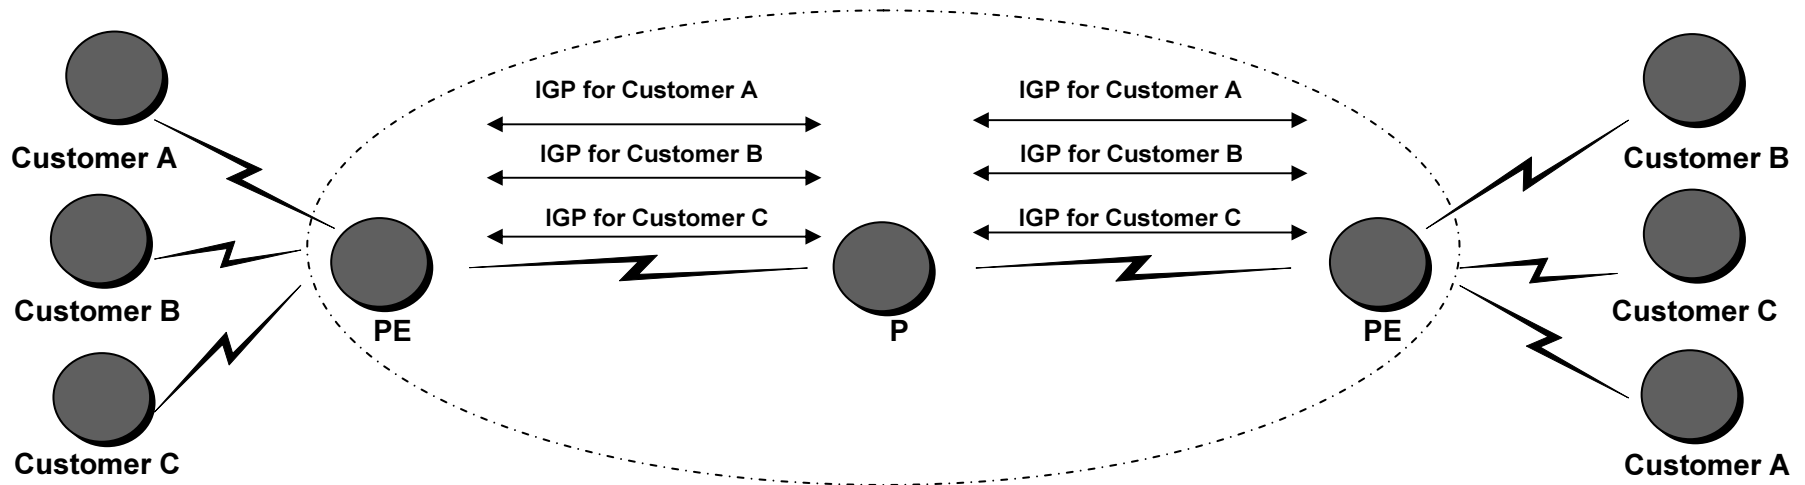
VRF and Attachments Circuits

Each PE has several routing Table:

- One Global Routing Table
- Multiple VPN Routing and Forwarding Table (VRFs)
 - An attachment circuit associated with a VRF is known as VRF attachment circuit.



Routing Information Propagation Across the P-Network



Q: How will PE routers exchange customer routing information?


A1: Run a dedicated Interior Gateway Protocol (IGP) for each customer across P-network.

Wrong answer:

- The solution does not scale.
- P routers carry all customer routers.

Routing Information Propagation Across the P-Network (cont.)

A dedicated routing protocol used to carry customer routes



Q: How will PE routers exchange customer routing information?


A2: Run a single routing protocol that will carry all customer routes inside the provider backbone.

Better answer, but still not good enough:

- P routers carry all customer routes.

Routing Information Propagation Across the P-Network (cont.)

A dedicated routing protocol used to carry customer routes between PE routers



Q: How will PE routers exchange customer routing information?


A3: Run a single routing protocol that will carry all customer routes between PE routers. Use MPLS labels to exchange packets between PE routers.

The best answer:

- P routers do not carry customer routes; the solution is scalable.

Routing Information Propagation Across the P-Network (cont.)

A dedicated routing protocol used to carry customer routes between PE routers



Q: Which protocol can be used to carry customer routes between PE routers?


A: The number of customer routes can be very large. BGP is the only routing protocol that can scale to a very large number of routes.

Conclusion:

BGP is used to exchange customer routes directly between PE routers.

Routing Information Propagation Across the P-Network (cont.)

A dedicated routing protocol used to carry customer routes between PE routers



Q: Customers can have overlapping address spaces. How will information about the same subnet of two customers be propagated via a single routing protocol?

