# Мрежи за достъп от следващо поколение
# MPLS VPN

маг. инж. Николай Милованов

Нов български университет

TCMM305 Feb-Jun 2010 г.

# Objectives

Upon completion of this lesson, you will be able to perform the following tasks:

- Identify major Virtual Private network topologies, their characteristics and usage scenarios
- Describe the differences between overlay VPN and peer-to-peer VPN
- List major technologies supporting overlay VPNs and peer-to-peer VPNs
- Position MPLS VPN in comparison with other peer-to-peer VPN implementations
- Describe major architectural blocks of MPLS VPN
- Describe MPLS VPN routing model and packet forwarding

# Introduction to Virtual Private Networks

TCMM305 Feb-Jun 2010 г.

# Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the concept of VPN
- Explain VPN terminology as defined by MPLS VPN architecture

# Traditional Router-Based Networks

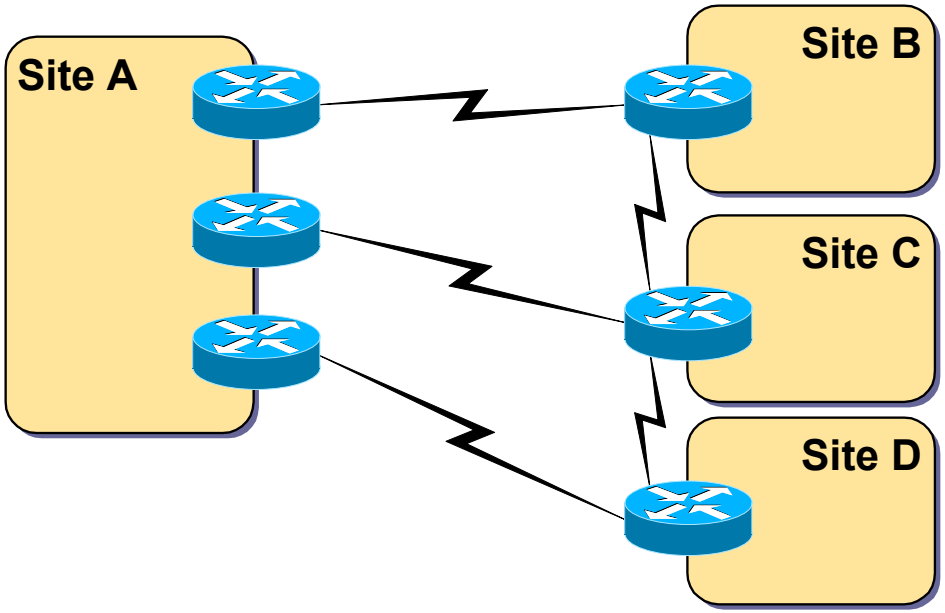


Traditional router-based networks connect customer sites through routers connected via dedicated point-to-point links.

# Virtual Private Networks

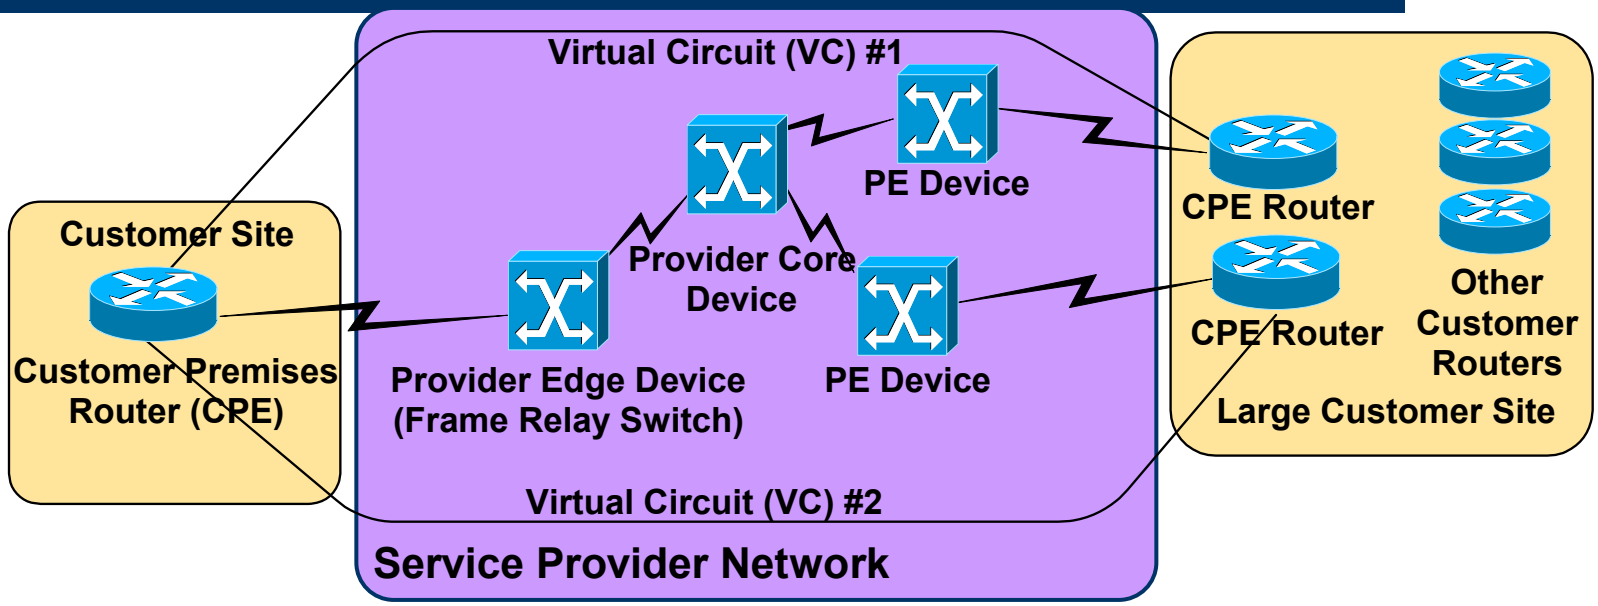


- Virtual Private Networks (VPNs) replace dedicated point-to-point links with emulated point-to-point links sharing common infrastructure.
- Customers use VPNs primarily to reduce their operational costs.

# VPN Terminology

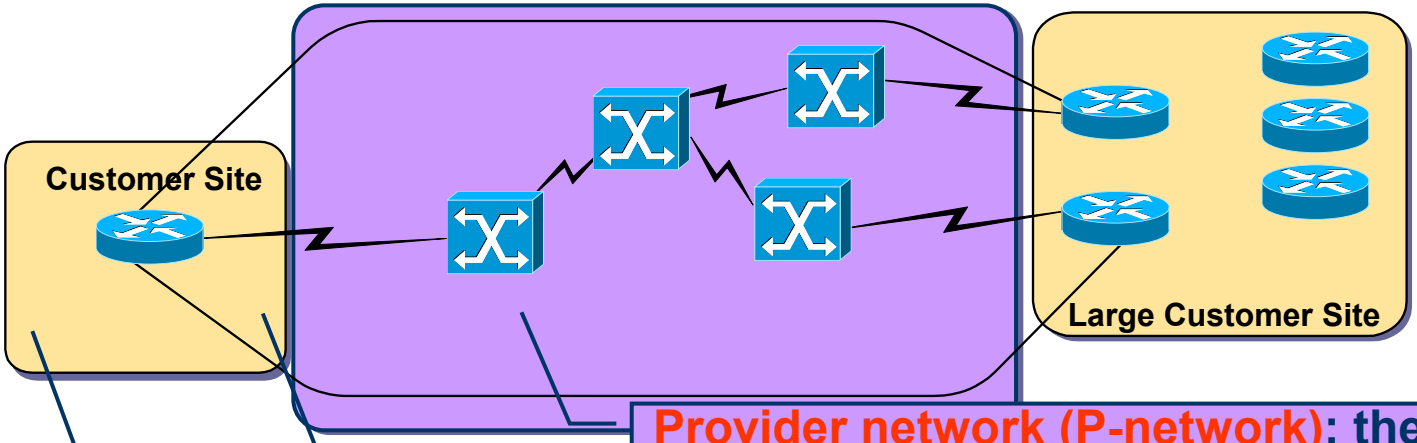


**Provider network (P-network)**: the service provider infrastructure used to provide VPN services

**Customer network (C-network)**: the part of the network still under customer control

**Customer Site**: a contiguous part of the customer network (can encompass many physical locations)

# VPN Terminology

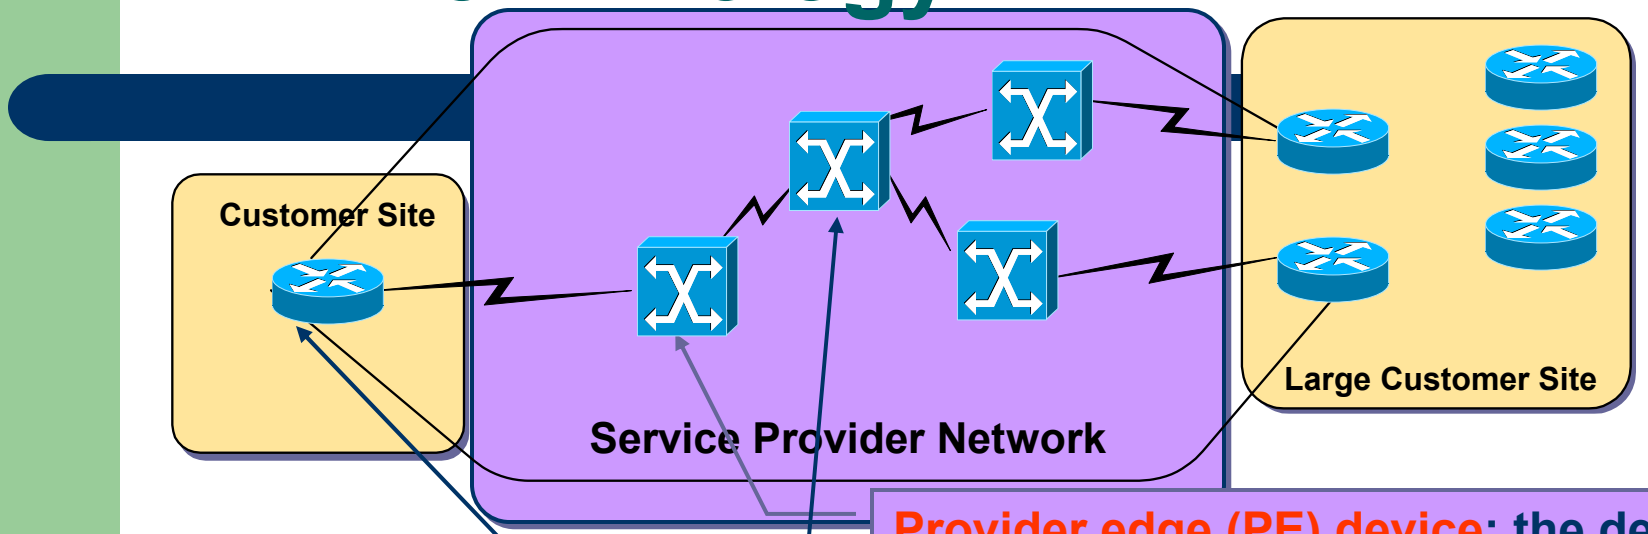


**Provider edge (PE) device**: the device in the P-network to which the CE devices are connected

**Customer edge (CE) device**: the device in the C-network that links to into P-network; also called **customer premises equipment (CPE)**

**Provider (P) device**: the device in the P-network with no customer connectivity

# VPN Implementation Technologies

- VPN services can be offered based on two major paradigms:
  - Overlay VPNs, in which the service provider provides virtual point-to-point links between customer sites
  - Peer-to-Peer VPNS, in which the service provider participates in the customer routing

# Overlay VPN Implementation (Frame Relay Example)

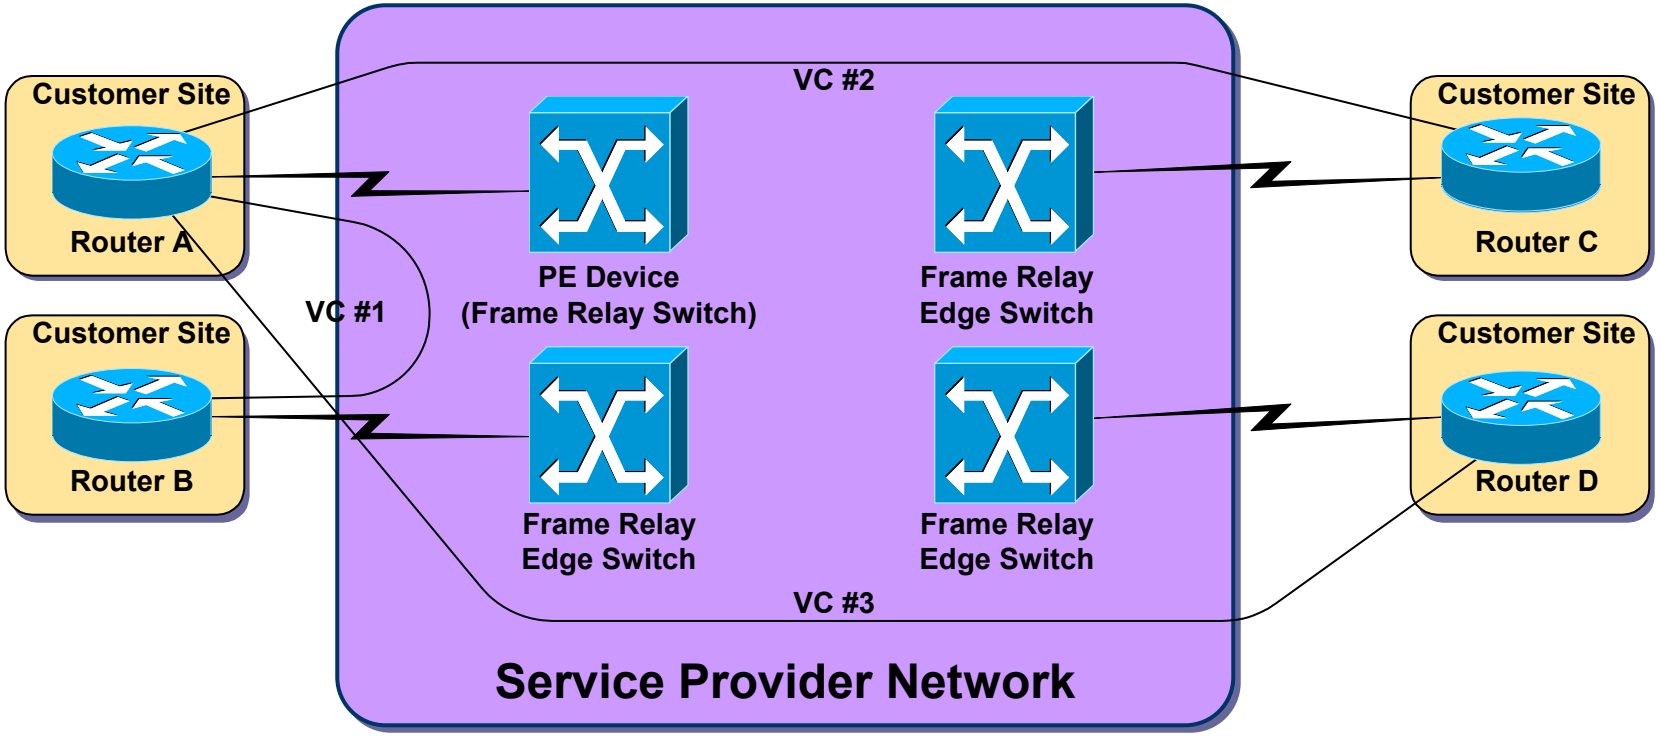

# Layer 3 Routing in Overlay VPN Implementation

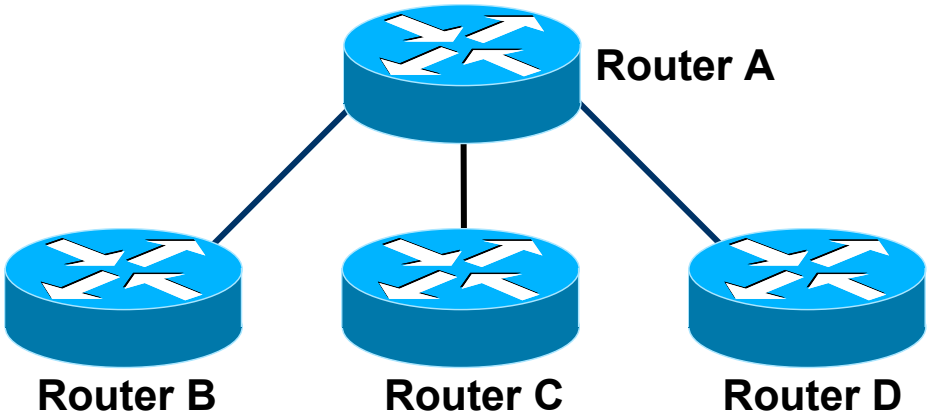


- Service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
- Service provider does not see customer routes and is responsible only for providing point-to-point transport of customer data.

# Overlay VPN
# Layer 1 Implementation

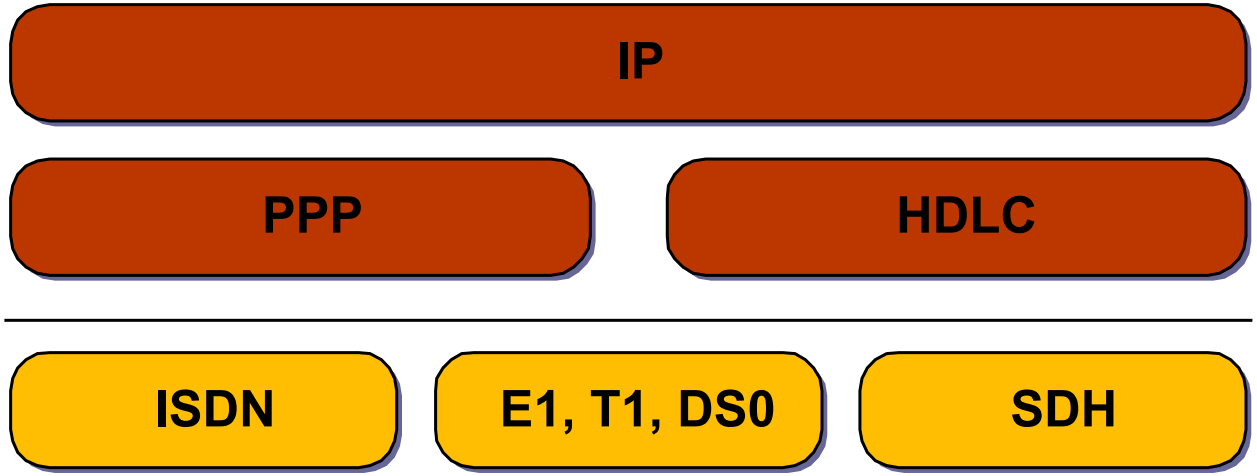


- This is the traditional TDM solution:
- Service provider establishes physical-layer connectivity between customer sites.
- Customer takes responsibility for all higher layers.

# Overlay VPN
# Layer 2 Implementation

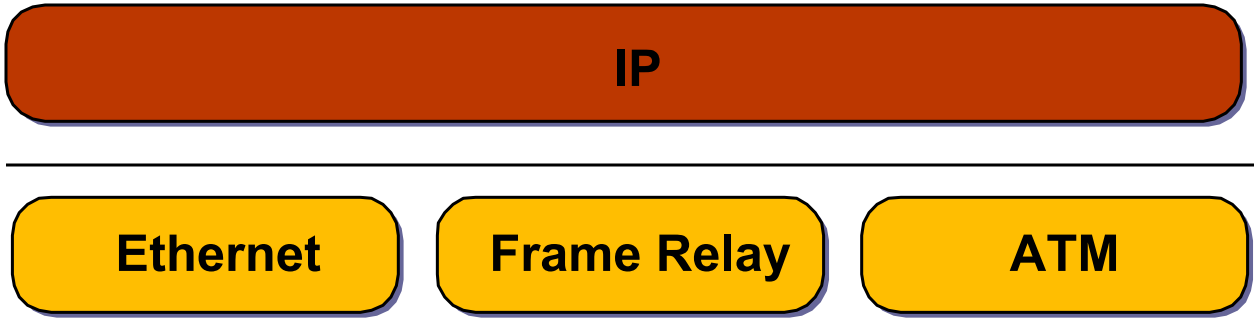


- This is the traditional switched WAN solution:
- Service provider establishes Layer 2 virtual circuits between customer sites.
- Customer takes responsibility for all higher layers.

# Overlay VPN
# IP Tunneling

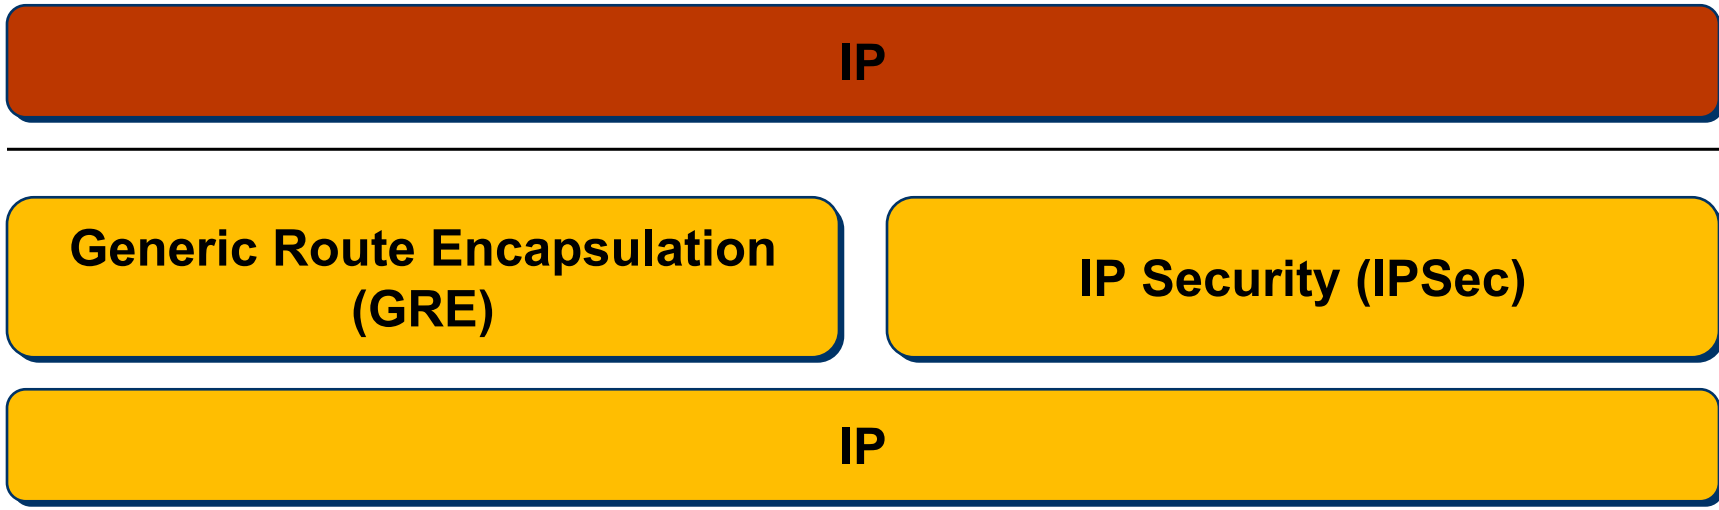


- VPN is implemented with IP-over-IP tunnels:
- Tunnels are established with GRE, IPtoIP or IPSec.
- GRE is simpler (and quicker); IPSec provides authentication and security.

