WORKSHOP MPLS

fbolanos@cisco.com

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS Concepts

- **MPLS: Multi Protocol Label Switching**
- **Packet forwarding is done based on Labels Multi-protocol Label Switching (MPLS) is a**
- **Labels may correspond to IP destination**
- **Labels can also correspond to other parameters (QoS, source address, ...).**
- **MPLS was designed to support forwarding of other protocols as well.**

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS/VPN Terminology

- **Provider Network (P-network)**

  backbone under the control of the Service Provider
- **Customer Network (C-network)**

  network under VPN customer control
- **CE Router**

  part of the C-network and interfaces to a PE router
- **PE Router**

  part of the P-network and interfaces to CE routers

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS/VPN Terminology

- **P Router**

  provider (core) router without knowledge of VPN

- **Site**

  set of (sub)networks which are part of the Customer network and co-located

  connected to the MPLS/VPN backbone through one of more PE/CE links

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS Concepts

- Unlike IP, classification/label can be based on:
  - Destination Unicast address
  - Traffic Engineering
  - VPN
  - QoS
- FEC: Forwarding Equivalence Class
- A FEC can represent a: Destination address prefix, VPN, Traffic Engineering tunnel, Class of Service.

© 2001, Cisco Systems, Inc. All rights reserved.

Control Component

Several protocols for label exchange:

Control Component is responsible for binding between labels and routes

LDP (646)/ TDP (711)

Maps unicast IP destinations into labels

- RSVP, CR-LDP

  Used in traffic engineering

- BGP/M-BGP

  External labels (VPN)

- PIM

  For multicast states label mapping

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS Specific Tables

- **Each LSR will use a LIB**

  Label Information Base

  Contains all label/prefix mappings from all TDP/LDP neighbours

- **Each LSR will also use an LFIB**

  Label Forwarding Information Base

  Contains only label/prefix mappings that are currently in use for label forwarding