A: Customer addresses are extended with a 64-bit prefix (route distinguisher—RD) to make them unique. Unique 96-bit addresses are exchanged between PE routers.

Address Extension

- **How to support VPNs without imposing constraints on address allocation/management within VPNs (e.g., allowing private address space [RFC1918]) ?**
 - constrained distribution of routing information uses BGP
 - BGP is designed with the assumption that addresses are unique

Address Extensions

VPN-IPv4 address

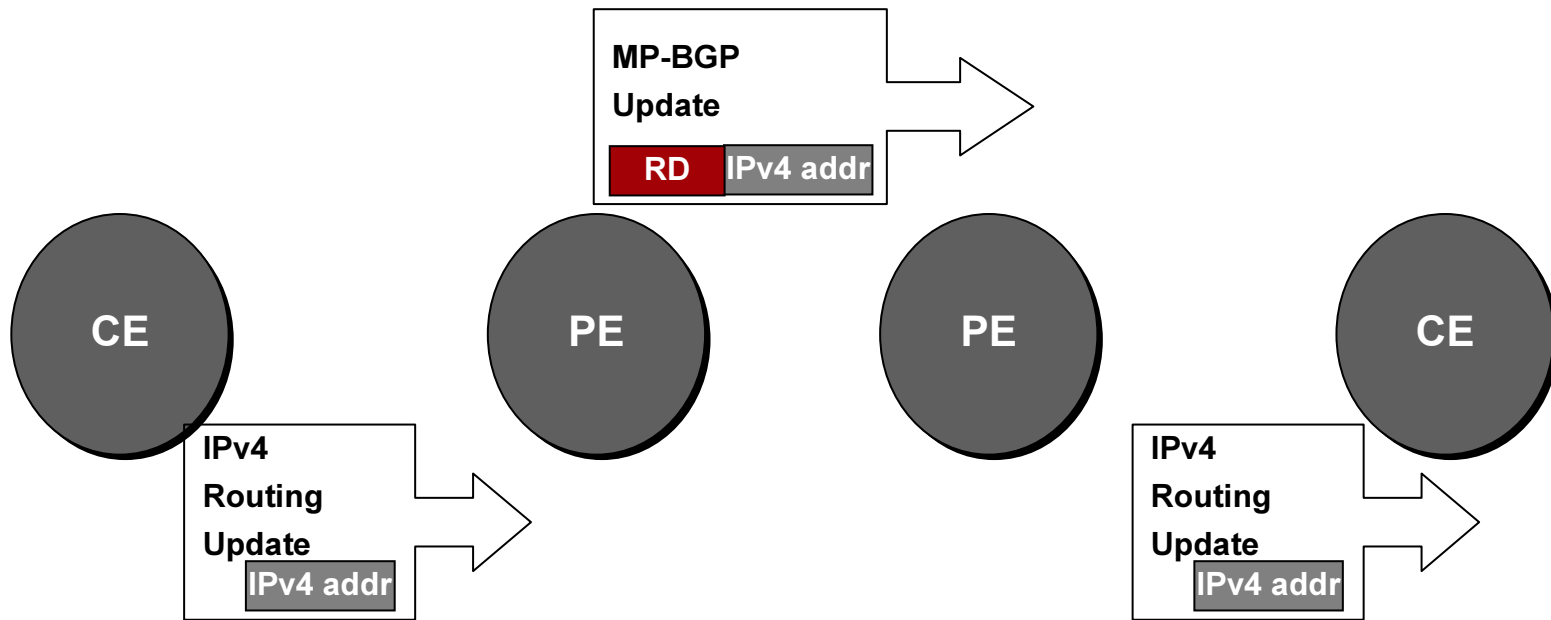
- **MP-BGP Extensions** allow BGP to carry routes from multiple address family
- **VPN-IPv4 address family** “make” the IPv4 address unique inside the Provider Backbone
- The VPN-IPv4 address is 12 bytes long

Route Distinguisher 8 bytes	IPv4 Address 4 bytes
---------------------------------------	--------------------------------

Address Extensions

VPN-IPv4 address

CONTROL PLANE OPERATION ONLY



Address Extensions

Route Distinguisher

- Route Distinguishers are assigned by a Service Provider
- Route Distinguishers are globally unique (by virtue of assignment)
- Route Distinguisher is simply a number and does not contain any inherent information
- Simple VPN topologies require one The RD per customer.
- The RD could serve as a VPN identifier for simple VPN topologies, but this design could not support all topologies required by the customers.
- **Not used for constrained distribution of routing information (route filtering)**
 - Route filtering is based on Route Target and Site of Origin Communities Attributes

Route Targets

- Some sites have to participate in more than one VPN—RD cannot identify participation in more than one VPN.
- A different method is needed in which a set of identifiers can be attached to a route.
- RTs were introduced in the MPLS VPN architecture to support complex VPN topologies.