# Overlay VPN
# Layer 2 Forwarding

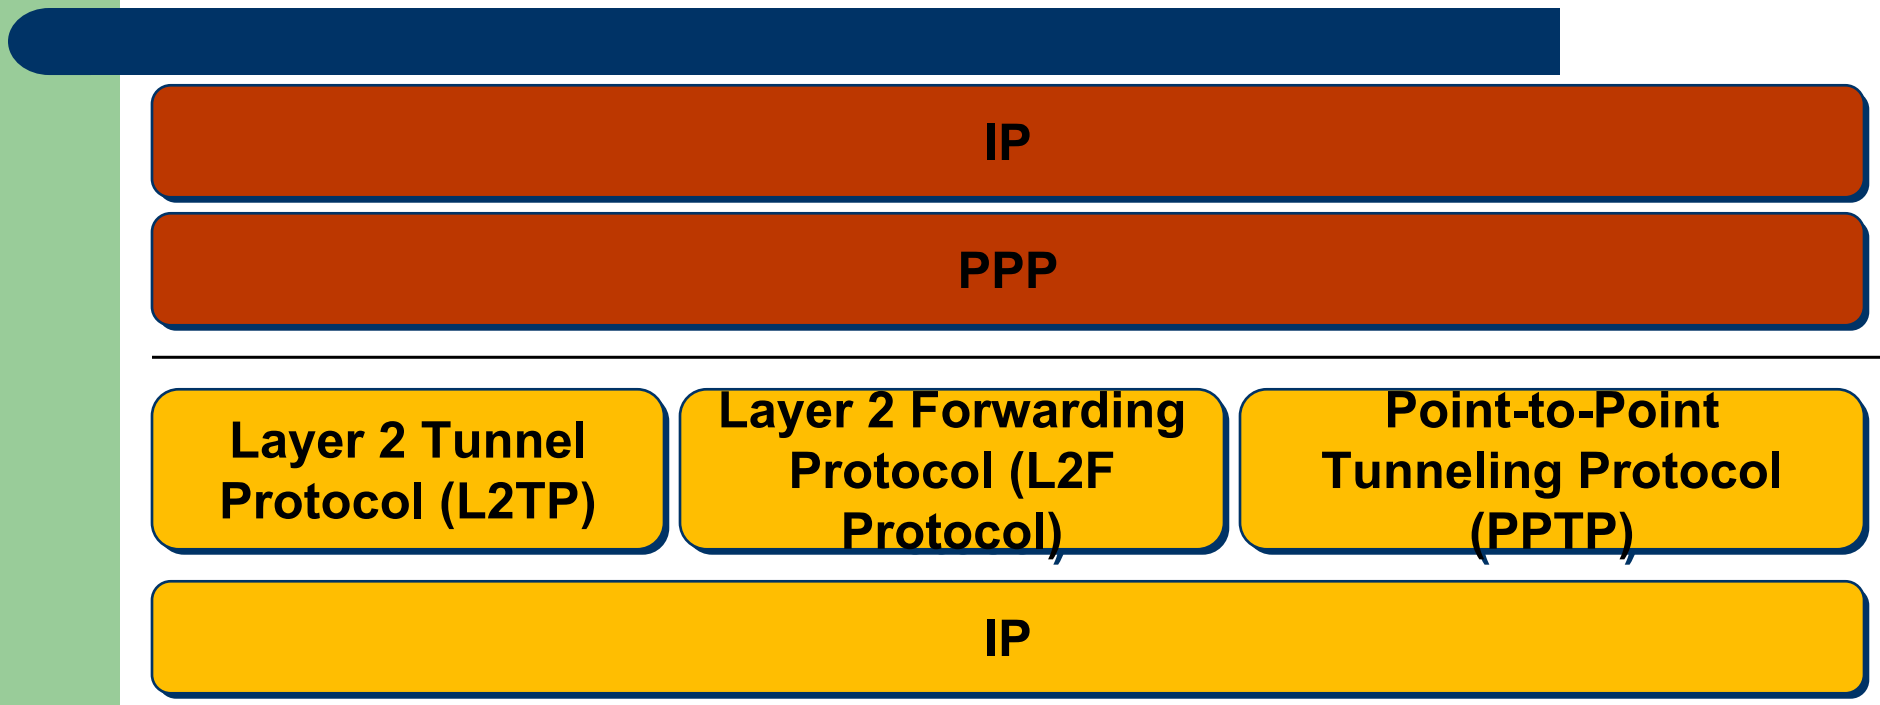


- VPN is implemented with PPP-over-IP tunnels:
- Usually used in access environments (dialup, digital subscriber line)

# Peer-to-Peer VPN Concept

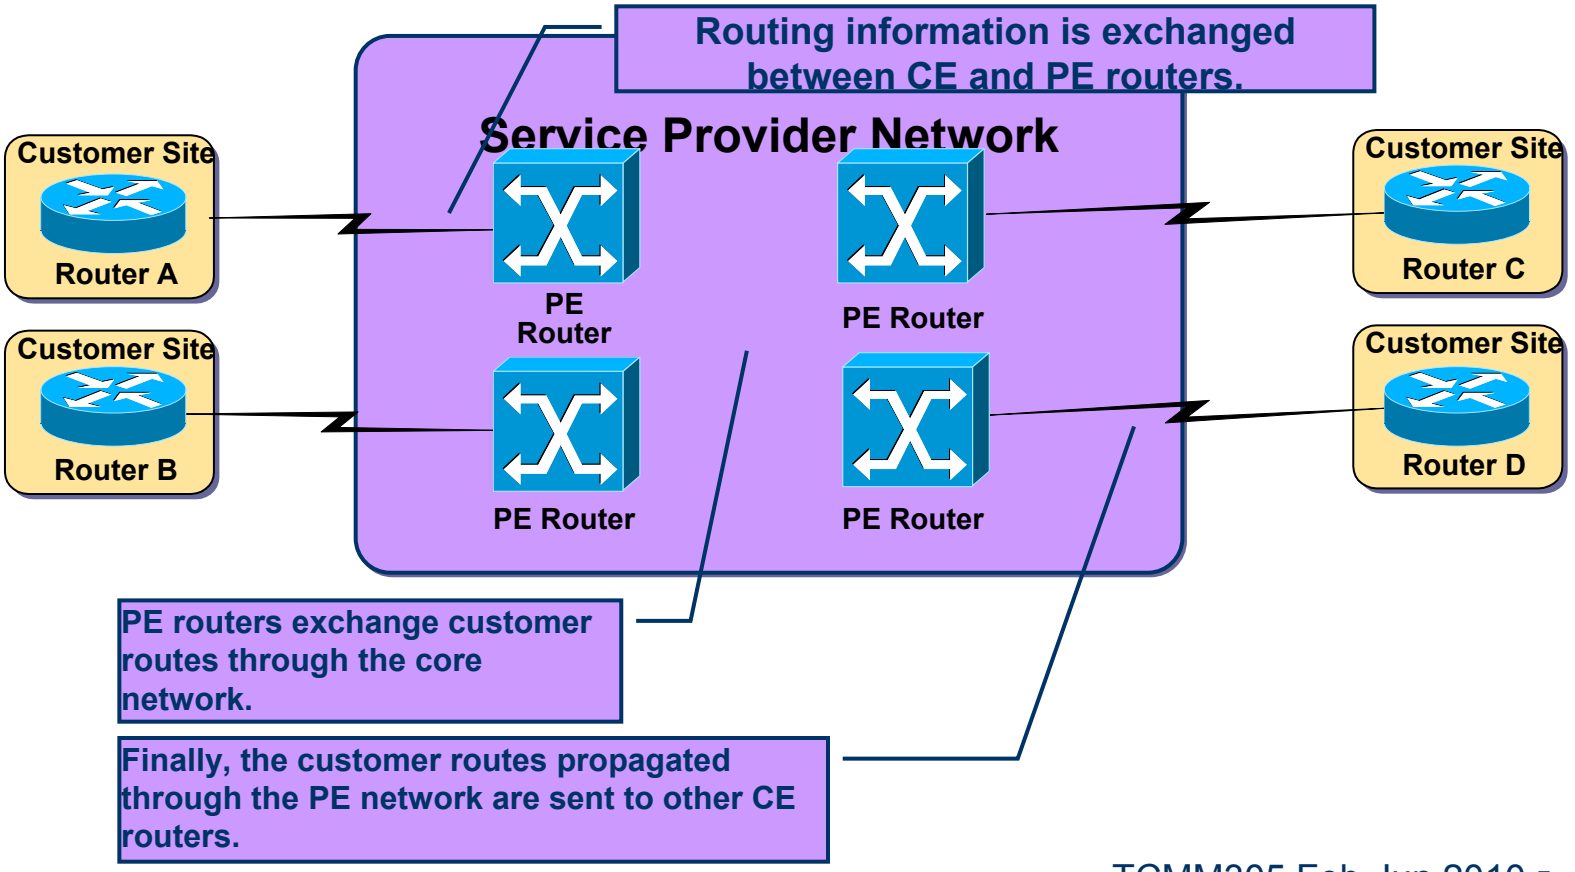

# Peer-to-Peer VPN with Packet Filters

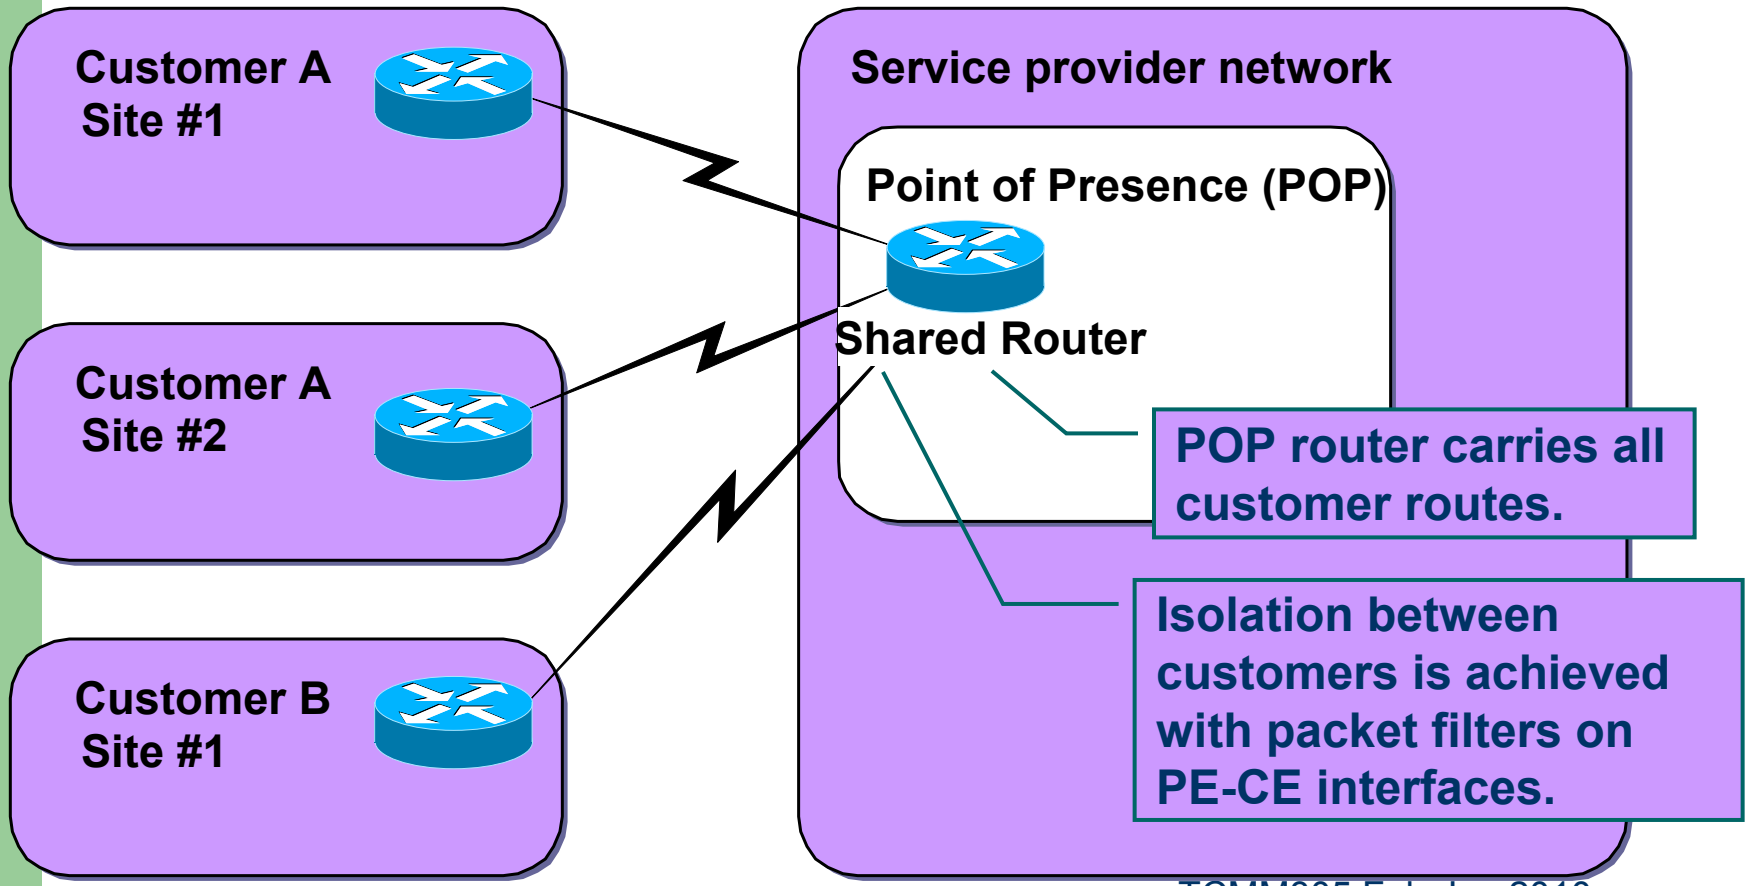

# Peer-to-Peer VPN with Controlled Route Distribution

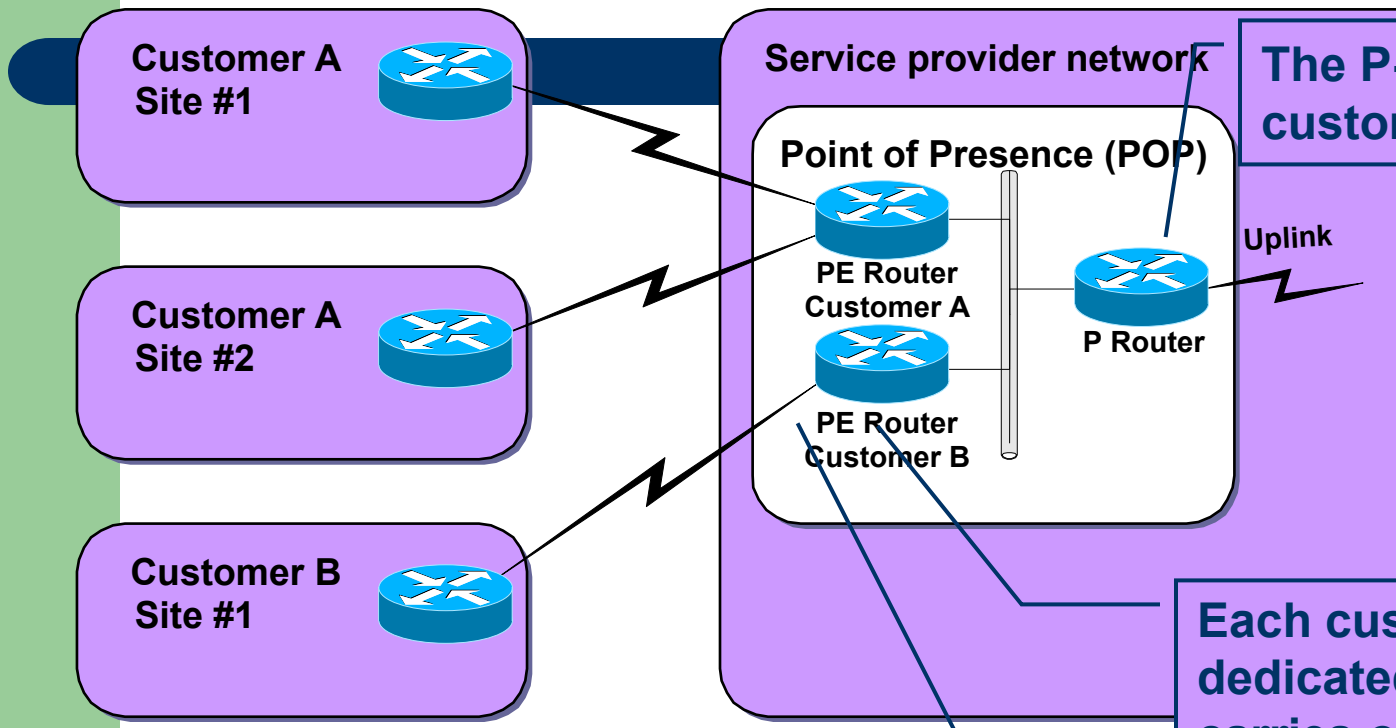


The P-router contains all customer routes.

Each customer has a dedicated PE router that carries only its routes.

Customer isolation is achieved through lack of routing information on the PE router.

# Benefits of Various VPN Implementations

- Overlay VPN:
  - Well-known and easy to implement.
  - Service provider does not participate in customer routing.
  - Customer network and service provider network are well isolated.
- Peer-to-peer VPN:
  - Guarantees optimum routing between customer sites.
  - Easier to provision an additional VPN.
  - Only the sites are provisioned, not the links between them.

# Drawbacks of Various VPN Implementations

Overlay VPN:

- Implementing optimum routing requires full mesh of virtual circuits.
- Virtual circuits have to be provisioned manually.
- Bandwidth must be provisioned on a site-to-site basis.
- Overlay VPNs always incur encapsulation overhead.

Peer-to-peer VPN:

- Service provider participates in customer routing.
- Service provider becomes responsible for customer convergence.
- PE routers carry all routes from all customers.
- Service provider needs detailed IP routing knowledge.

# Drawbacks of Traditional Peer-to-Peer VPNs

- Shared PE router:
  - All customers share the same (provider-assigned or public) address space.
  - High maintenance costs are associated with packet filters.
  - Performance is lower— each packet has to pass a packet filter.
- Dedicated PE router:
  - All customers share the same address space.
  - Each customer requires a dedicated router at each POP.

# VPN Taxonomy

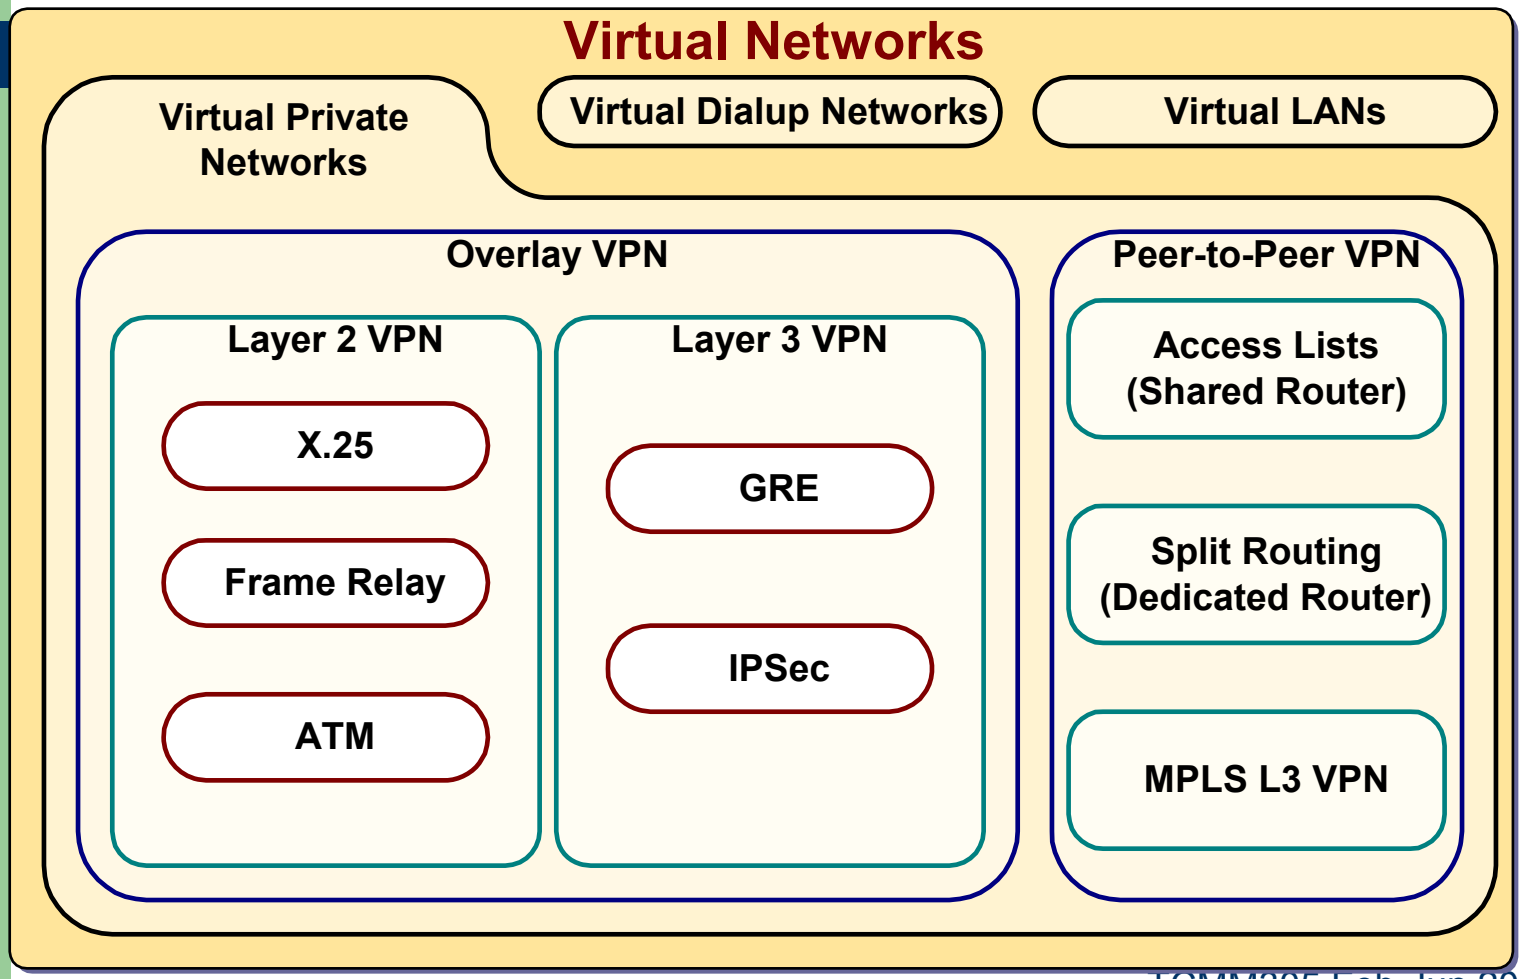

# VPN Topology Categorization

Overlay VPNs are categorized based on the topology of the virtual circuits:

- Hub and spoke
- (Redundant) Hub and spoke topology
- Partial mesh topology
- Full mesh topology
- Multilevel topology—combines several levels of overlay VPN topologies

# Overlay VPN
# Hub-and-Spoke Topology

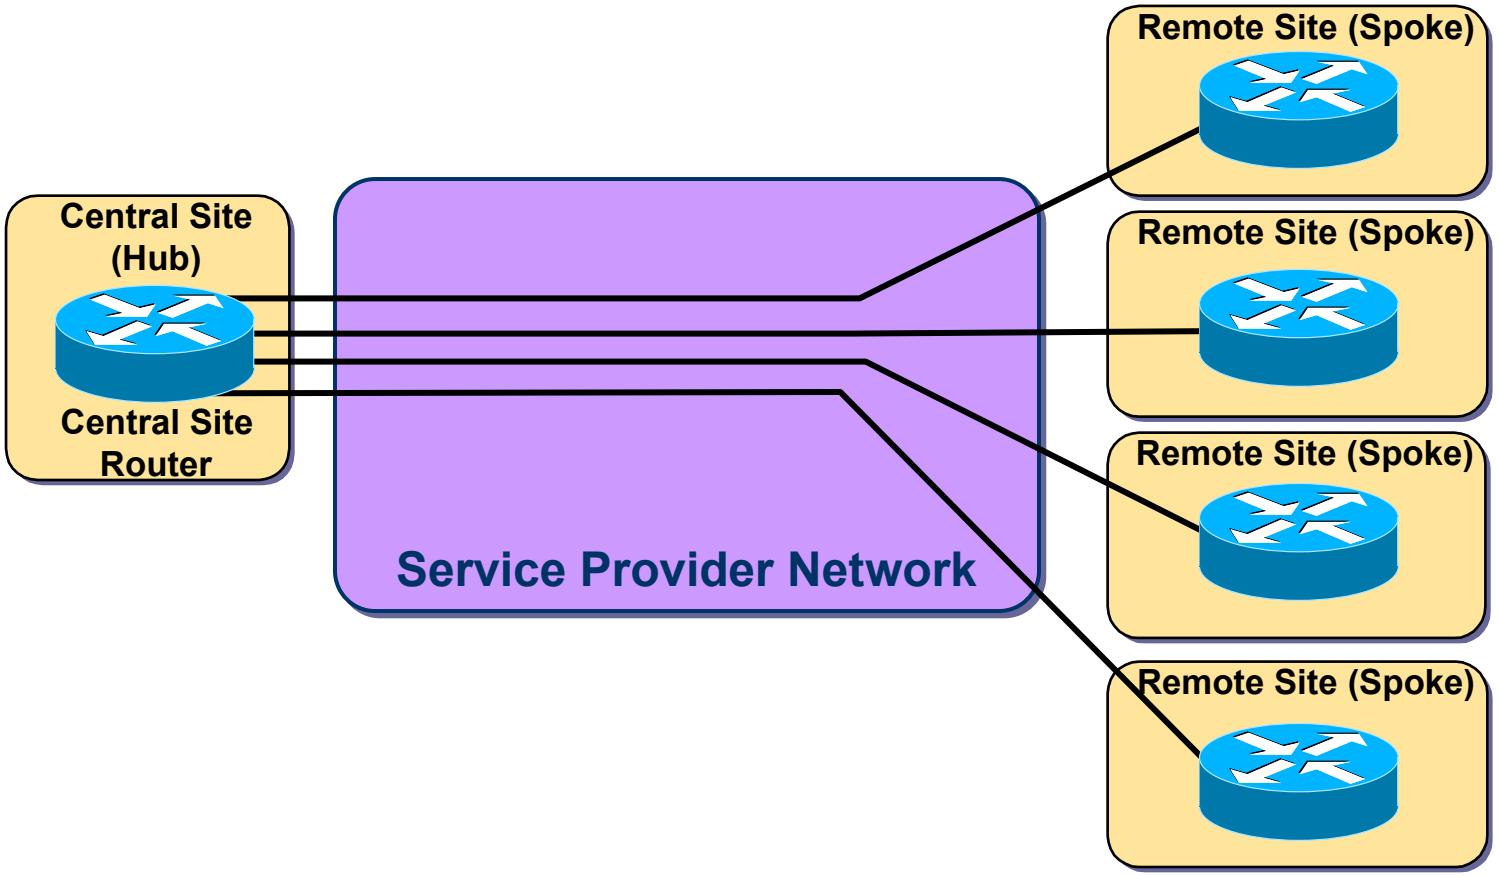

# Overlay VPN Redundant Hub and Spoke Topology

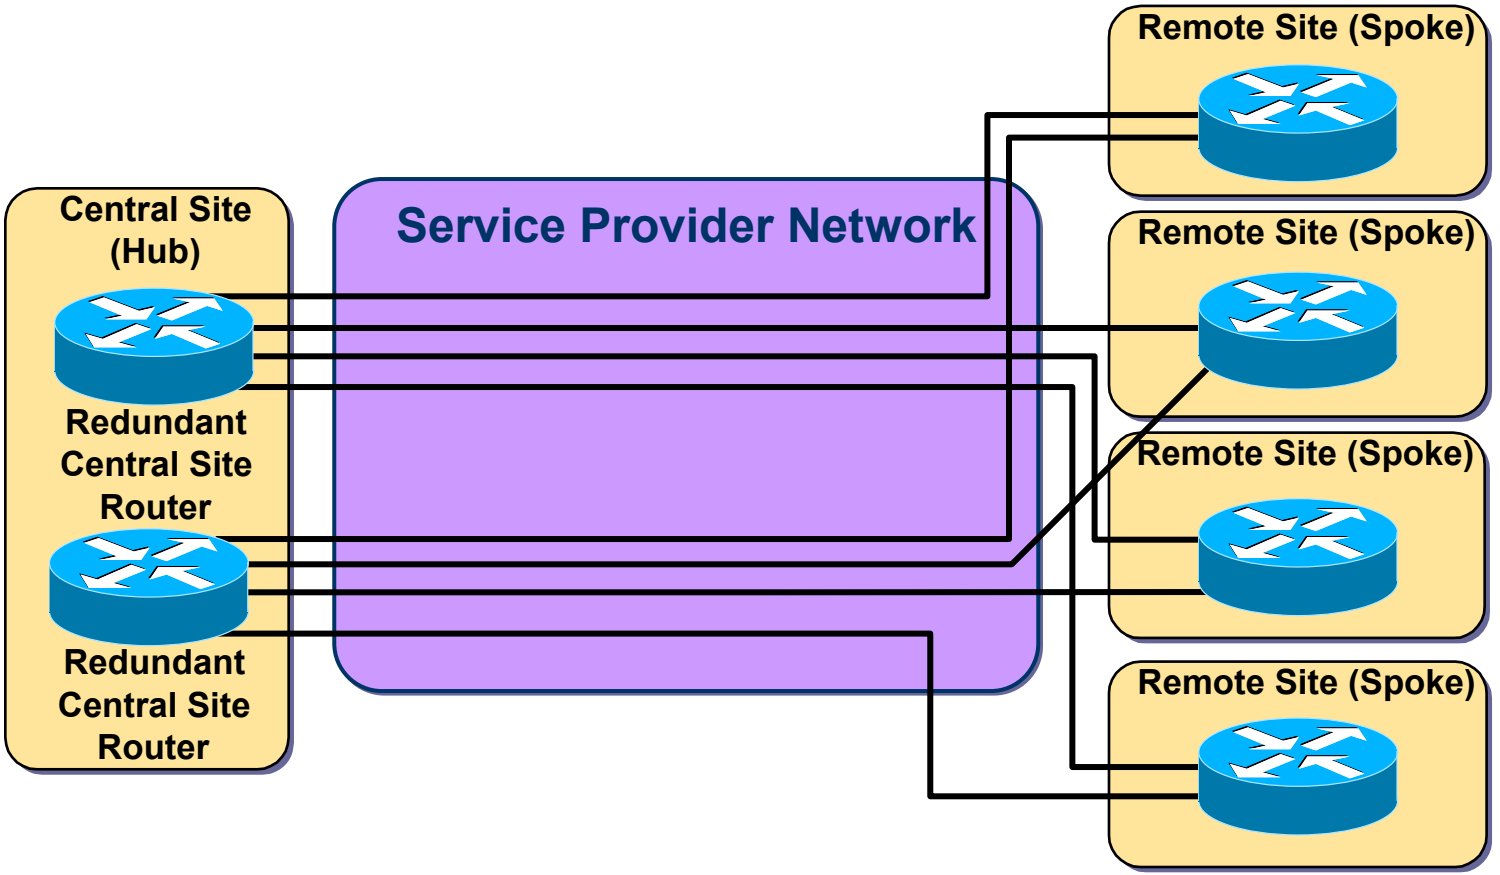

# Overlay VPN Partial Mesh

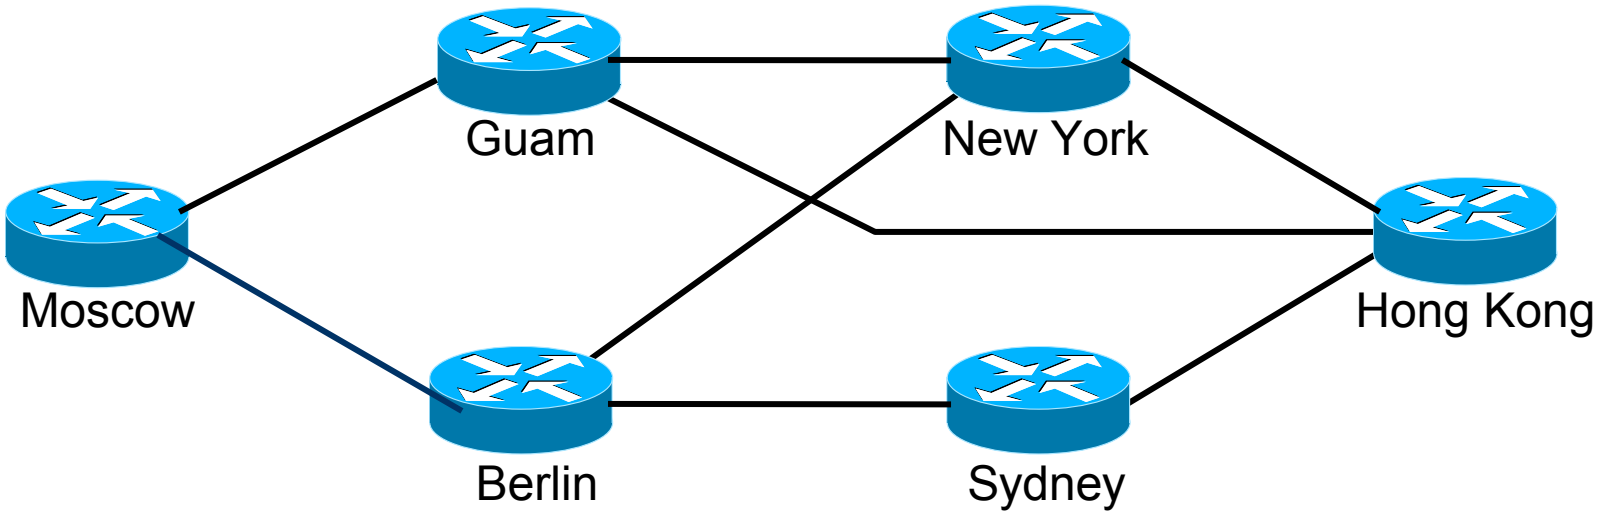


**Virtual circuits (Frame Relay Data-Link Connection Identifier)**

# Overlay VPN Multilevel Hub-and-Spoke

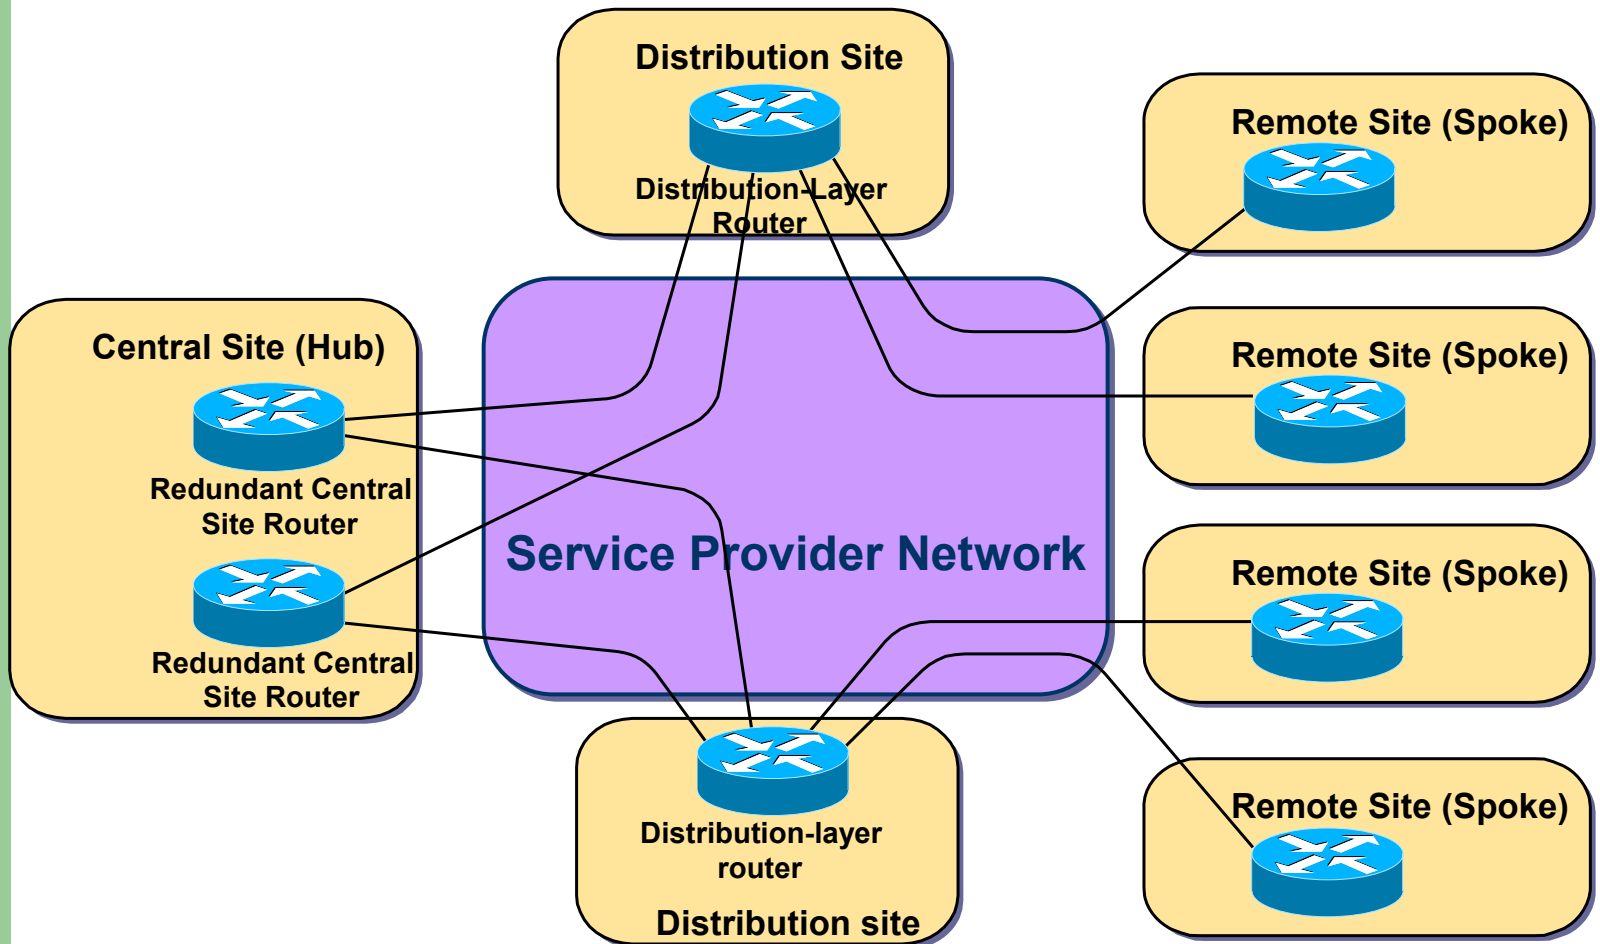

# VPN Business Categorization

VPNs can be categorized on the business needs they fulfill:

- Intranet VPN—connects sites within an organization.
- Extranet VPN—connects different organizations in a secure way.
- Access VPN — VPDN provides dialup access into a customer network.

# Extranet VPN—Overlay VPN Implementation

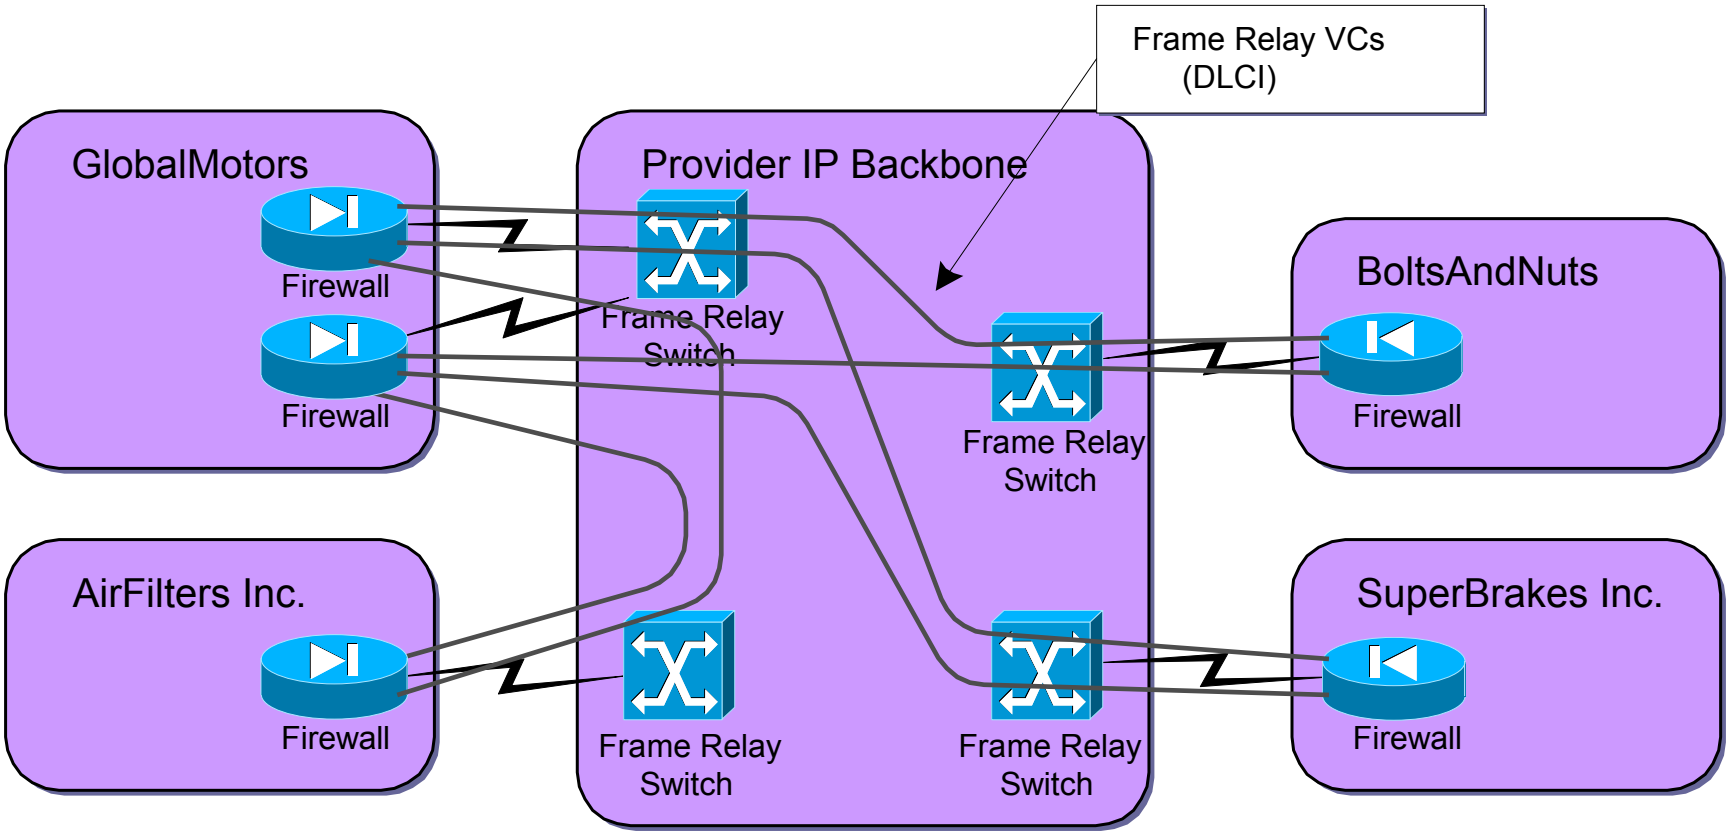

# Extranet VPN—Peer-to-Peer VPN Implementation

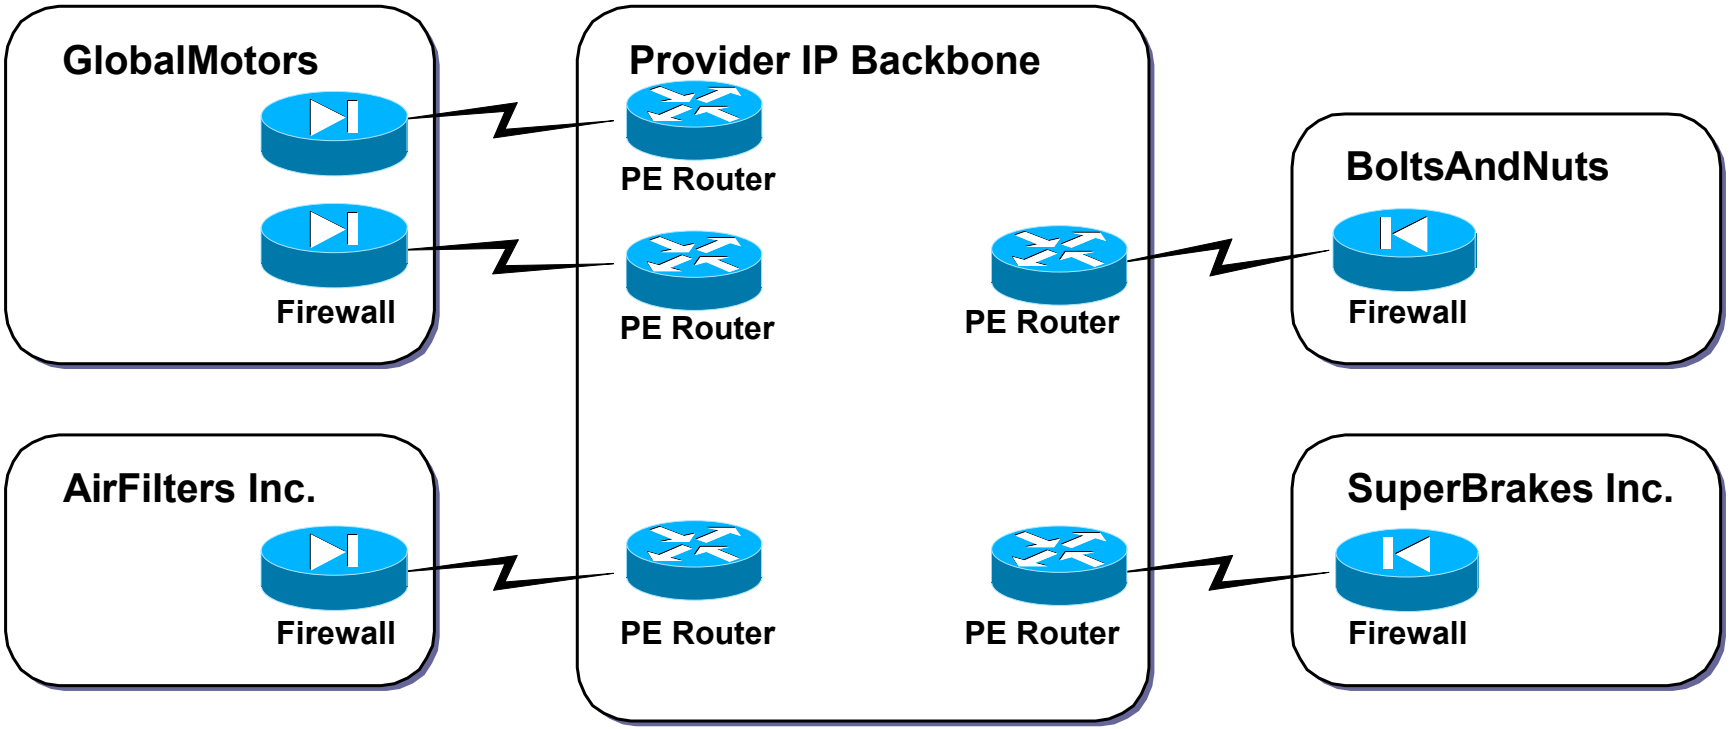

# VPN Connectivity Categorization

VPNs can also be categorized by the connectivity required between sites:

- Simple VPN—every site can communicate with every other site.
- Overlapping VPN—some sites participate in more than one simple VPN.
- Central Services VPN—all sites can communicate with central servers, but not with each other.
- Managed Network—a dedicated VPN is established to manage CE routers.