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS/VPN Model

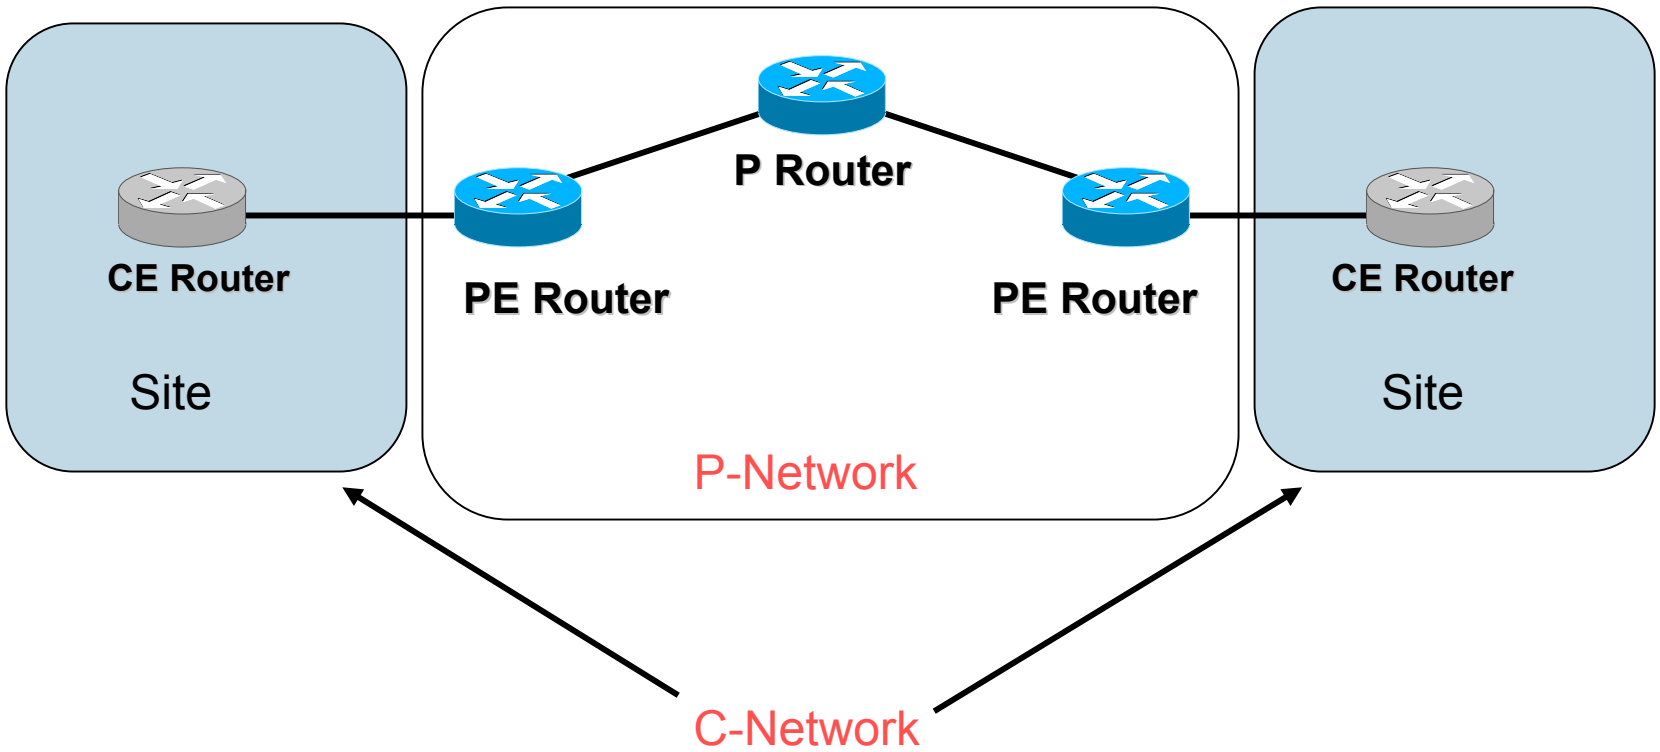