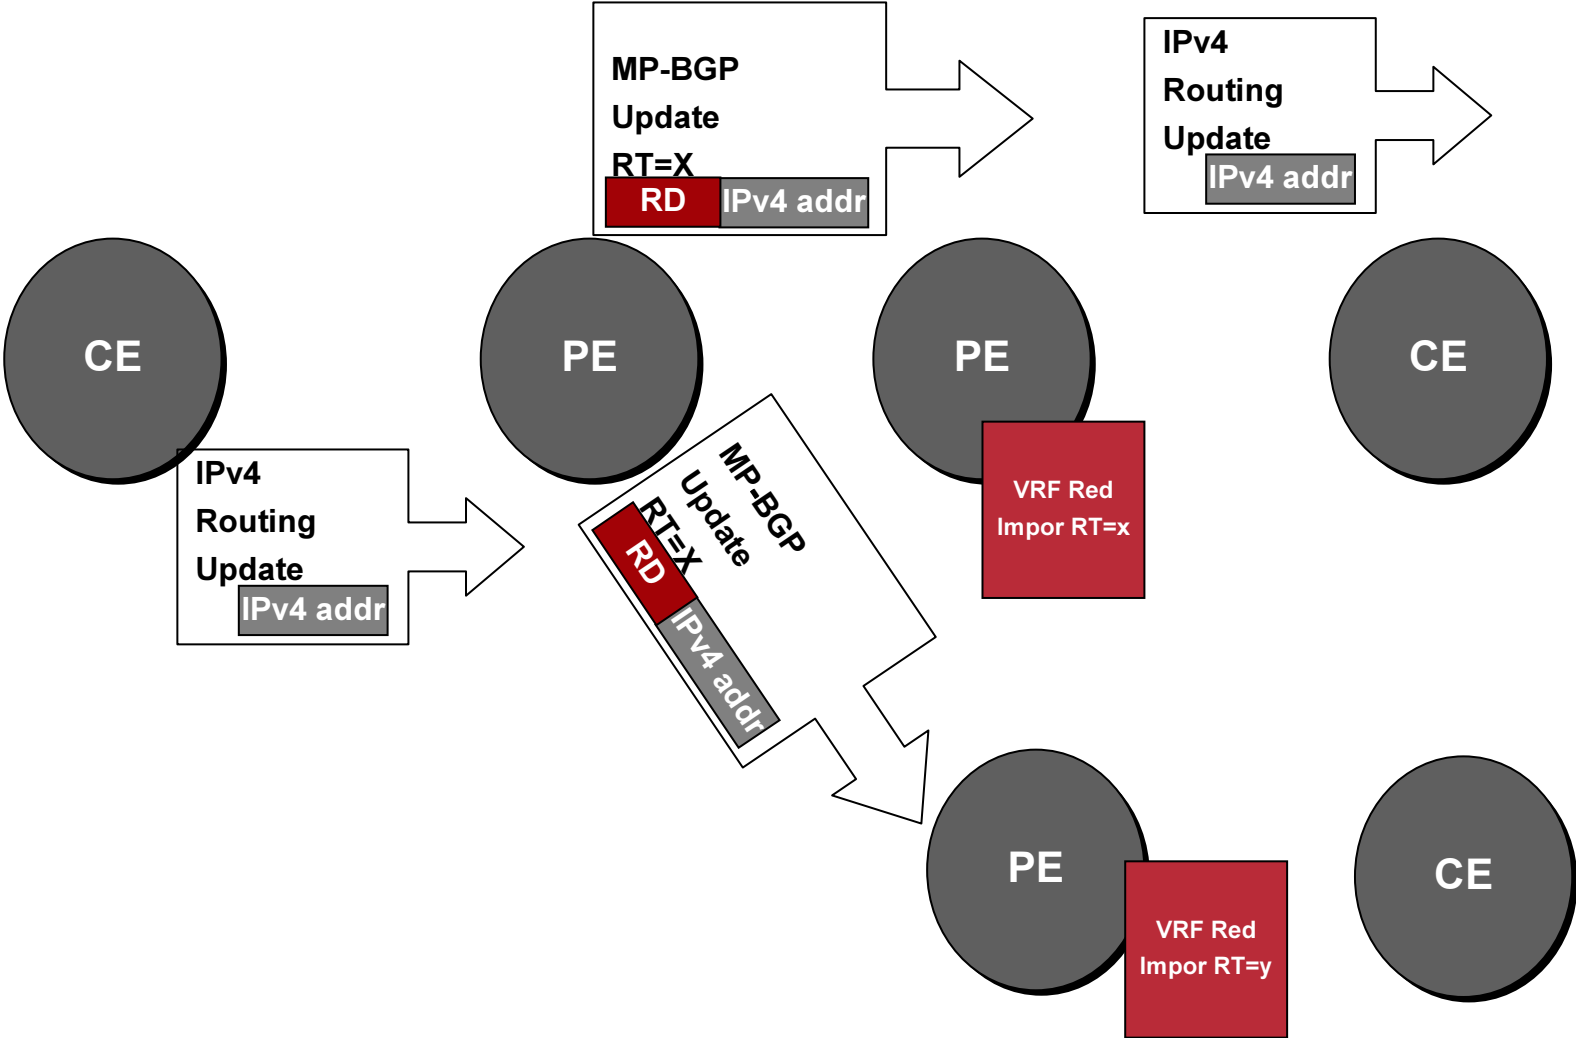
What Are Route Targets?