# Central Services Extranet

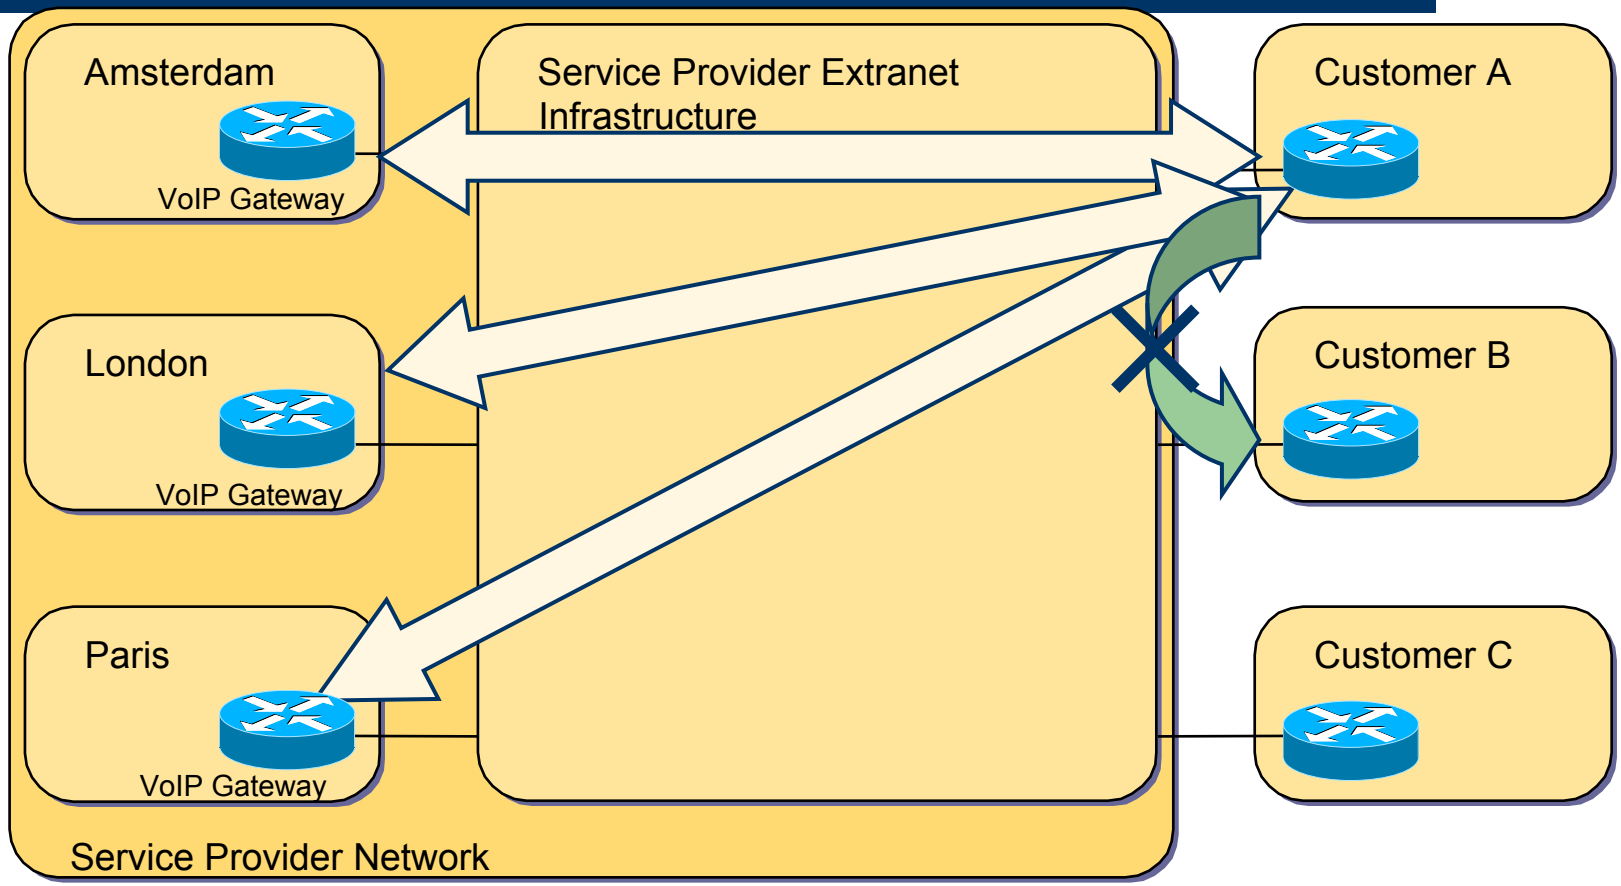

# Central Services Extranet—Hybrid (Overlay + Peer-to-Peer) Implementation

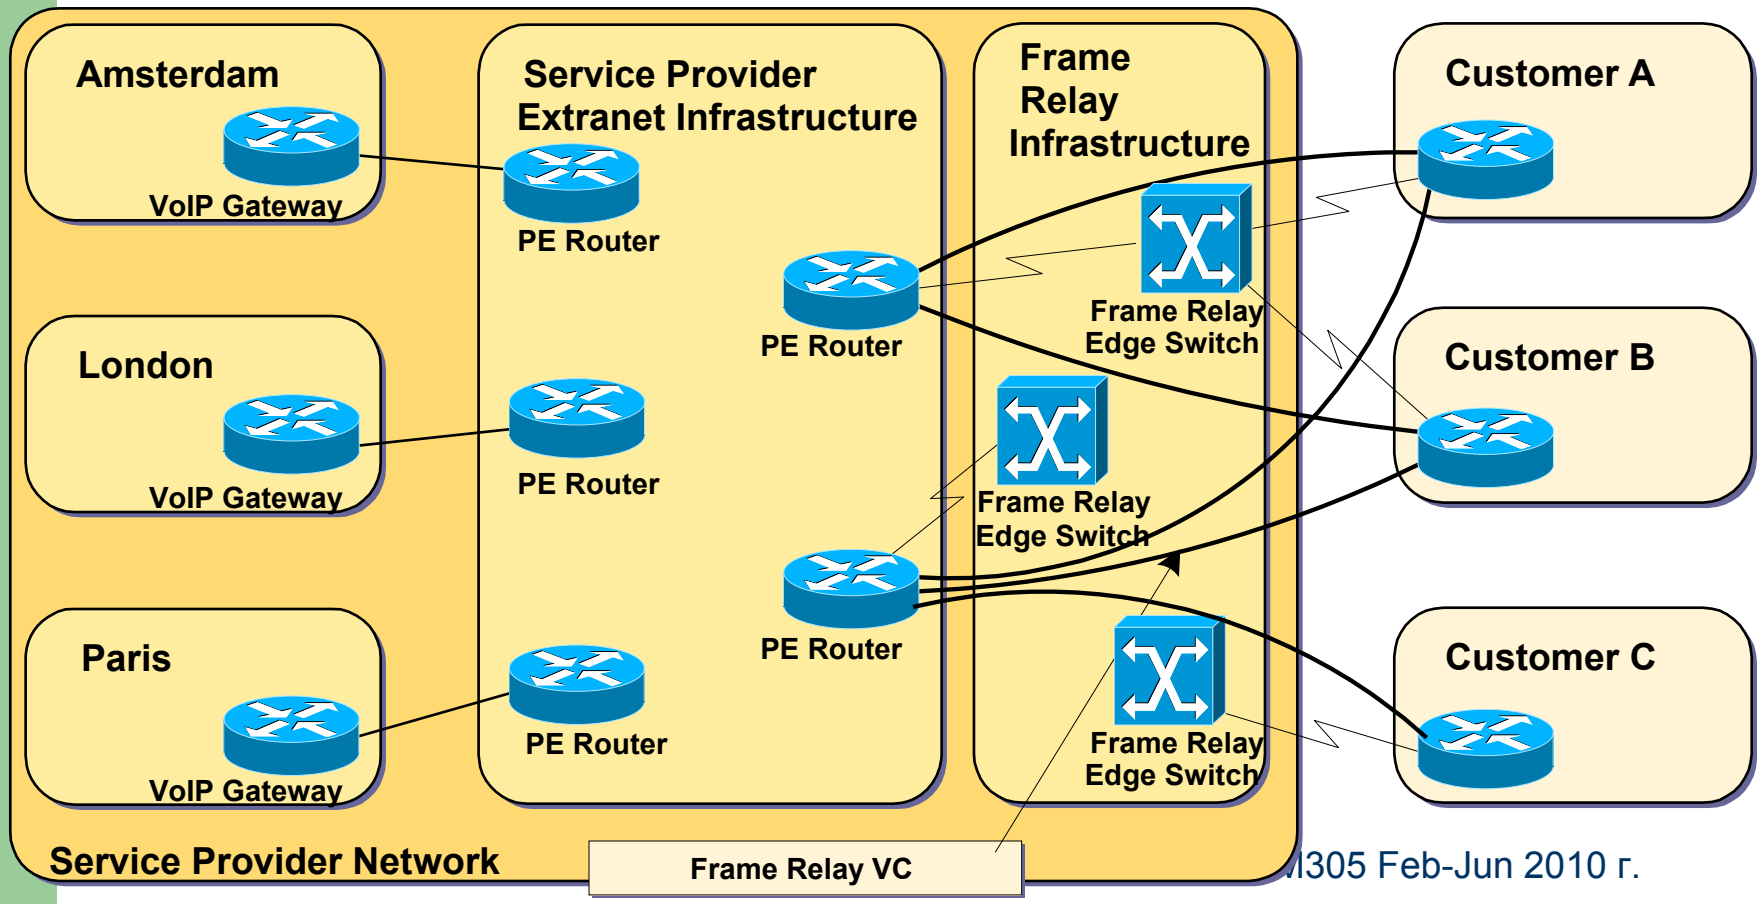

# Managed Network Overlay VPN Implementation

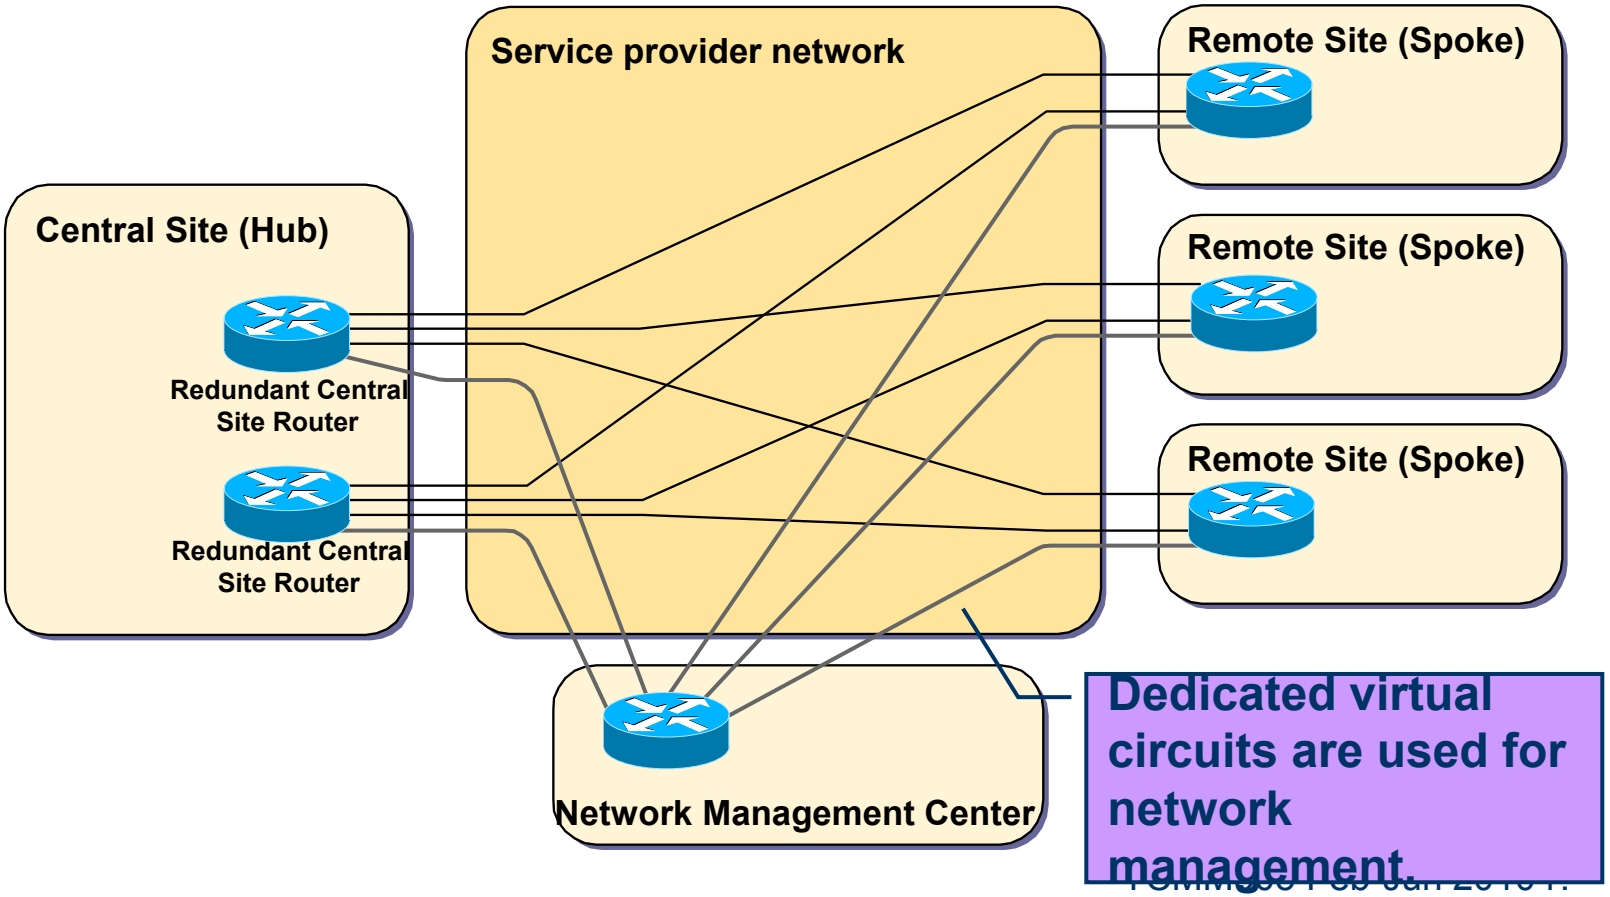

# Summary

After completing this section, you should be able to perform the following tasks:

- Identify major VPN topologies
- Describe the implications of using overlay VPN or peer-to-peer VPN approach with each topology
- List sample usage scenarios for each topology

# MPLS VPN Architecture

TCMM305 Feb-Jun 2010 г.

# Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the difference between traditional peer-to-peer models and MPLS VPN
- List the benefits of MPLS VPN
- Describe major architectural blocks of MPLS VPN
- Explain the need for route distinguisher and route target

# MPLS VPN Architecture

MPLS VPN combines the best features of overlay VPN and peer-to-peer VPN:

- PE routers participate in customer routing, guaranteeing optimum routing between sites and easy provisioning.
- PE routers carry a separate set of routes for each customer (similar to the dedicated PE router approach).
- Customers can use overlapping addresses.