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS Architecture

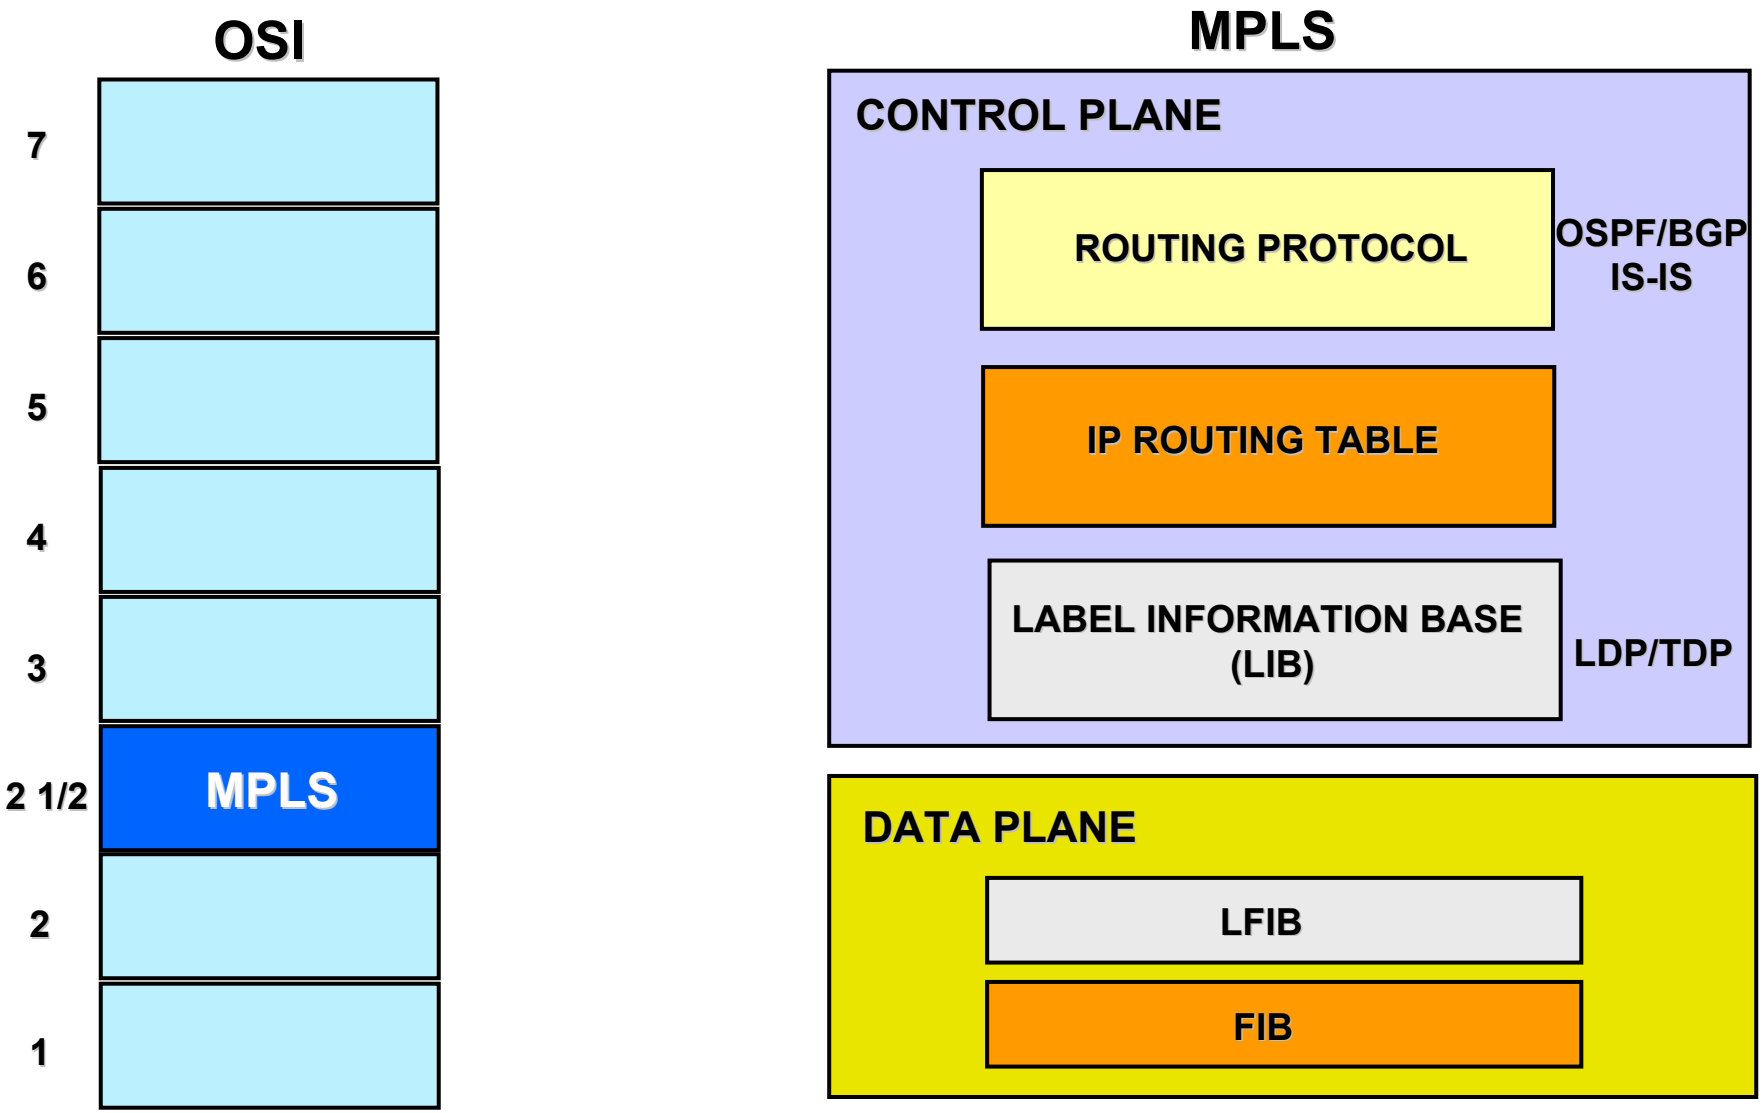