- Route targets (RTs) are BGP attributes of a route
- **Extended BGP communities** are used to encode these attributes.
 - Extended communities carry the meaning of the attribute together with its value.
- Every VRF is associated with one or more Route Targets
- Any number of RTs can be attached to a single route.
- Associating a particular Route Target attribute with a route allows that route to be placed in the VRFs

How Do Route Targets Work?

- **Export RTs** are attached by a PE router to a route received from site S (at the time of conversion into a VPNv4 route).
- **Import RTs** are used to determine whether a route received from another PE could be placed in the VRF associated with Site S
- Route targets usually identify VPN membership, but they can also be used in more complex scenarios.
- *VPN-IPv4 route is eligible for installation in VRF if there is some RT which is BOTH one of the route RTs and one of the VRF Import RTs*

How Do Route Targets Work? (cont.)



Route Targets to improve scalability.

- A route can only have one RD but multiple RTs.
- BGP scalability is improved with one route with multiple attributes instead of multiple routes.
- RTs could have been avoided creating more routes (thus using more RDs) but scalability would have suffered from that.

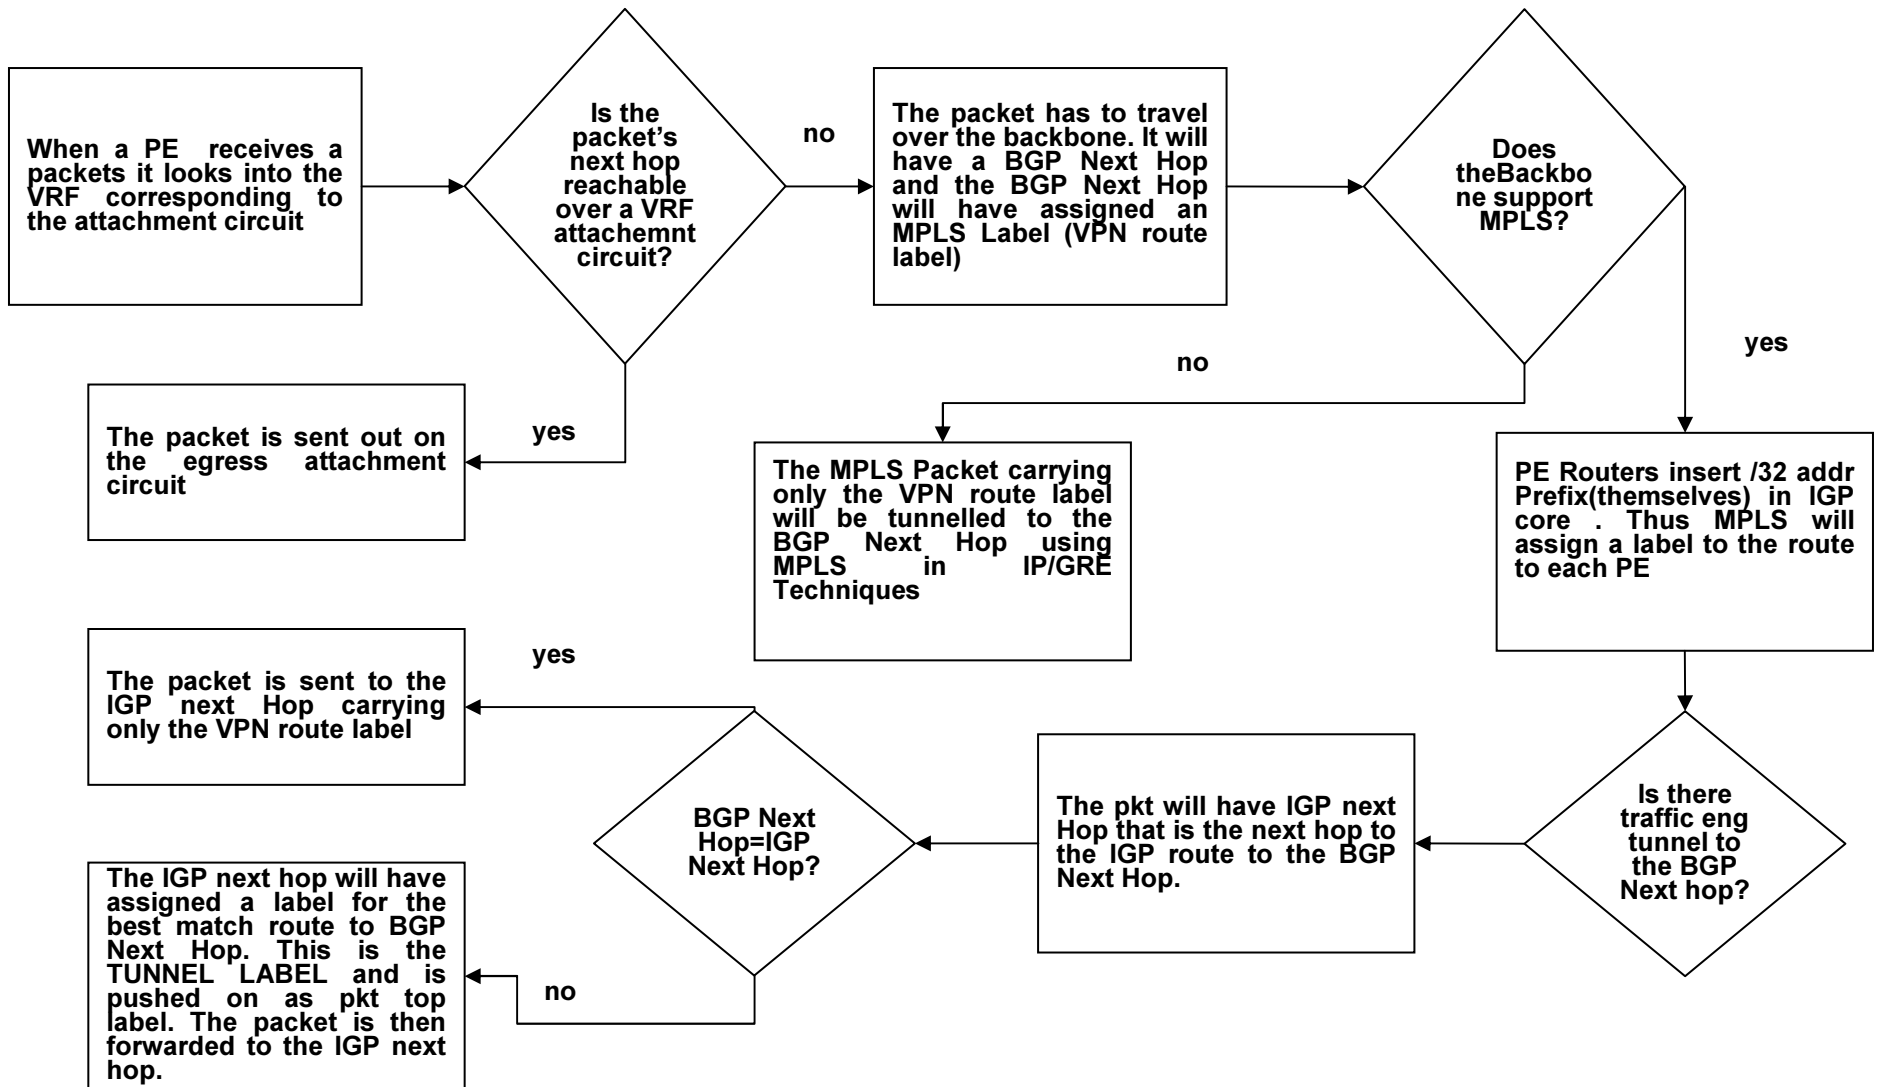
How MPLS Forwarding works for VPNs?

MPLS Forwarding

Packet forwarding

The backbone routers (P) do not have any information about VPNs -> How are packets forwarded from one side to another?

The Forwarding Flow



Two-level label stack: How it Works

Egress PE receives the packets with the label corresponding to the outgoing interface (VRF)
 One single lookup
 Label is popped and packet sent to IP neighbor