# MPLS VPN Terminology

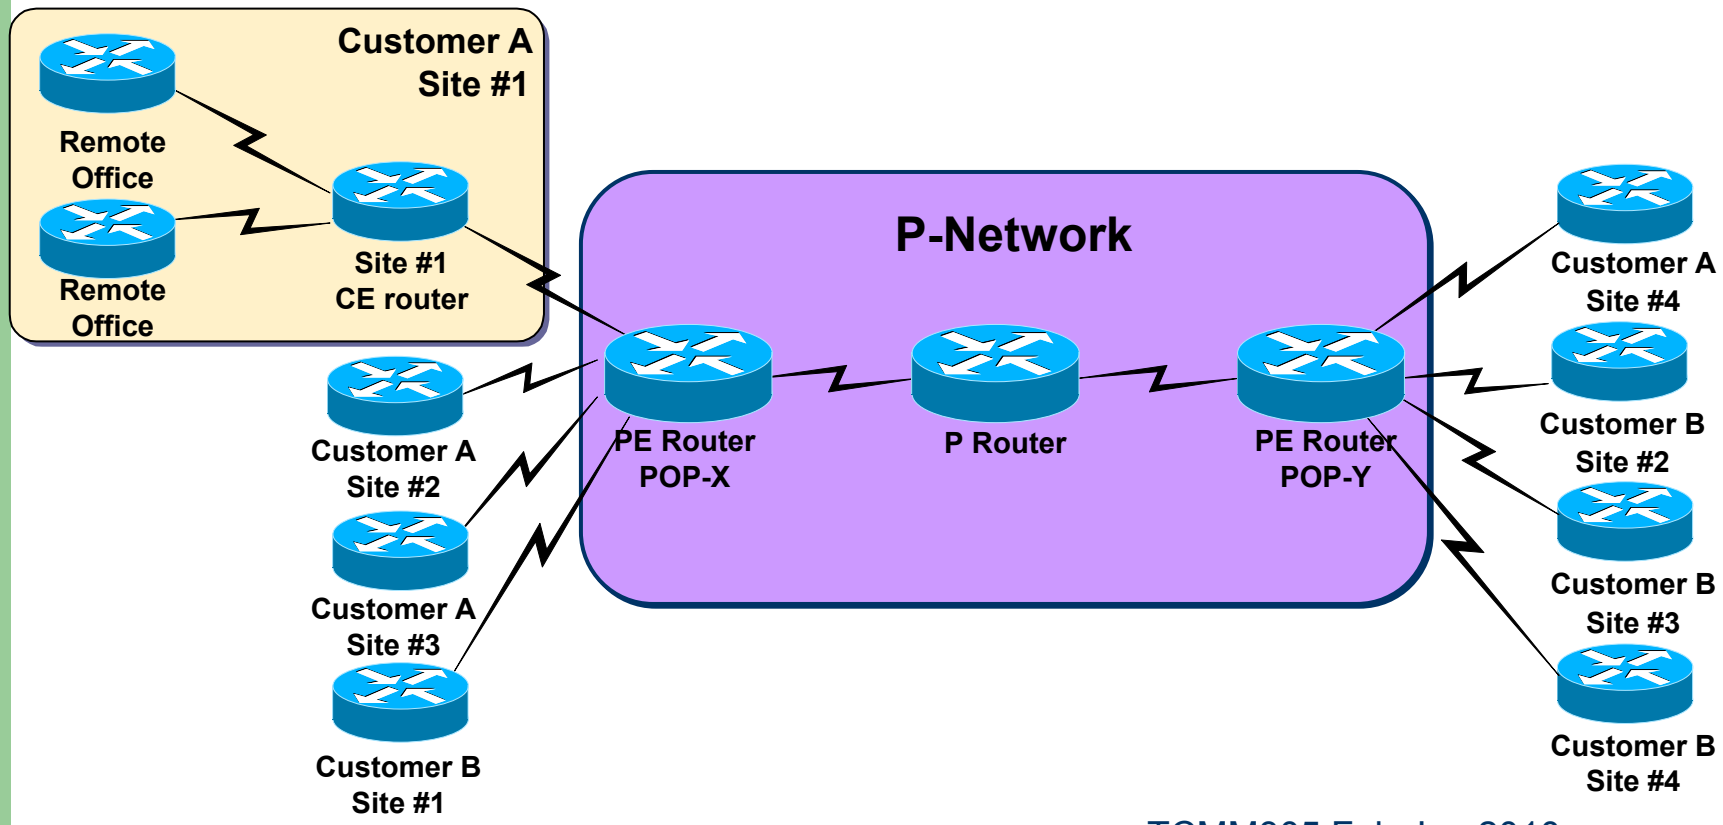

# PE Router Architecture

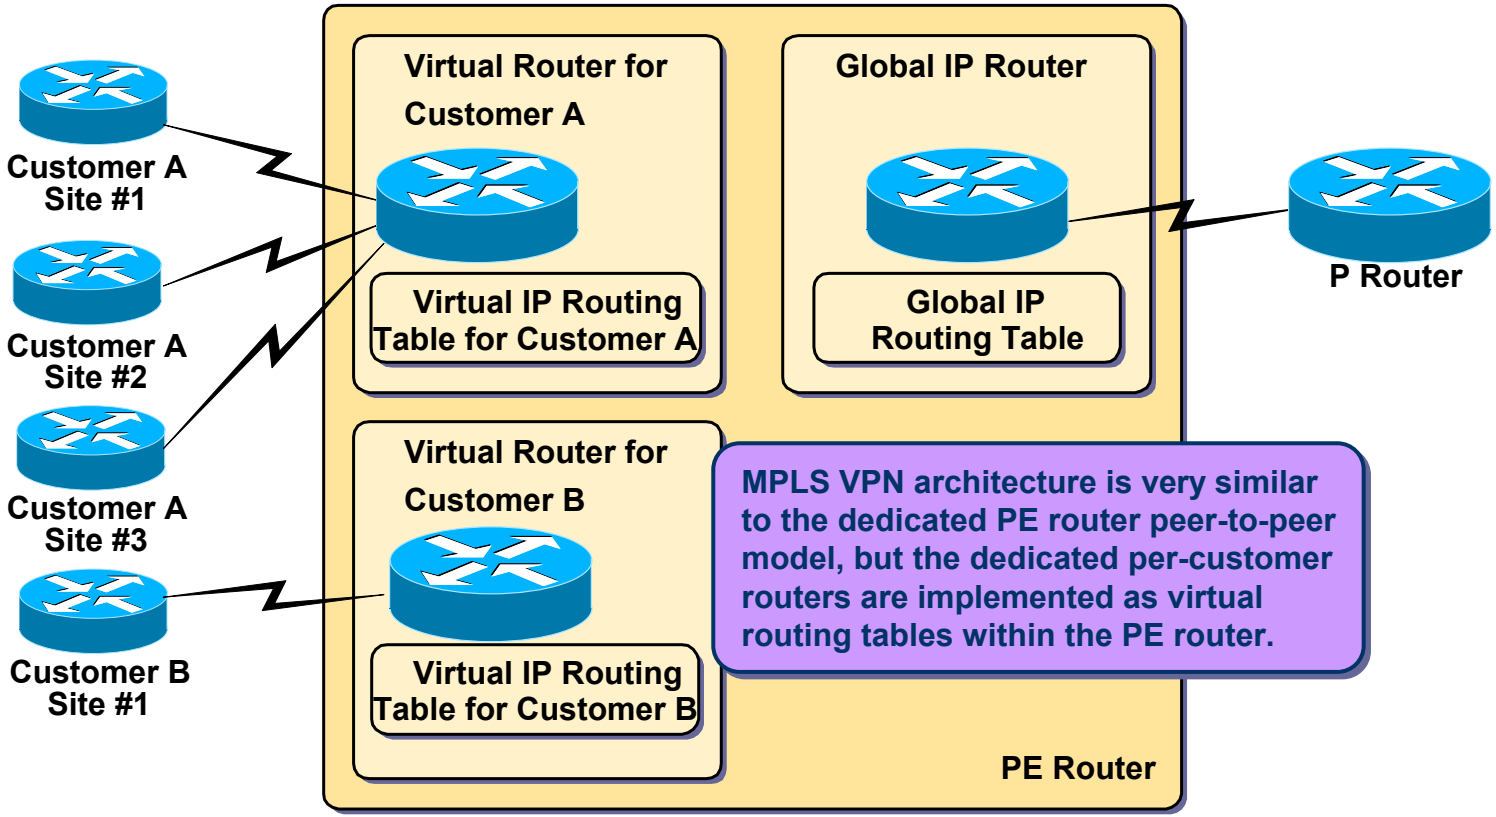

# Routing Information Propagation Across the P-Network

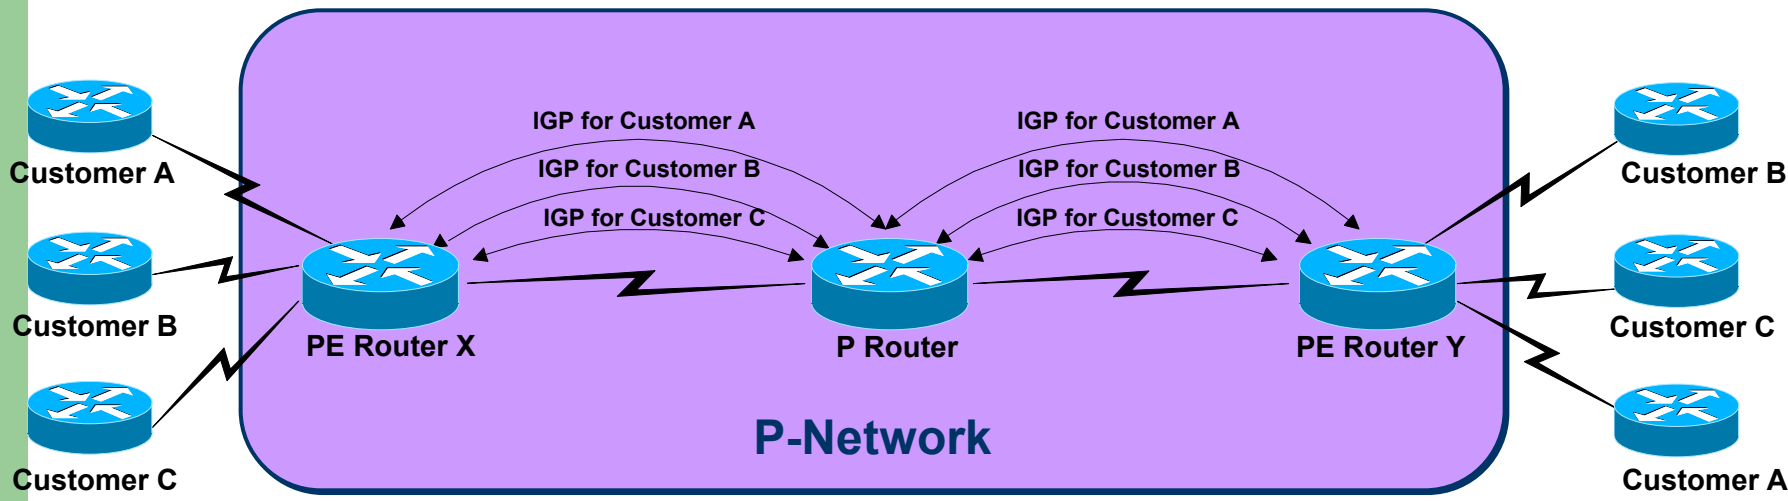


**Q:** How will PE routers exchange customer routing information?

**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.)

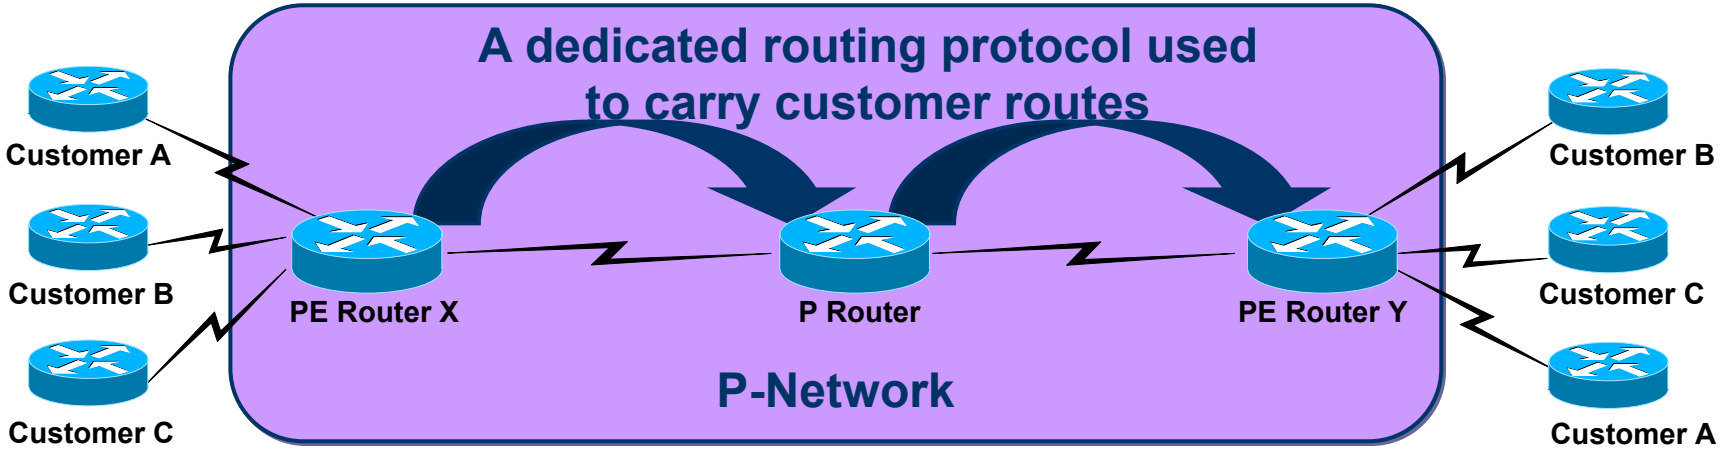


**Q:** How will PE routers exchange customer routing information?

**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 routers.

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

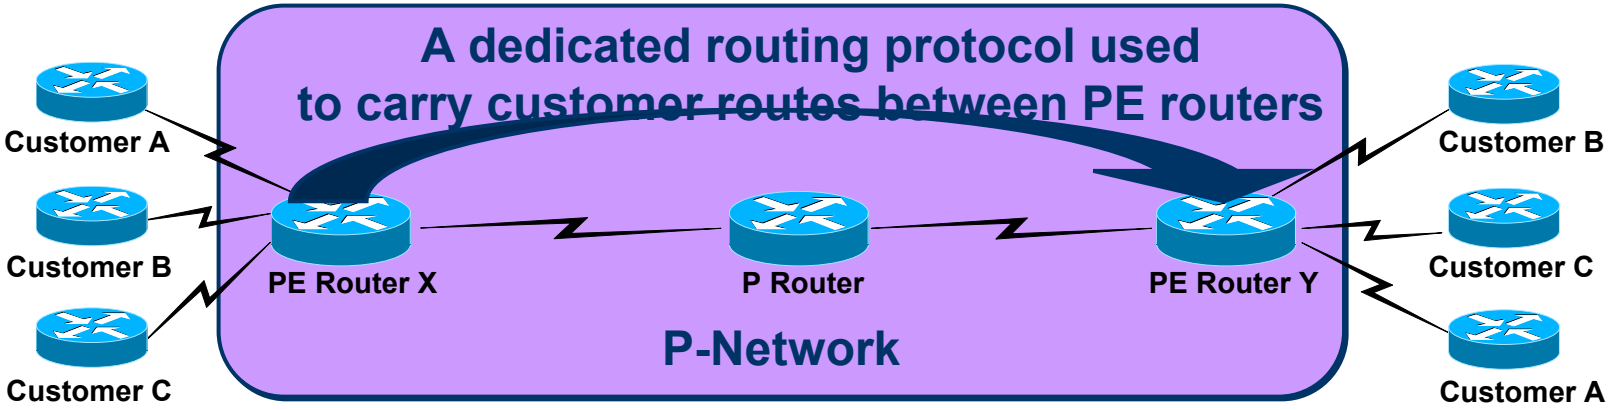


**Q: How will PE routers exchange customer routing information?**

**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.)

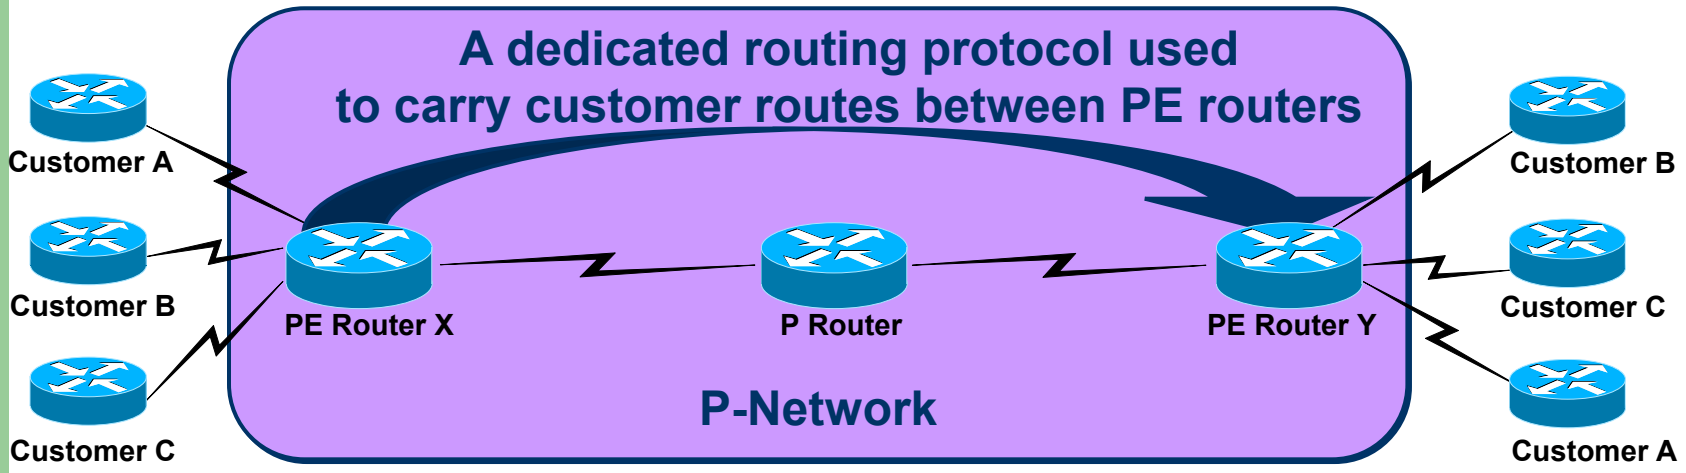


**Q:** Which protocol can be used to carry customer routes between PE route

**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.)

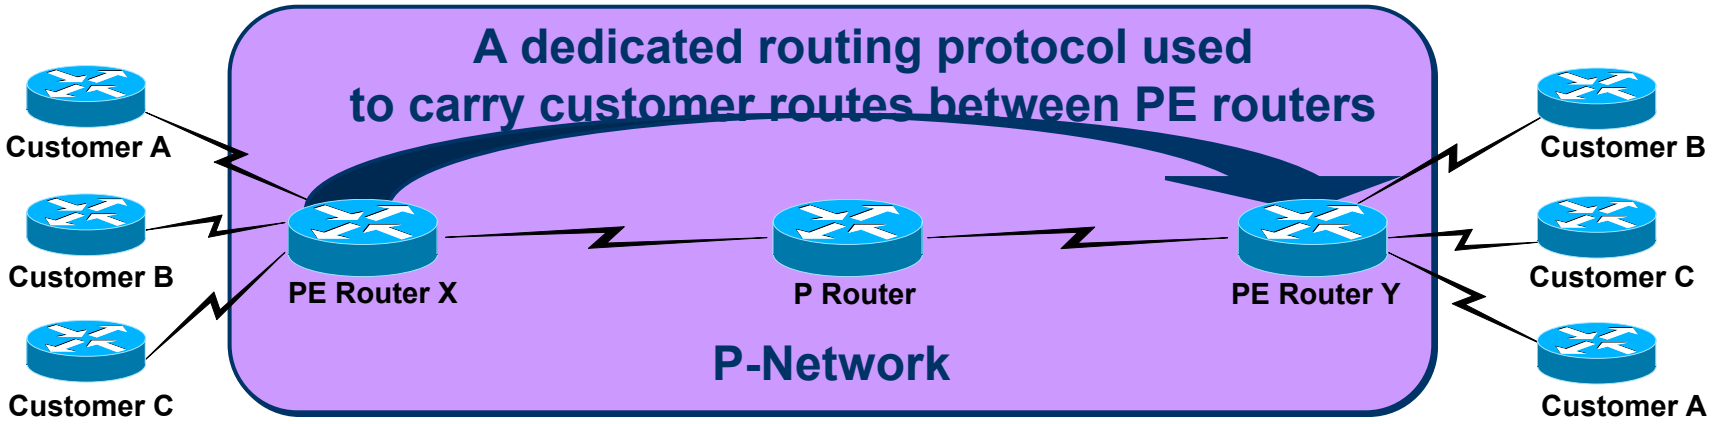


**Q:** Customers can have overlapping address spaces. How will information about the same subnet of two customers be propagated via a single routing protocol?

**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.

# Route Distinguisher

- The RD is a 64-bit quantity prepended to an IP version 4 (IPv4) address to make it globally unique.
- The resulting 96-bit address is called a VPNv4 address.
- VPNv4 addresses are exchanged only via BGP between PE routers.
  - BGP that supports address families other than IPv4 addresses is called multiprotocol BGP (MP-BGP).

# Route Distinguisher Usage in an MPLS VPN

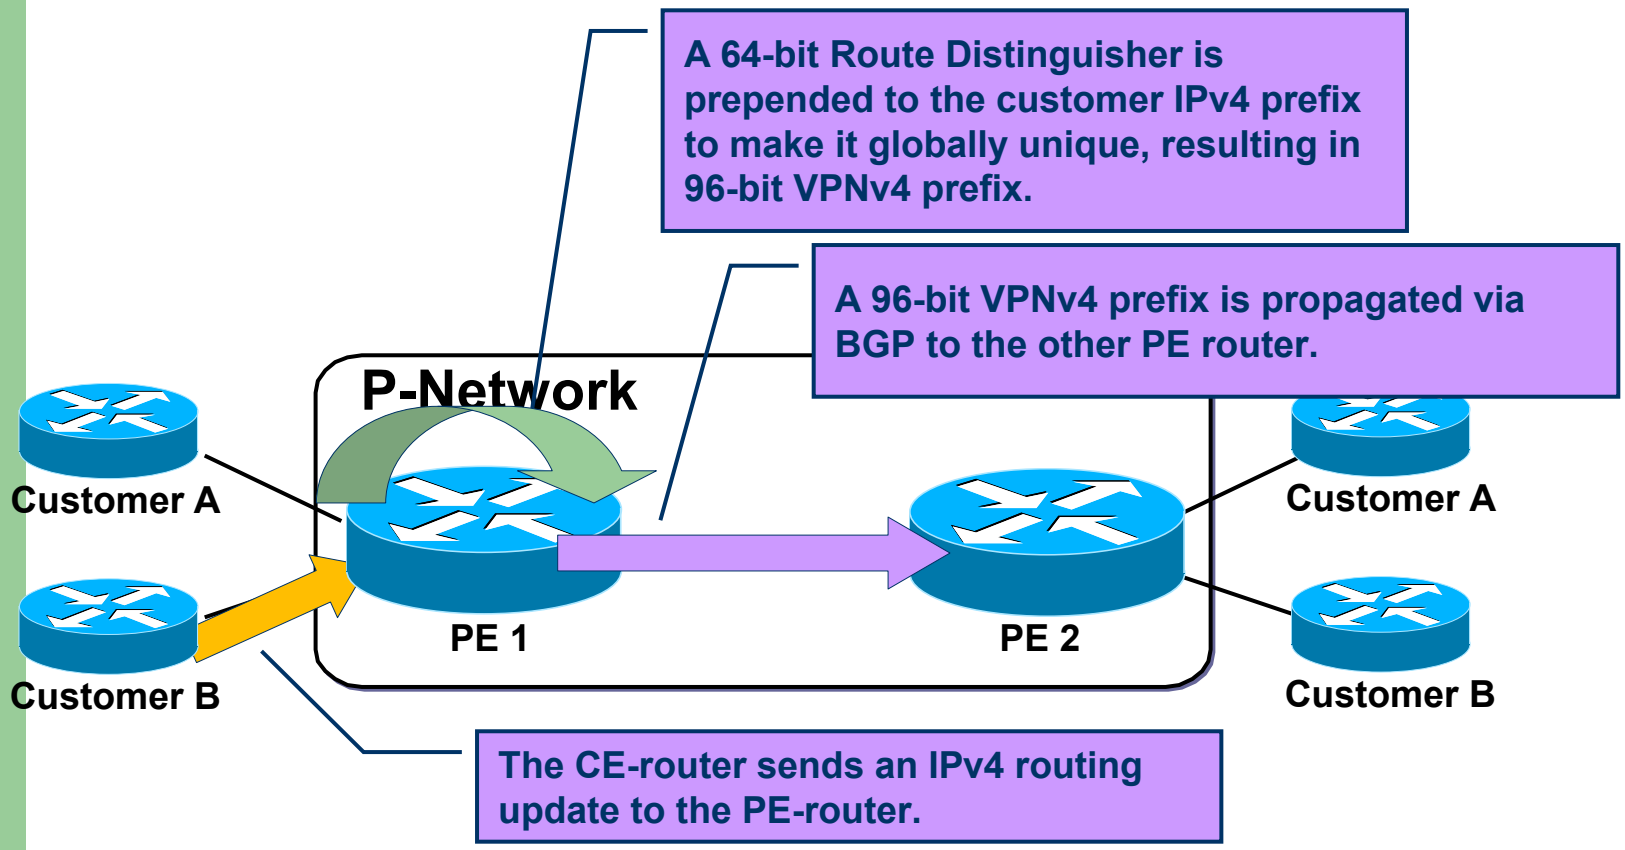

# Route Distinguisher Usage in an MPLS VPN

The RD is removed from the VPNv4 prefix, resulting in a 32-bit IPv4 prefix.

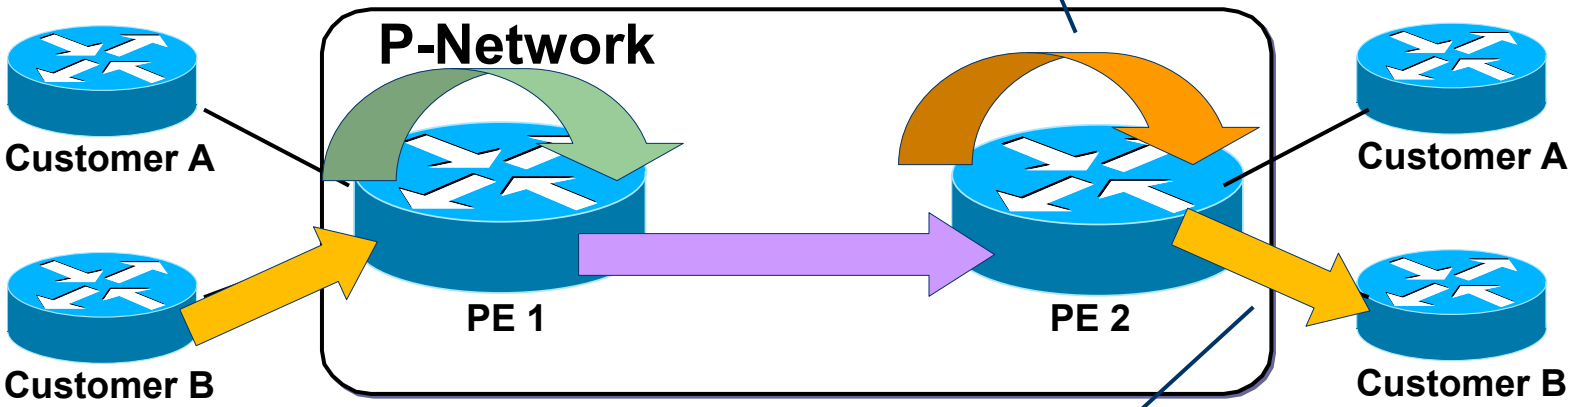


The PE router sends the resulting IPv4 prefix to the CE router.

# Route Distinguisher Usage in an MPLS VPN

- The RD has no special meaning — it is used only to make potentially overlapping IPv4 addresses globally unique.
- 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.