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS Shim header structure

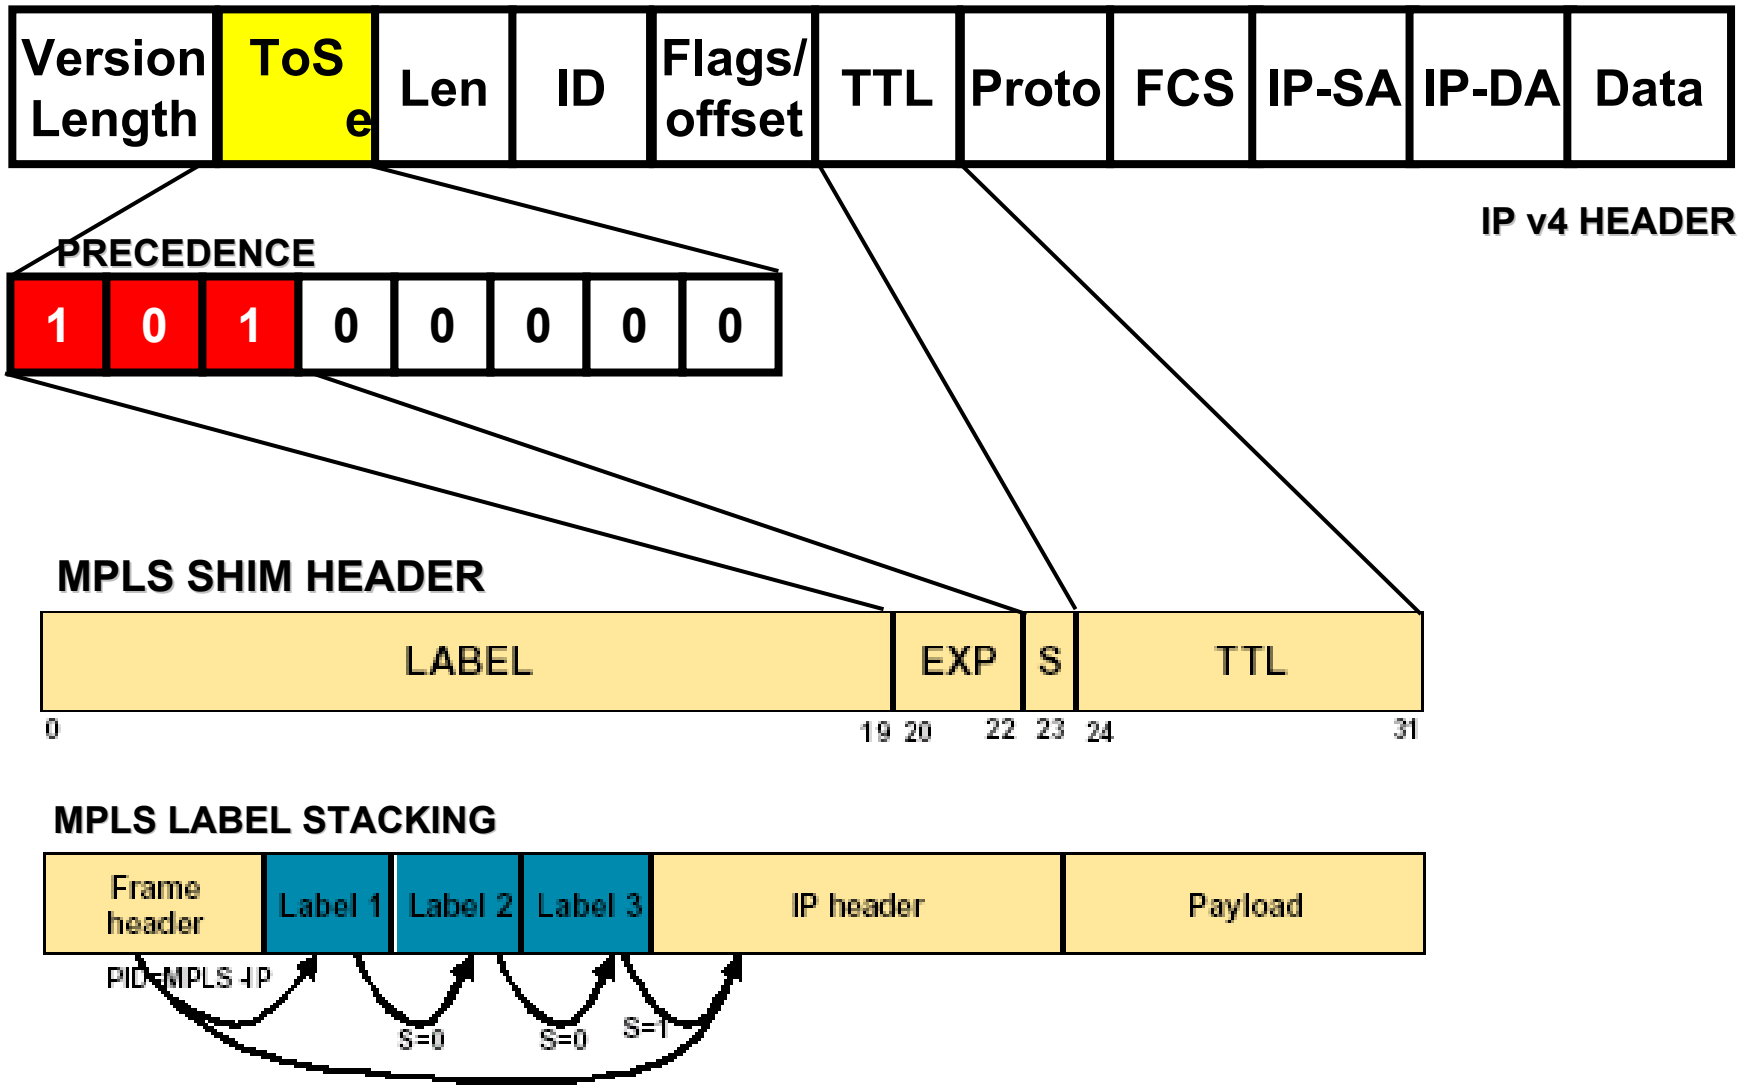

© 2001, Cisco Systems, Inc. All rights reserved.

Label Structure

PPP Header(Packet over SONET/SDH)

Ethernet

Frame Relay

ATM Cell Header

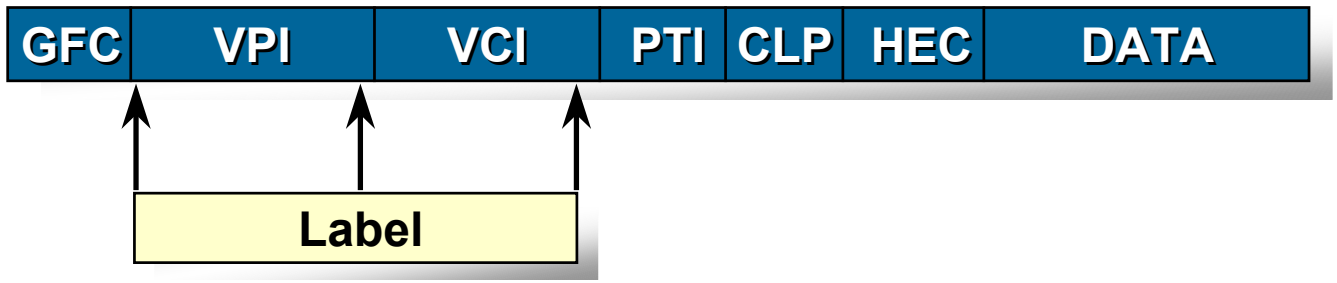

Subsequent cells

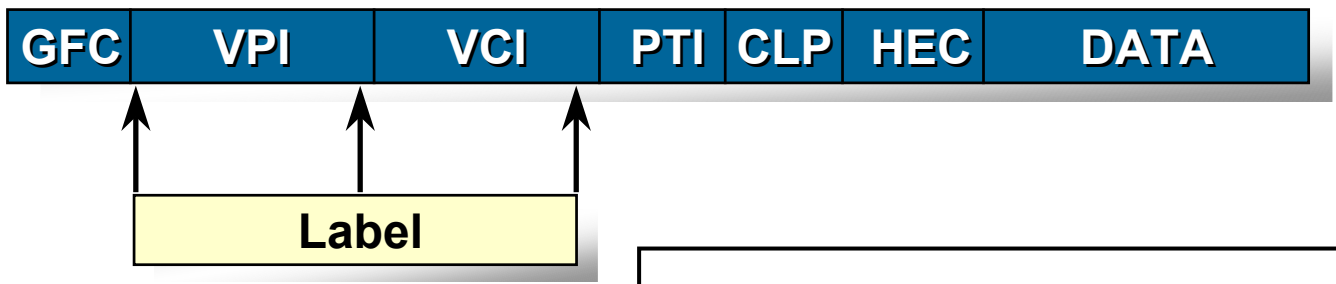

Ether Type 0x8000 for unlabeled IP packet
0x8847 for labeled IP unicast packet
0x8848 for labeled IP multicast packet

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

Label Imposition

- CEF must be enabled on the interface where labels are first ‘imposed’. This is because the CEF mechanisms queries the LFIB in order to find which labels to apply.

- If an interface is purely passing MPLS frames, CEF can be disabled.

- Default label selection is based on the unicast destination IP address

© 2001, Cisco Systems, Inc. All rights reserved.

Label Imposition

Address	Prefix	Local Label
10.1.0.0	/16	1
10.2.0.0	/16	2
20.1.0.0	/16	3
20.2.0.0	/16	4
1.0.0.0	/16	5
2.0.0.0	/16	6
3.0.0.0	/16	7

Address	Prefix	Local Label
10.1.0.0	/16	8
10.2.0.0	/16	9
20.1.0.0	/16	10
20.2.0.0	/16	11
1.0.0.0	/16	12
2.0.0.0	/16	13
3.0.0.0	/16	14

Address	Prefix	Local label
10.1.0.0	/16	15
10.2.0.0	/16	16
20.1.0.0	/16	17
20.2.0.0	/16	18
1.0.0.0	/16	19
2.0.0.0	/16	20
3.0.0.0	/16	21

Address	Prefix	Local Label
10.1.0.0	/16	22
10.2.0.0	/16	23
20.1.0.0	/16	24
20.2.0.0	/16	25
1.0.0.0	/16	26
2.0.0.0	/16	27
3.0.0.0	/16	28