# Complex VPN—Sample VoIP Service

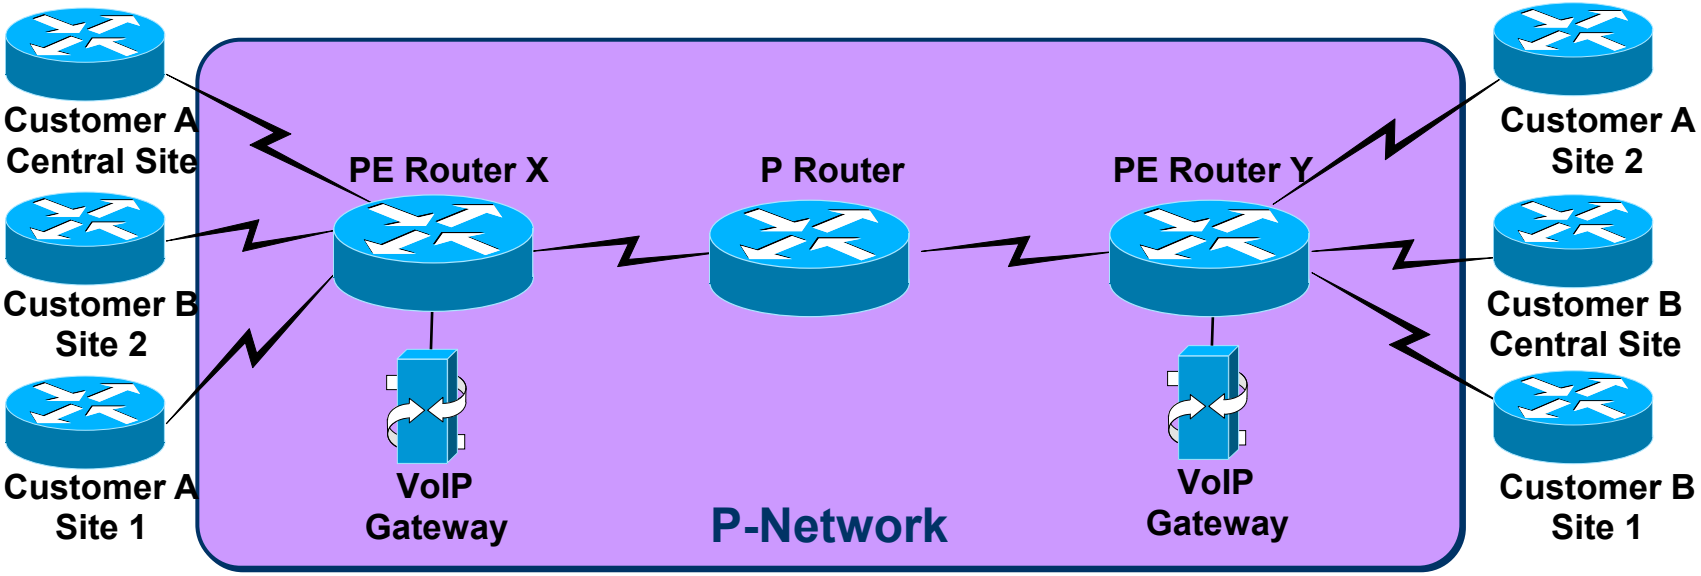


**Requirements:**

- All sites of one customer need to communicate.
- Central sites of both customers need to communicate with VoIP gateways and other central sites.
- Other sites from different customers do not communicate with each other.

# Sample VoIP Service Connectivity Requirements

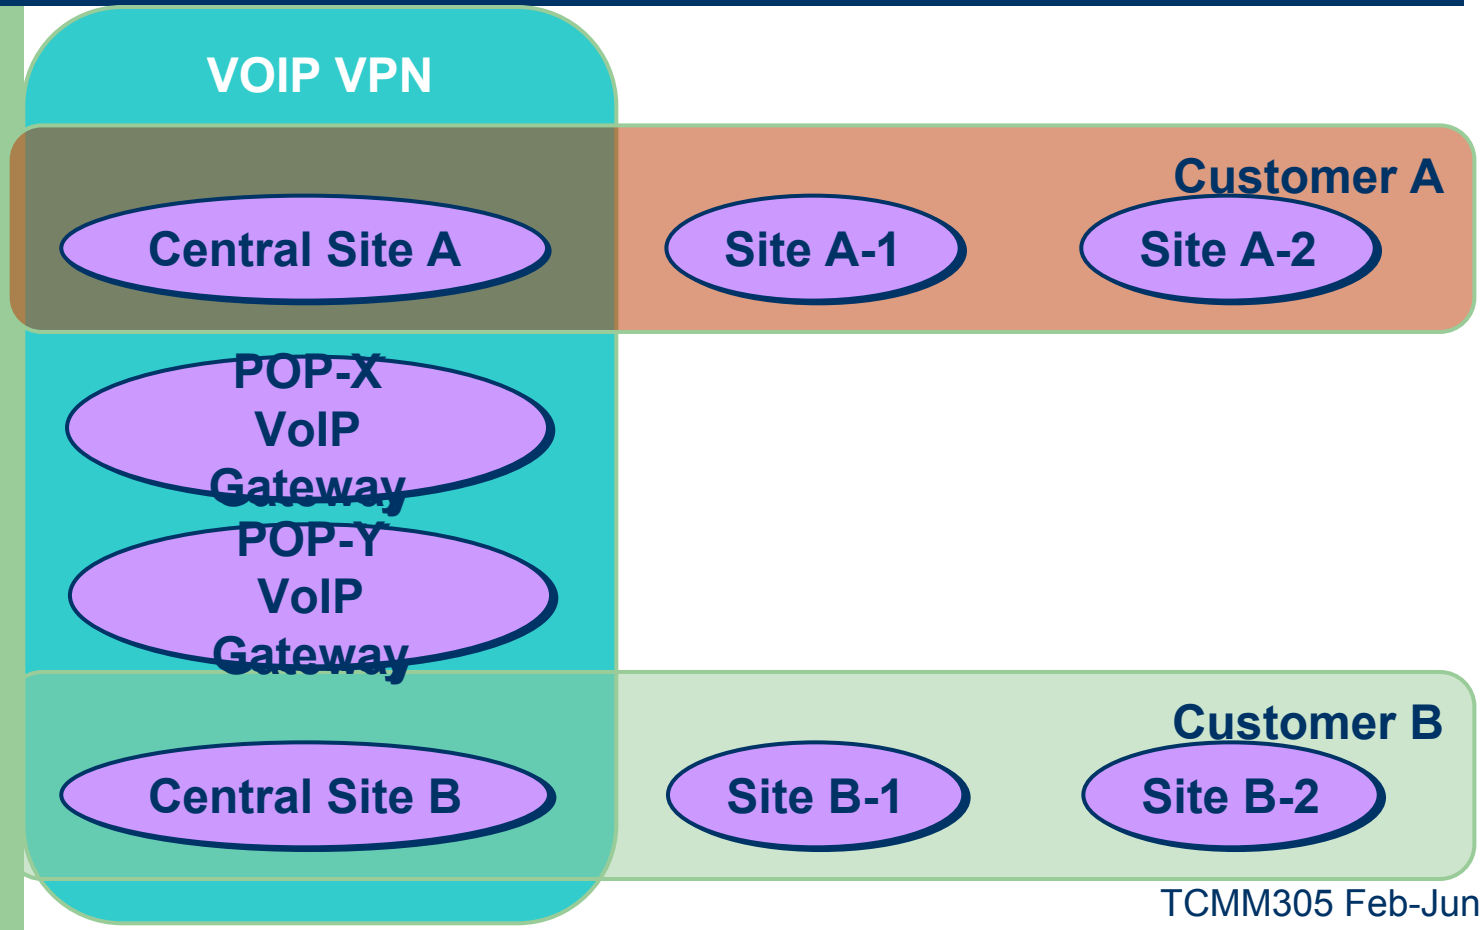

# Route Targets

- 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.
- RDs were introduced in the MPLS VPN architecture to support complex VPN topologies.

# What Are Route Targets?

- Route targets (RTs) are additional attributes attached to VPNv4 BGP routes to indicate VPN membership.
- Extended BGP communities are used to encode these attributes.
  - Extended communities carry the meaning of the attribute together with its value.
- Any number of RTs can be attached to a single route.

# How Do Route Targets Work?

- Export RTs identifying VPN membership are appended to the customer route when it is converted into a VPNv4 route.
- Each virtual routing table has a set of associated import RTs that select routes to be inserted into the virtual routing table.
- Route targets usually identify VPN membership, but they can also be used in more complex scenarios.

# Virtual Private Networks Redefined

**With the support of complex VPN topologies, VPNs have to be redefined**

- A VPN is a collection of sites sharing common routing information.
- A site can be part of different VPNs.
- A VPN can be seen as a community of interest (closed user group—CUG).
- Complex VPN topologies are supported by multiple virtual routing tables on the PE routers.

# Impact of Complex VPN Topologies on Virtual Routing Tables

- A virtual routing table in a PE router can be used only for sites with identical connectivity requirements.
- Complex VPN topologies require more than one virtual routing table per VPN.
- As each virtual routing table requires a distinct RD value, the number of RDs in the MPLS VPN network increases.

# Sample VoIP Service Virtual Routing Tables

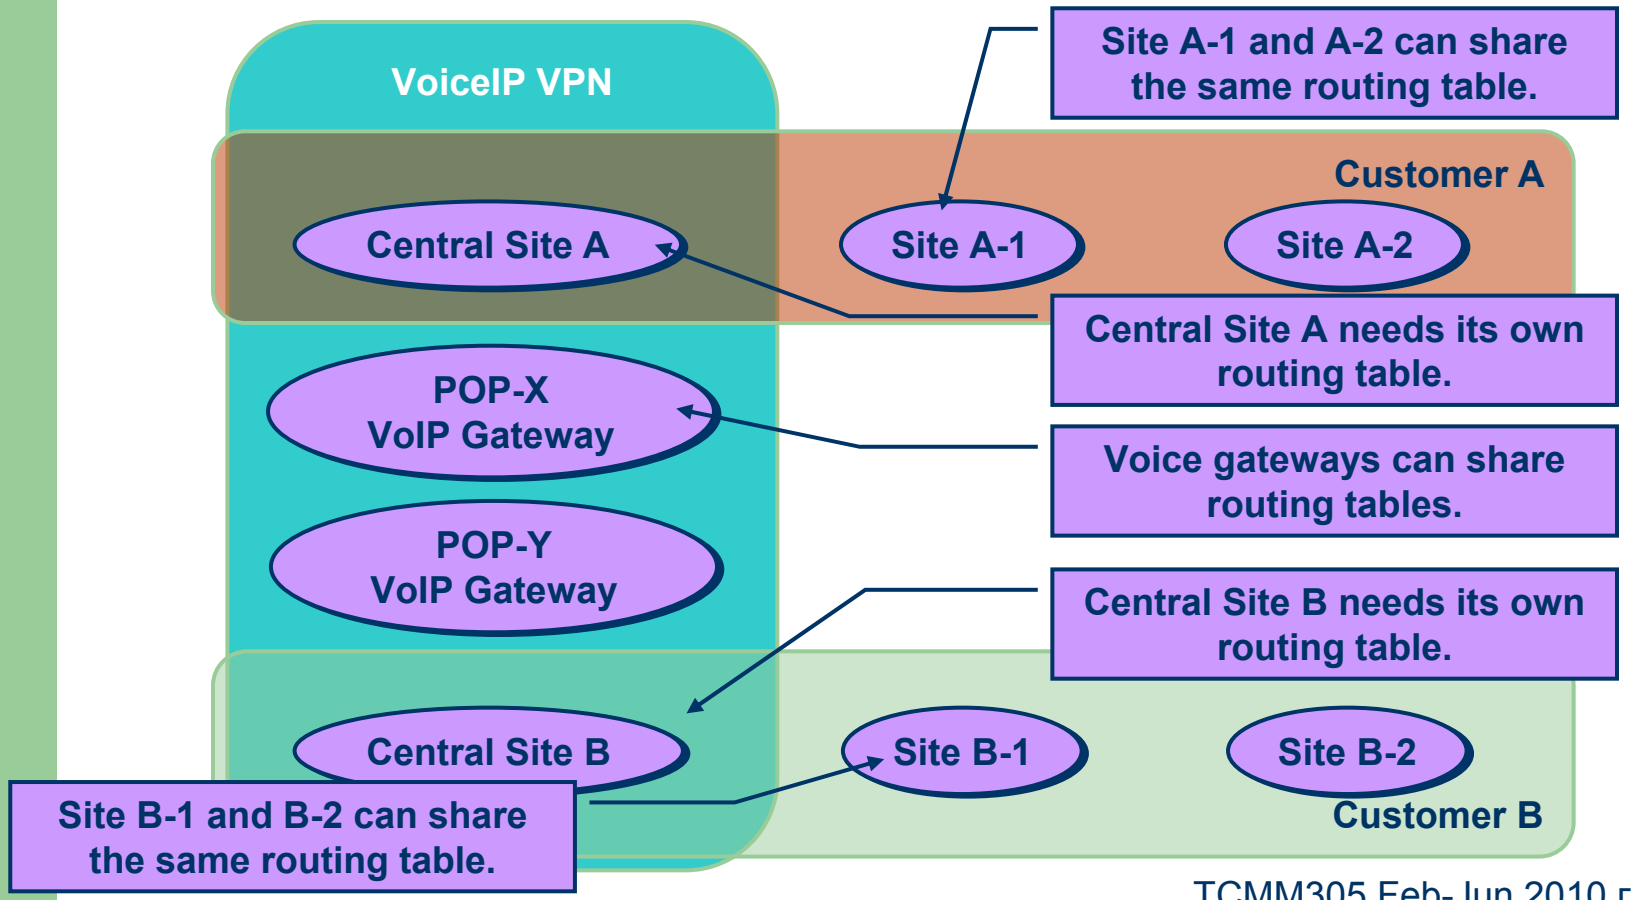

# Benefits of MPLS VPN Technology

- MPLS VPN technology has all the benefits of peer-to-peer VPN technology:
  - Easy provisioning
  - Optimal routing
- It also bypasses most drawbacks of traditional peer-to-peer VPN technologies:
  - RDs enable overlapping customer address spaces.
  - RTs enable topologies that were hard to implement with other VPN technologies.

# Summary

After completing this section, you should be able to perform the following tasks:

- Describe the difference between traditional peer-to-peer models and MPLS VPN
- List the benefits of MPLS VPN
- Describe major architectural blocks of MPLS VPN
- Explain the need for route distinguishers and route targets

# Review Questions

- How does MPLS VPN support overlapping customer address spaces?
- How are customer routes exchanged across the P-network?
- What is a route distinguisher?
- Why is the RD not usable as VPN identifier?
- What is a route target?
- Why were the route targets introduced in MPLS VPN architecture?
- How are route targets used to build virtual routing tables in the PE routers?
- What is the impact of complex VPN topologies on virtual routing tables in the PE routers?

# MPLS VPN Routing Model

TCMM305 Feb-Jun 2010 г.

# Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the routing model of MPLS VPN
- Describe the MPLS VPN routing model from customer and provider perspectives
- Identify the routing requirements of CE-routers, PE-routers and P-routers

# MPLS VPN Routing Requirements

- Customer routers (CE routers) have to run standard IP routing software.
- Provider core routers (P routers) have no VPN routes.
- Provider edge routers (PE routers) have to support MPLS VPN and Internet routing.

# MPLS VPN Routing—CE Router Perspective

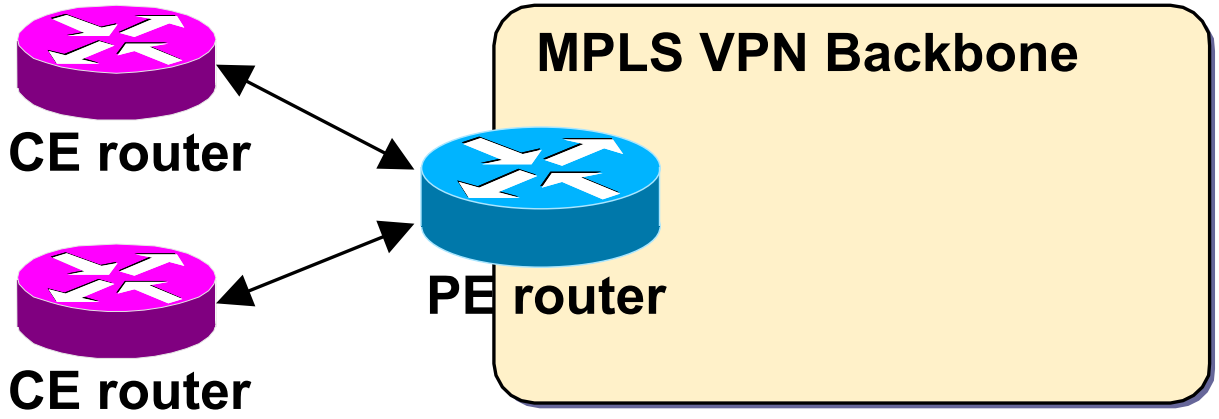


- The CE routers run standard IP routing software and exchange routing updates with the PE router.
  - External BGP (EBGP), Open Shortest Path First (OSPF), RIP version 2 (RIPv2), and static routes are supported.
- The PE router appears as another router in the C-network.

# MPLS VPN Routing—Overall Customer Perspective

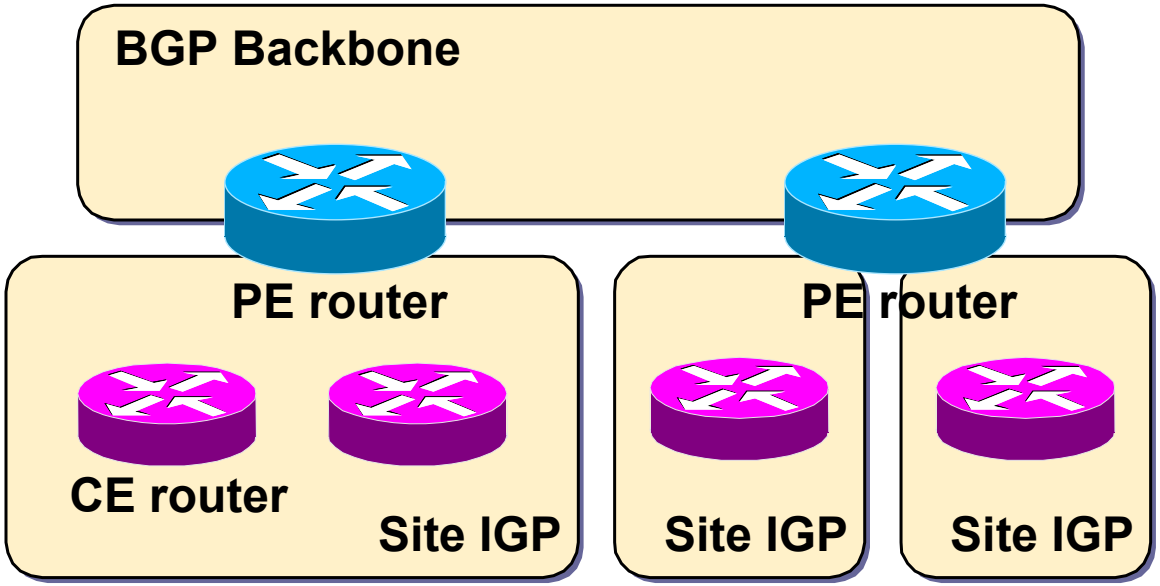


- To the customer, the PE routers appear as core routers connected via a BGP backbone.
- The usual BGP and IGP design rules apply.
- The P routers are hidden from the customer.

# MPLS VPN Routing—P Router Perspective

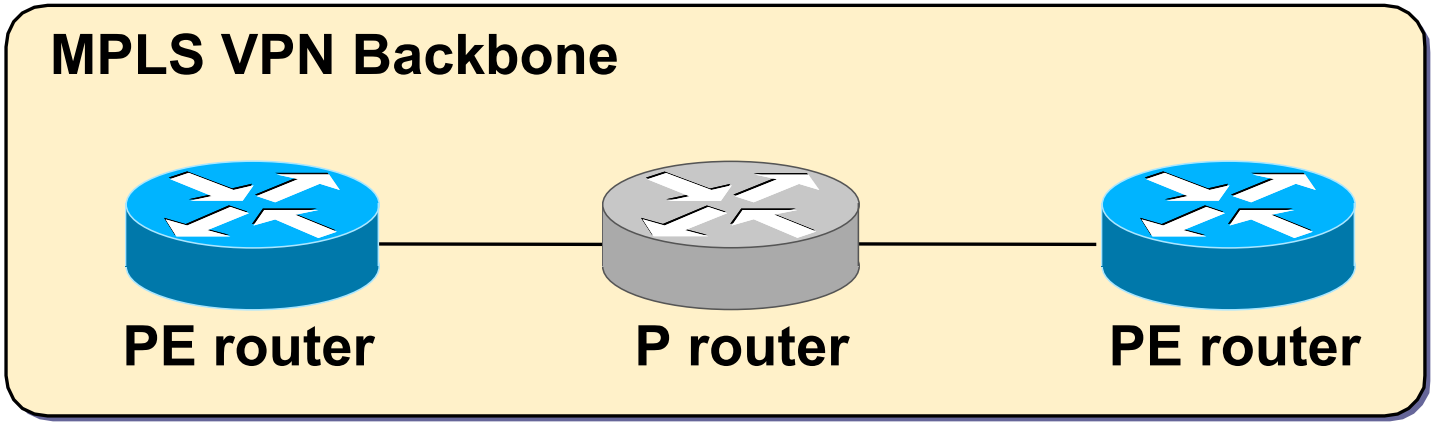


- P routers do not participate in MPLS VPN routing and do not carry VPN routes.
- P routers run backbone IGP with the PE routers and exchange information about global subnets (core links and loopbacks).