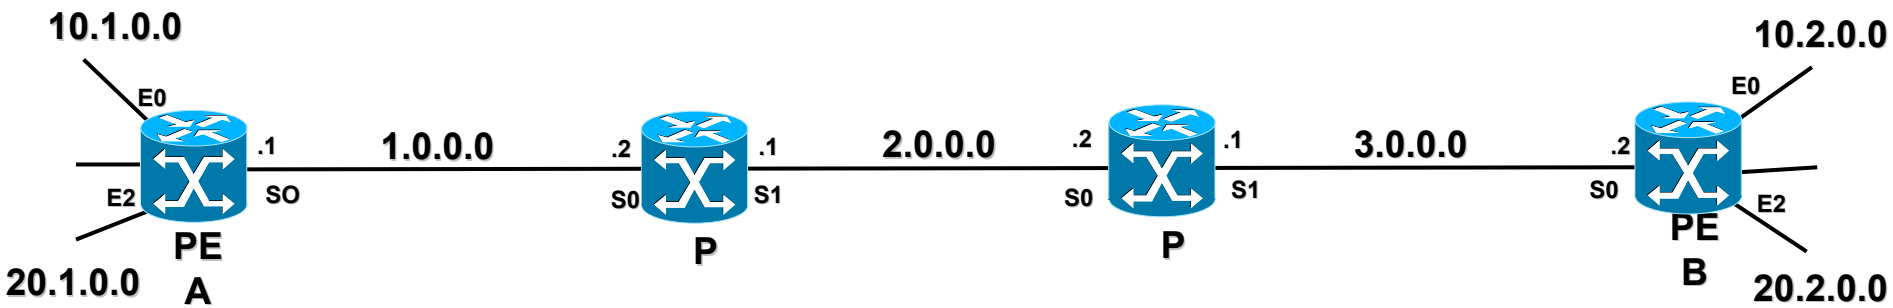

Address	Prefix	Next Hop
10.1.0.0	/16	Connected
10.2.0.0	/16	1.0.0.2
20.1.0.0	/16	Connected
20.2.0.0	/16	1.0.0.2
1.0.0.0	/16	Connected
2.0.0.0	/16	1.0.0.2
3.0.0.0	/16	1.0.0.2

Address	Prefix	Next Hop
10.1.0.0	/16	1.0.0.1
10.2.0.0	/16	2.0.0.2
20.1.0.0	/16	1.0.0.1
20.2.0.0	/16	2.0.0.2
1.0.0.0	/16	Connected
2.0.0.0	/16	Connected
3.0.0.0	/16	2.0.0.2

Address	Prefix	Next Hop
10.1.0.0	/16	2.0.0.1
10.2.0.0	/16	3.0.0.2
20.1.0.0	/16	2.0.0.1
20.2.0.0	/16	3.0.0.2
1.0.0.0	/16	2.0.0.1
2.0.0.0	/16	Connected
3.0.0.0	/16	Connected

Address	Prefix	Next Hop
10.1.0.0	/16	3.0.0.1
10.2.0.0	/16	Connected
20.1.0.0	/16	3.0.0.1
20.2.0.0	/16	Connected
1.0.0.0	/16	3.0.0.1
2.0.0.0	/16	3.0.0.1
3.0.0.0	/16	Connected

© 2001, Cisco Systems, Inc. All rights reserved.

LIB Content

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	1	Null
10.2.0.0	/16	2	2
20.1.0.0	/16	3	Null
20.2.0.0	/16	4	4
1.0.0.0	/16	5	Null
2.0.0.0	/16	6	POP
3.0.0.0	/16	7	7

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	8,15	POP
10.2.0.0	/16	9,2	9
20.1.0.0	/16	10,17	POP
20.2.0.0	/16	11,4	11
1.0.0.0	/16	12,19	Null
2.0.0.0	/16	13,6	Null
3.0.0.0	/16	14,7	POP

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	15,22	15
10.2.0.0	/16	9,16	POP
20.1.0.0	/16	17,24	17
20.2.0.0	/16	11,18	POP
1.0.0.0	/16	19,26	POP
2.0.0.0	/16	20,27	Null
3.0.0.0	/16	14,21	Null

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	22	22
10.2.0.0	/16	23	Null
20.1.0.0	/16	24	24
20.2.0.0	/16	25	Null
1.0.0.0	/16	26	26
2.0.0.0	/16	27	POP
3.0.0.0	/16	28	Null