# MPLS VPN Routing—PE Router Perspective

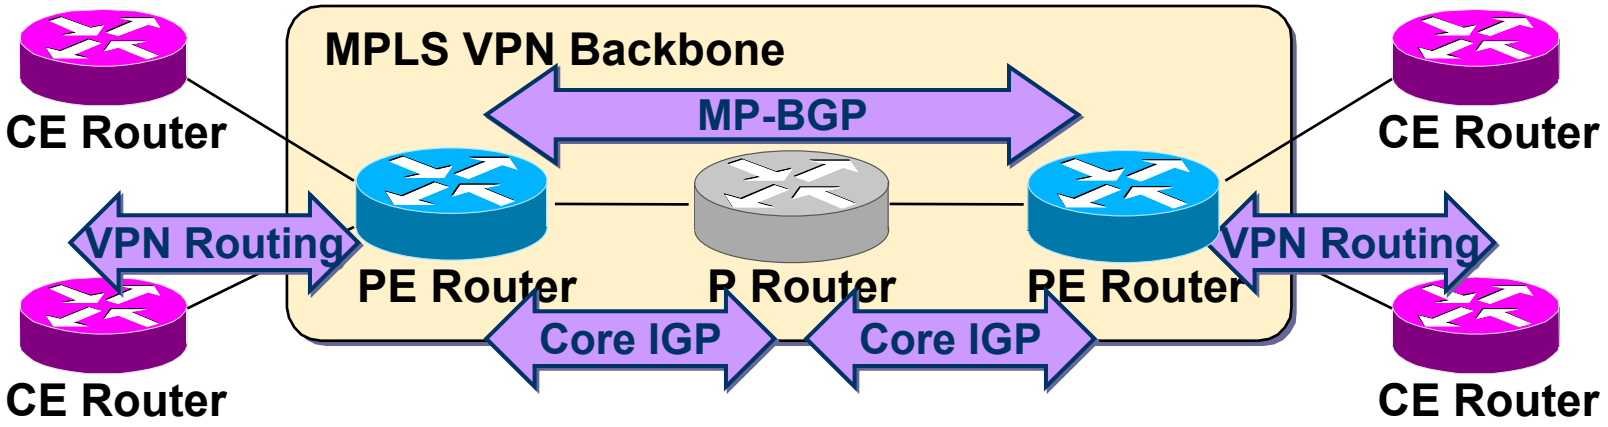


- PE routers:
  - Exchange VPN routes with CE routers via per-VPN routing protocols
  - Exchange core routes with P routers and PE-routers via core IGP
  - Exchange VPNv4 routes with other PE routers via MP- IBGP sessions

# MPLS VPN Support for Internet Routing

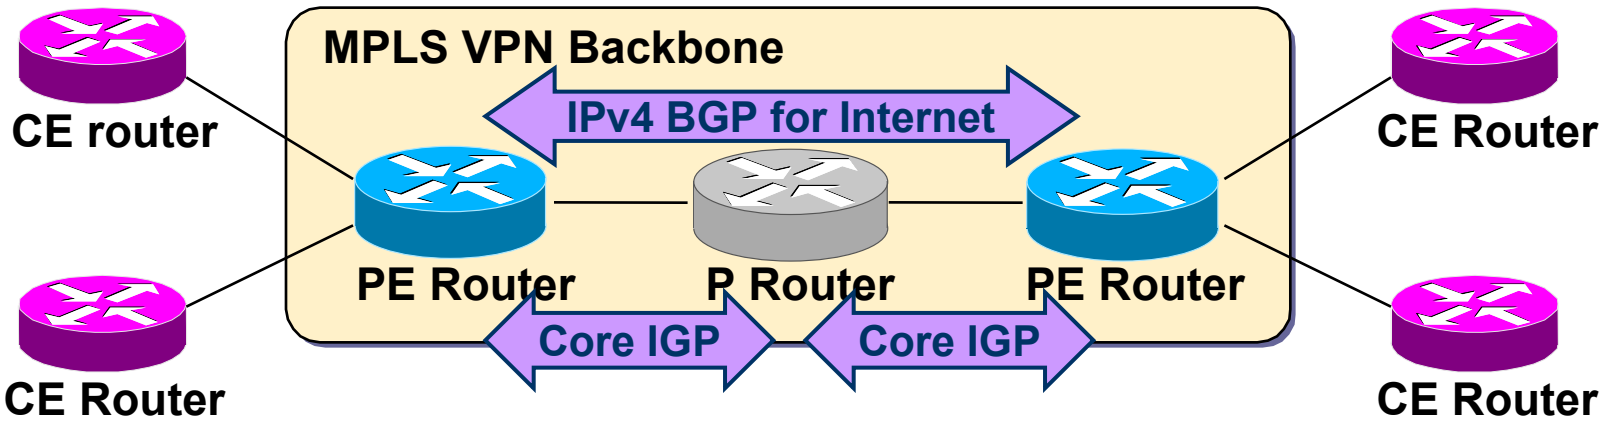


PE routers can run standard IPv4 BGP in the global routing table:

- PE routers exchange Internet routes with other PE routers.
- CE routers do not participate in Internet routing.
- P routers do not need to participate in Internet routing.

# Routing Tables on PE Routers

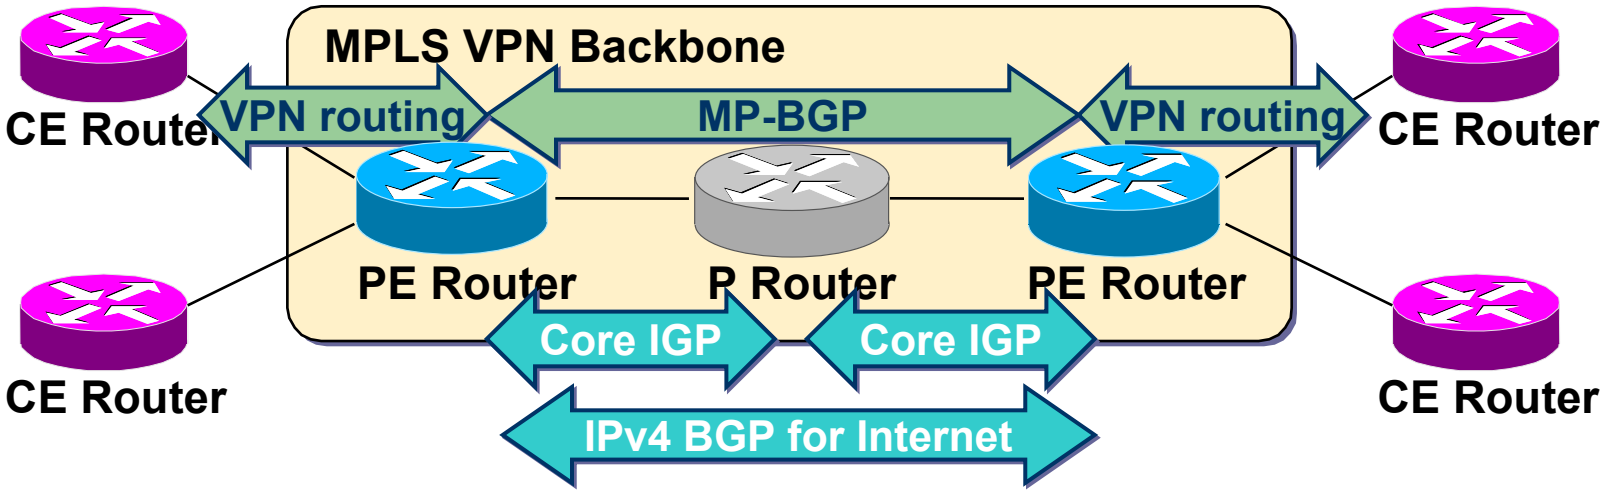


- PE routers contain a number of routing tables:
  - Global routing table that contains core routes (filled with core IGP) and Internet routes (filled with IPv4 BGP).
  - VRF tables for sets of sites with identical routing requirements.
  - VRFs filled with information from CE routers and MP-BGP information from other PE routers.

# MPLS VPN End-to-End Routing Information Flow (1/3)

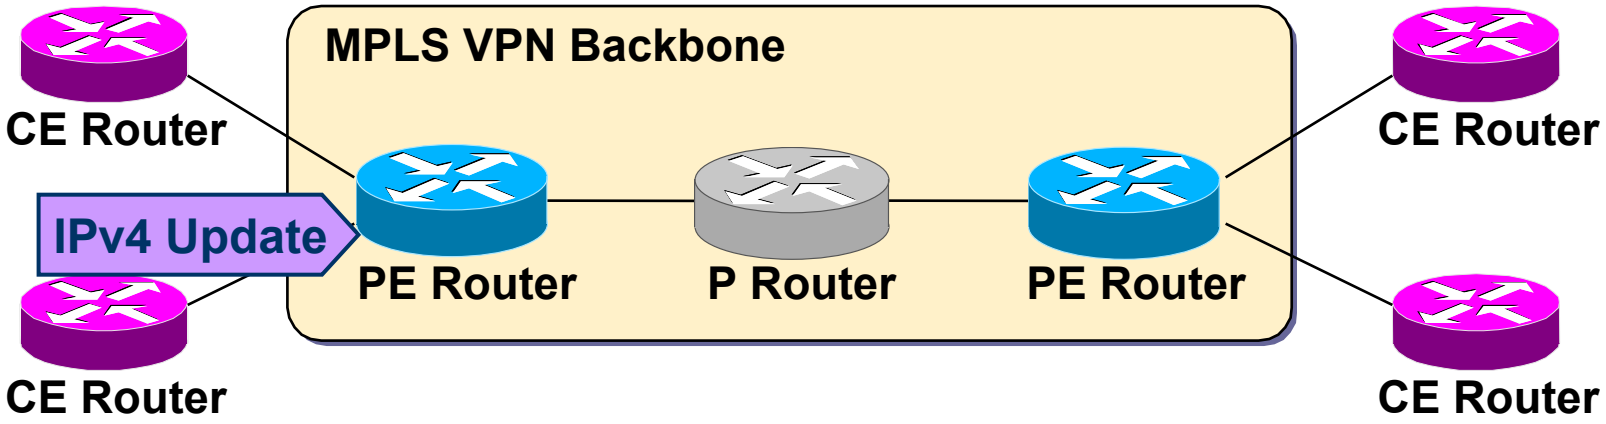


- PE routers receive IPv4 routing updates from CE routers and install them in the appropriate VRF table.

# MPLS VPN End-to-End Routing Information Flow (2/3)

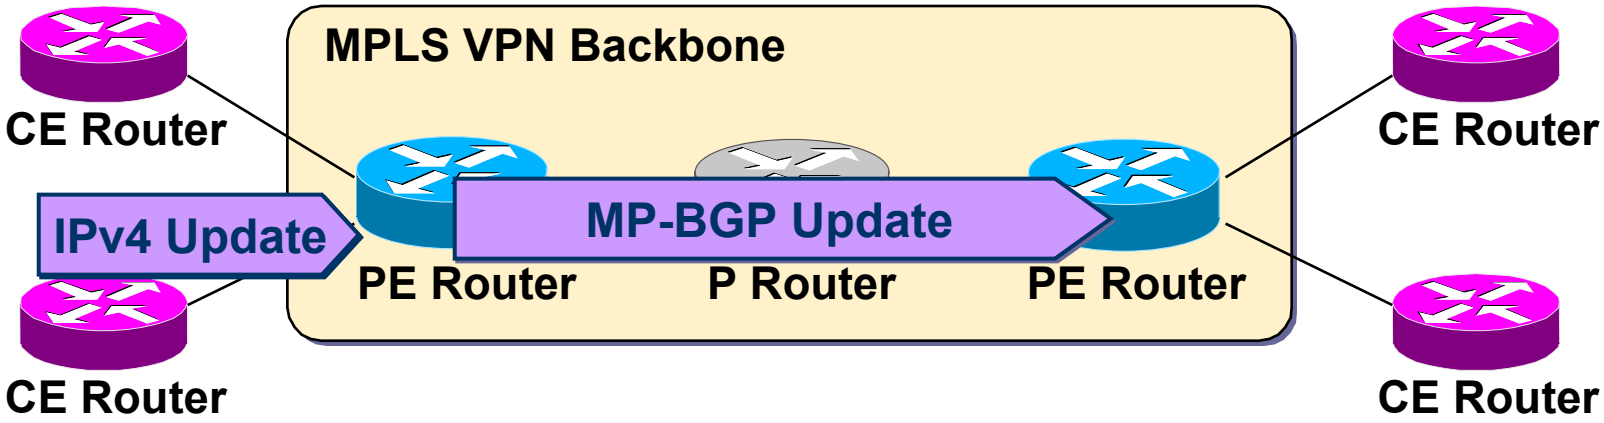


- PE routers export VPN routes from VRF tables into MP-BGP and propagate them as VPNv4 routes to other PE routers.
- A full mesh of MP-IBGP sessions is needed between PE routers.

# MP-BGP Update

- An MP-BGP update contains:
  - VPNv4 address
  - Extended communities (route targets, optionally Site-of-Origin, or SOO)
  - Label used for VPN packet forwarding
  - Any other BGP attribute (for example, AS path, local preference, multi-exit discriminator (MED), standard community)

# MP-BGP Update - VPNv4 Address

- A VPN IPv4 address contains:
  - RD
  - 64 bits
  - Makes the IPv4 route globally unique
  - RD is configured in the PE  for each VRF
  - RD may or may not be related to a site or a VPN
  - IPv4 address (32 bits)

# MP-BGP Update - Extended Communities

- 64-bit attribute attached to a route
- Set of communities can be attached to a single route
- High-order 16 bits identify extended community type
  - RT: identifies the set of sites to which the route must be advertised
  - SOO: identifies the originating site
  - OSPF route type: identifies the link-state advertisement (LSA) type of OSPF route redistributed into MP-BGP

# Extended BGP Community Display Format

- Two display formats are supported:
  - <16bits type>:<ASN>:<32 bit number>
    -Uses registered AS number
  - <16bits type>:<IP address>:<16 bit number>
    -Uses registered IP address

# MPLS VPN End-to-End Routing Information Flow (3/3)

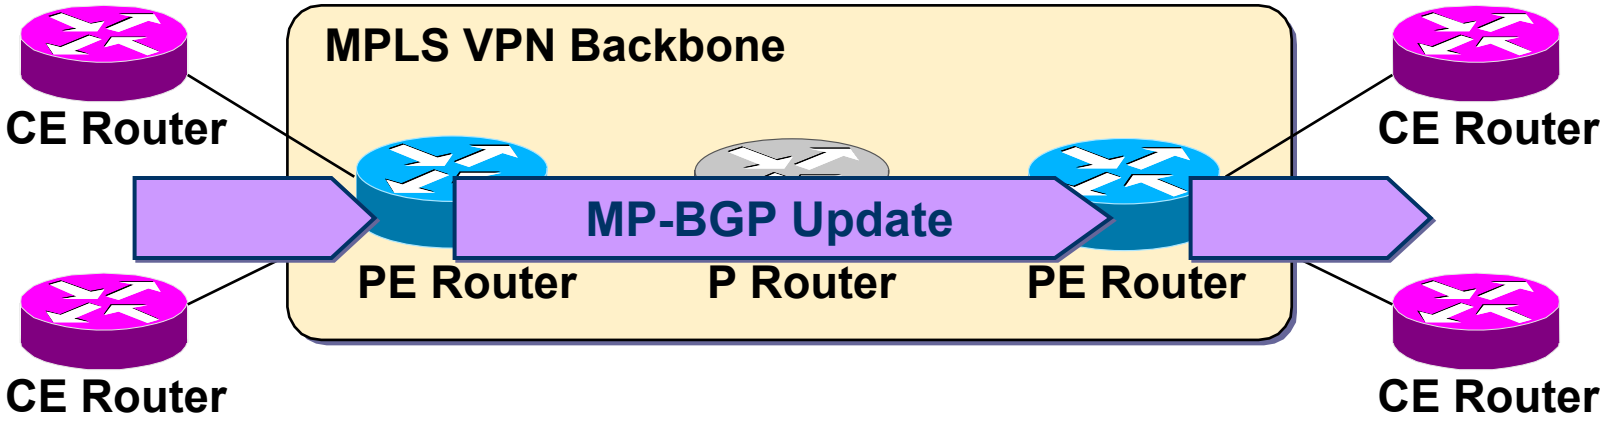


- Receiving PE router imports incoming VPNv4 routes into the appropriate VRF based on route targets attached to the routes
- Routes installed in VRF are propagated to CE routers

# Route Distribution to CE Routers

- Route distribution to sites is driven by the SOO and RT EBGP communities.
- A route is installed in the site VRF that matches the RT attribute.
  - A PE router that connects sites belonging to multiple VPNs will install the route into the site VRF if the RT attribute contains one or more VPNs to which the site is associated.

# Summary

- After completing this section, you should be able to perform the following tasks:
  - Describe the routing model of MPLS VPN
  - Describe the MPLS VPN routing model from customer and provider perspective
  - Identify the routing requirements of CE-routers, PE-routers and P-routers

# Review Questions

- What is the impact of MPLS VPN on CE-routers?
- What is the customer's perception of end-to-end MPLS VPN routing?
- What is the P-router perception of end-to-end MPLS VPN routing?
- How many routing tables does a PE-router have?
- How many routing tables reside on a P-router?
- Which routing protocols fill the global routing table of a PE-router?
- Which routing protocols fill the Virtual Routing table of a PE-router?

# More Review Questions

- How is the Internet routing supported by MPLS VPN architecture?
- How is the VPN routing information exchanged between the PE-routers?
- Which attributes are always present in a MP-BGP update?
- Which attributes can be optionally present in a MP-BGP update?
- Which BGP attributes drive the import of VPNv4 route into a VRF?
- Which BGP attributes control the VPN route distribution toward CE-routers?

# MPLS VPN Packet Forwarding

TCMM305 Feb-Jun 2010 г.

# Objectives

- Upon completion of this section, you will be able to perform the following tasks:
  - Describe the MPLS VPN forwarding mechanisms
  - Describe the VPN and backbone label propagation
  - Explain the need for end-to-end LSP between PE routers
  - Explain the implications of BGP next-hop on MPLS VPN forwarding

# VPN Packet Forwarding Across an MPLS VPN Backbone

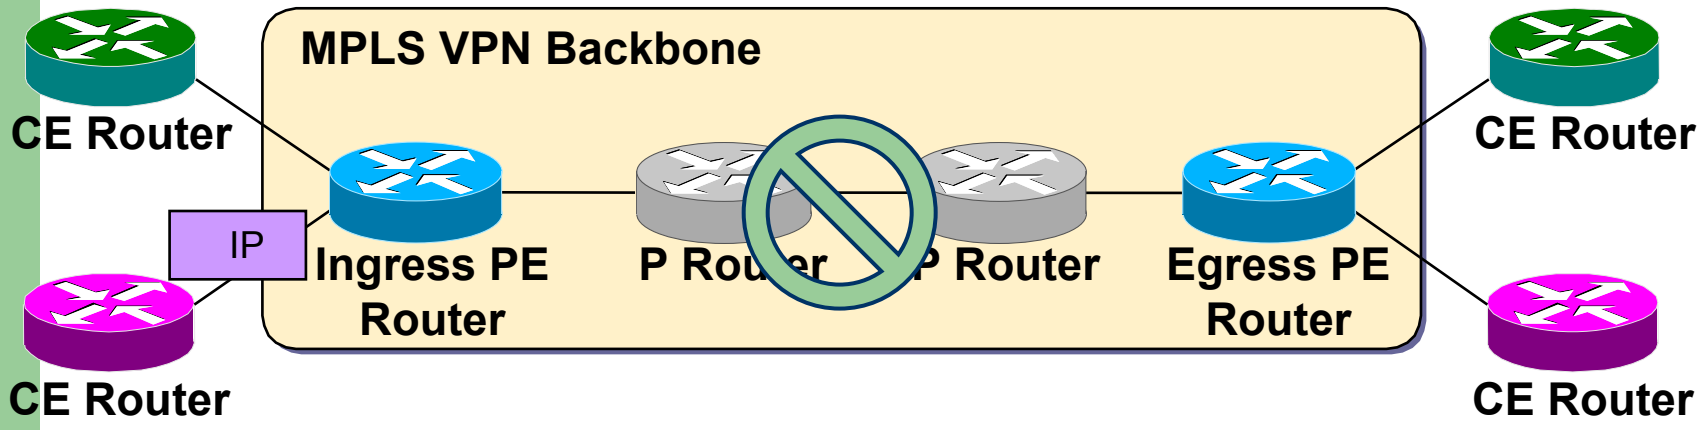


**Q:** How will the PE routers forward the VPN packets across the MPLS VPN backbone?

**A1:** They will forward pure IP packets.

Wrong answers:

- P routers do not have VPN routes; the packet is dropped on IP lookup.
- How about using MPLS for packet propagation across the backbone?

# VPN Packet Forwarding Across an MPLS VPN Backbone

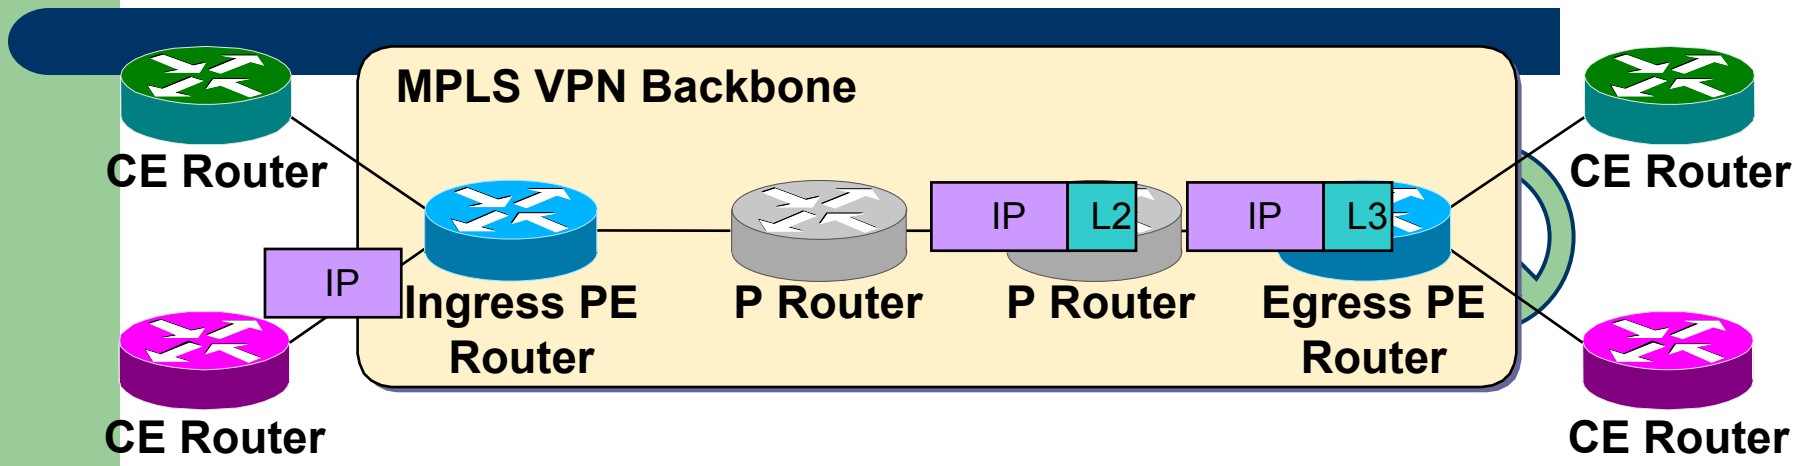


**Q:** How will the PE routers forward the VPN packets across the MPLS VPN backbone?

**A2:** They will label the VPN packets with a label distribution protocol (LDP) label for the egress PE router and forward the labeled packets across the MPLS backbone.

Better answers:

- The P routers perform the label switching and the packet reaches the egress PE router.
- However, the egress PE router does not know which VRF to use for packet switching, so the packet is dropped.
- How about using a label stack?

# VPN Packet Forwarding Across an MPLS VPN Backbone

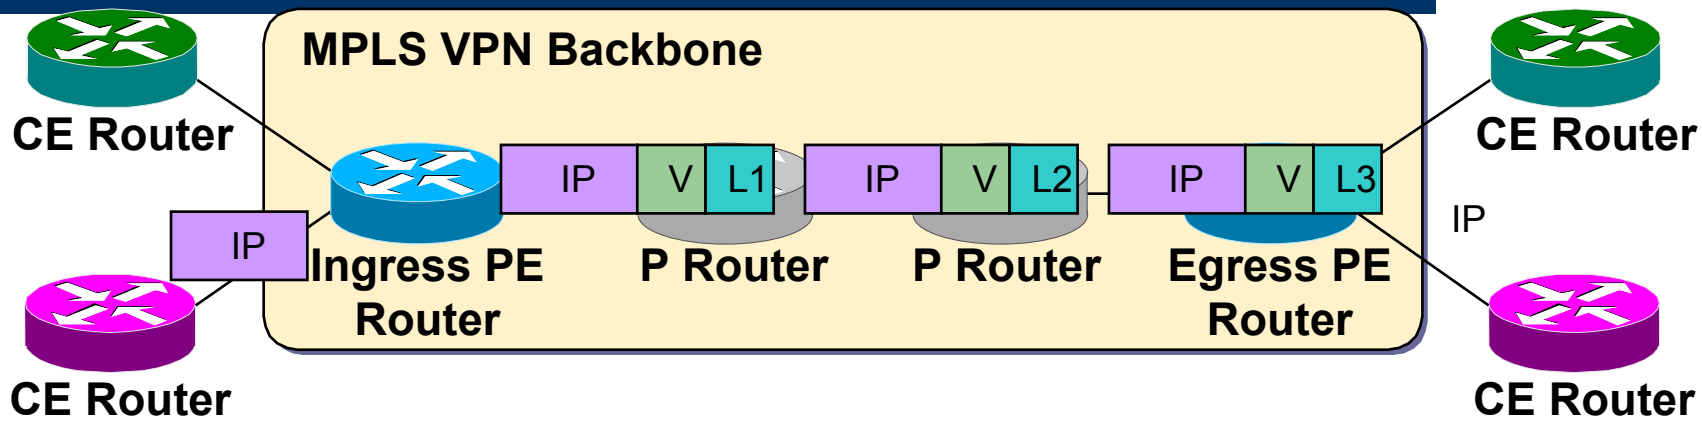


**Q:** How will the PE routers forward the VPN packets across the MPLS VPN backbon

**A3:** They will label the VPN packets with a label stack, using the LDP label for the egress PE router as the top label and the VPN label assigned by the egress PE router as the second label in the stack.

# VPN Packet Forwarding - Penultimate Hop Popping

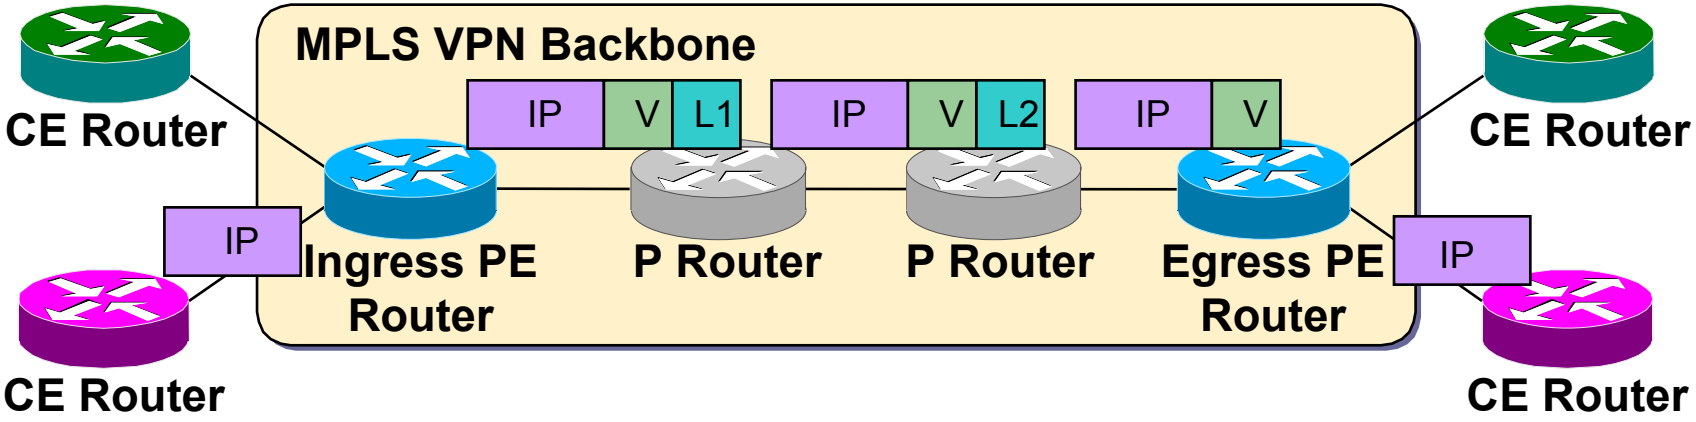


- Penultimate hop popping on the LDP label can be performed on the last P router.
- The egress PE router performs label lookup only on the VPN label, resulting in faster and simpler label lookup.
- IP lookup is performed only once—in the ingress PE router.

# VPN Label Propagation

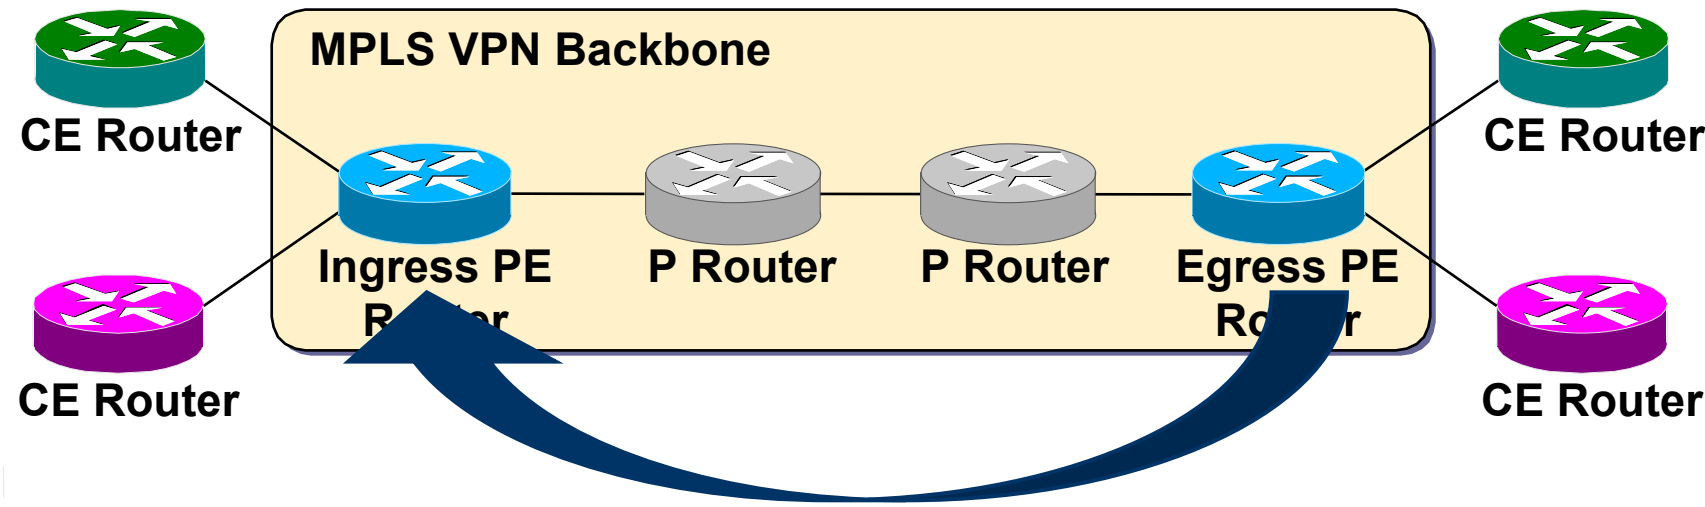


**Q:** How will the ingress PE router get the second label in the label stack from the egress PE router?

**A:** Labels are propagated in MP BGP VPNv4 routing updates.

# VPN Label Propagation

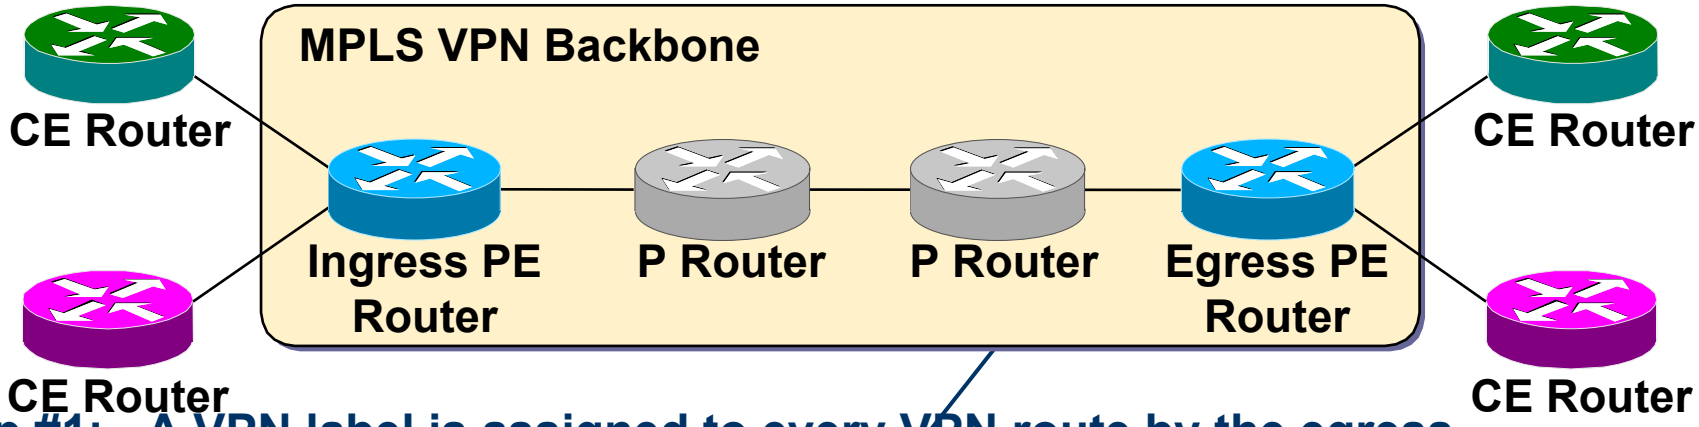


**Step #1:** A VPN label is assigned to every VPN route by the egress PE router.

```
Egress-PE#show tag-switching forwarding vrf SiteA2
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
26     Aggregate   150.1.31.36/30[V] 0
37     Untagged    203.1.2.1/32[V]   0          Se1/0.20    point2point
38     Untagged    203.1.20.0/24[V]  0          Se1/0.20    point2point
```

# VPN Label Propagation

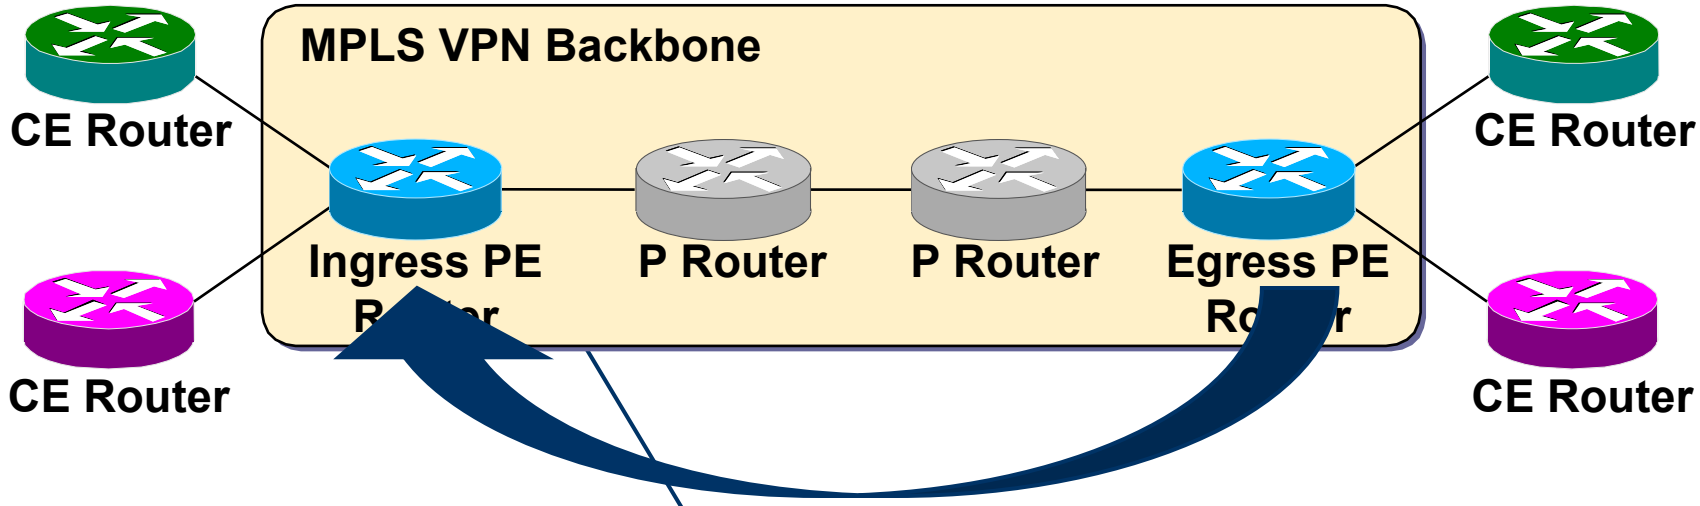


**Step #2:** The VPN label is advertised to all other PE routers in an MP-BGP update.

```
Ingress-PE#show ip bgp vpnv4 all tags
   Network                Next Hop      In tag/Out tag
Route Distinguisher: 100:1 (vrf1)
   12.0.0.0          10.20.0.60      26/notag
                     10.20.0.60      26/notag
   203.1.20.0        10.15.0.15      notag/38
```

# VPN Label Propagation

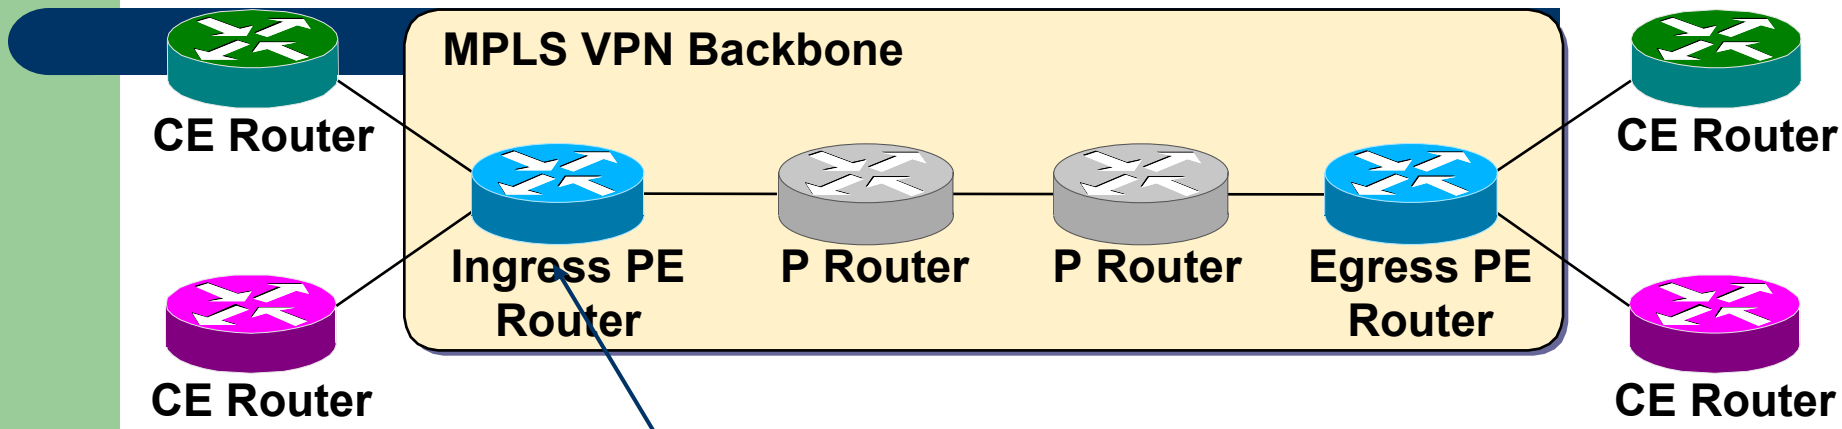


**Step #3: A label stack is built in VFR table.**

```
Ingress-PE#show ip cef vrf Vrf1 203.1.20.0 detail
203.1.20.0/24, version 57, cached adjacency to Serial1/0.2
0 packets, 0 bytes
  tag information set
    local tag: VPN-route-head
    fast tag rewrite with Se1/0.2, point2point, tags imposed: {26 38}
  via 192.168.3.103, 0 dependencies, recursive
    next hop 192.168.3.10, Serial1/0.2 via 192.168.3.103/32
    valid cached adjacency
    tag rewrite with Se1/0.2, point2point, tags imposed: {26 38}
```

# Effects of MPLS VPN Label Propagation

- The VPN label must be assigned by the BGP next hop.
- The BGP next hop should not be changed in the MP-IBGP update propagation.
  - Do not use next-hop-self on confederation boundaries.
- The PE router must be the BGP next hop.
  - Use next-hop-self on the PE router.
- The label must be reoriginated if the next hop is changed.
  - A new label is assigned every time the MP-BGP update crosses the AS boundary where the next hop is changed.
  - This functionality is supported by Cisco IOS Releases 12.1(4)T, 12.2, and later.

# Effects of MPLS VPN Packet Forwarding

- The VPN label is understood only by the egress PE router.
- An end-to-end LSP tunnel is required between the ingress and egress PE routers.
- BGP next hops must not be announced as BGP routes.
- LDP labels are not assigned to BGP routes.
- BGP next hops announced in IGP must not be summarized in the core network.
  - Summarization breaks the LSP tunnel.

# VPN Packet Forwarding with Summarization in the Core

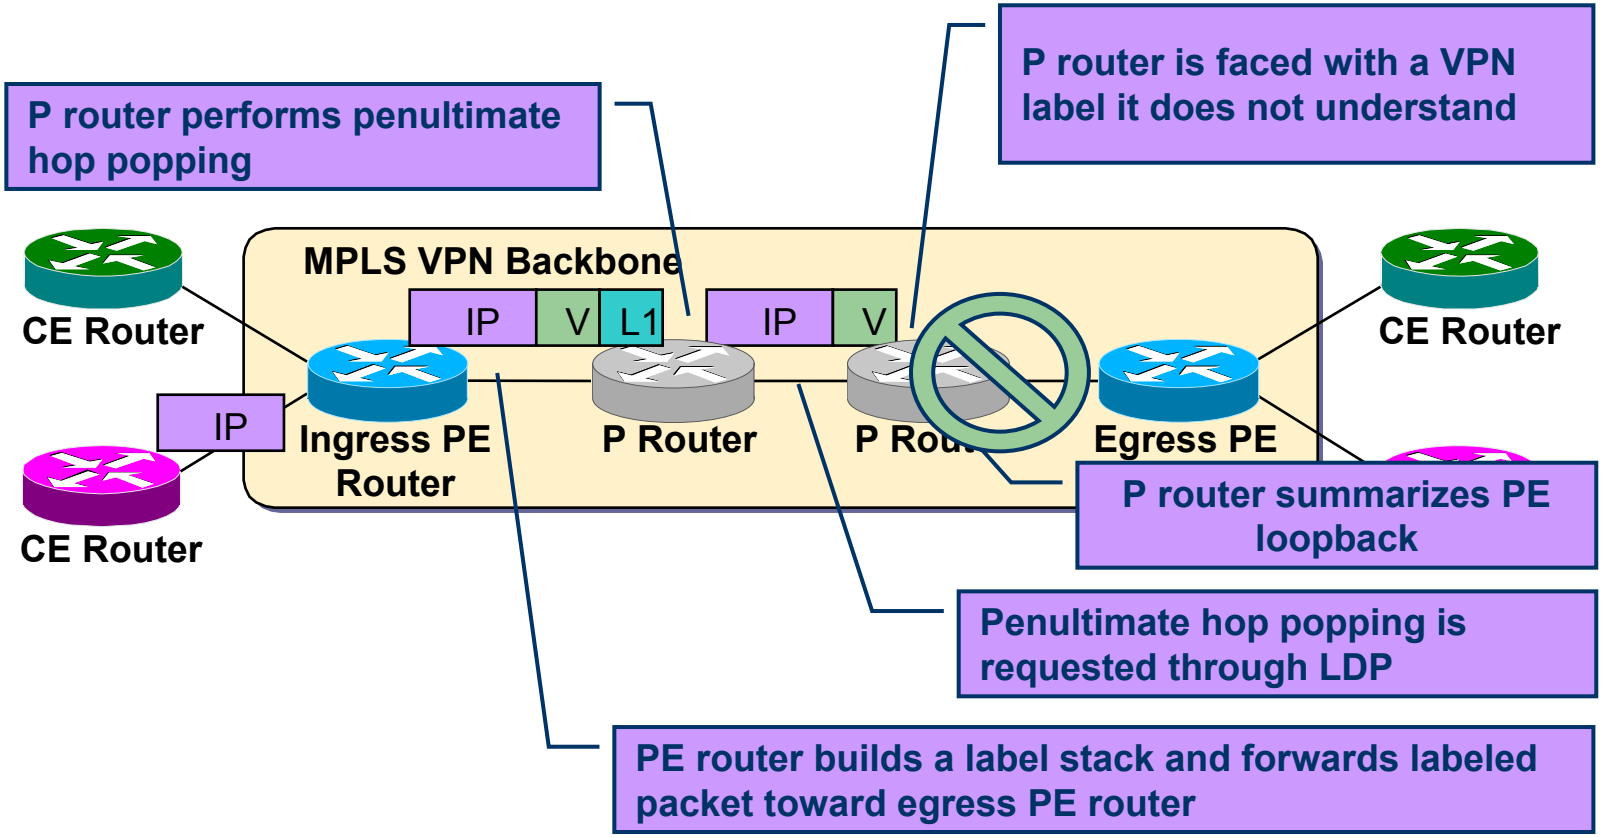

# Summary

After completing this section, you should be able to perform the following tasks:

- Describe the MPLS VPN forwarding mechanisms
- Describe the VPN and backbone label propagation
- Explain the need for end-to-end LSP between PE routers
- Explain the implications of BGP next-hop on MPLS VPN forwarding

# Review Questions

- How are VPN packets propagated across MPLS VPN backbone?
- How can P-routers forward VPN packets if they don't have VPN routes?
- How is the VPN label propagated between PE-routers?
- Which router assigns the VPN label?
- How is the VPN label used on other PE-routers?
- What is the impact of changing BGP next-hop on MP-BGP update?
- How are MP-BGP updates propagated across AS boundary?
- What is the impact of BGP next-hop summarization in the network core?

# Summary

After completing this lesson, you should be able to perform the following tasks:

- Identify major Virtual Private network topologies, their characteristics and usage scenarios
- Describe the differences between overlay VPN and peer-to-peer VPN
- List major technologies supporting overlay VPNs and peer-to-peer VPNs
- Position MPLS VPN in comparison with other peer-to-peer VPN implementations
- Describe major architectural blocks of MPLS VPN
- Describe MPLS VPN routing model and packet forwarding

# ???
# Мрежи за достъп от следващо поколение

**инж. Николай Милованов**
**email: *nmil@niau.org***
**Skype: niau33**

Нов български университет