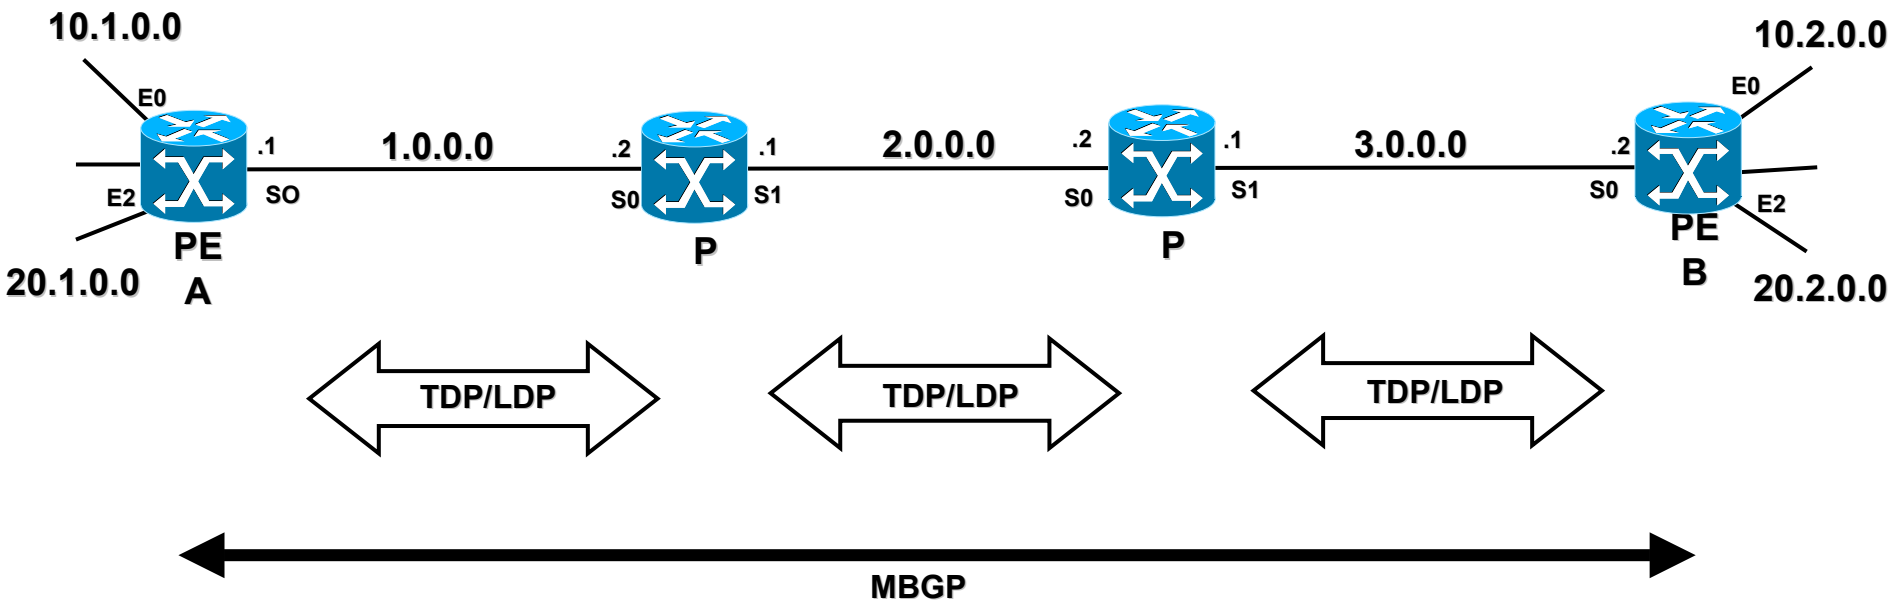

© 2001, Cisco Systems, Inc. All rights reserved.

LFIB Creation

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	1	Null
10.2.0.0	/16	23	2
20.1.0.0	/16	3	Null
20.2.0.0	/16	4	4
1.0.0.0	/16	5	Null
2.0.0.0	/16	6	POP
3.0.0.0	/16	7	7

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	8,15	POP
10.2.0.0	/16	9,2	9
20.1.0.0	/16	10,17	POP
20.2.0.0	/16	11,4	11
1.0.0.0	/16	12,19	Null
2.0.0.0	/16	13,6	Null
3.0.0.0	/16	14,7	POP

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	15,22	15
10.2.0.0	/16	9,16	POP
20.1.0.0	/16	17,24	17
20.2.0.0	/16	11,18	POP
1.0.0.0	/16	19,26	POP
2.0.0.0	/16	20,27	Null
3.0.0.0	/16	14,21	Null

LIB

Address	Prefix	LocLbl	NHLbl
10.1.0.0	/16	22	22
10.2.0.0	/16	23	Null
20.1.0.0	/16	24	24
20.2.0.0	/16	25	Null
1.0.0.0	/16	26	26
2.0.0.0	/16	27	POP
3.0.0.0	/16	28	Null

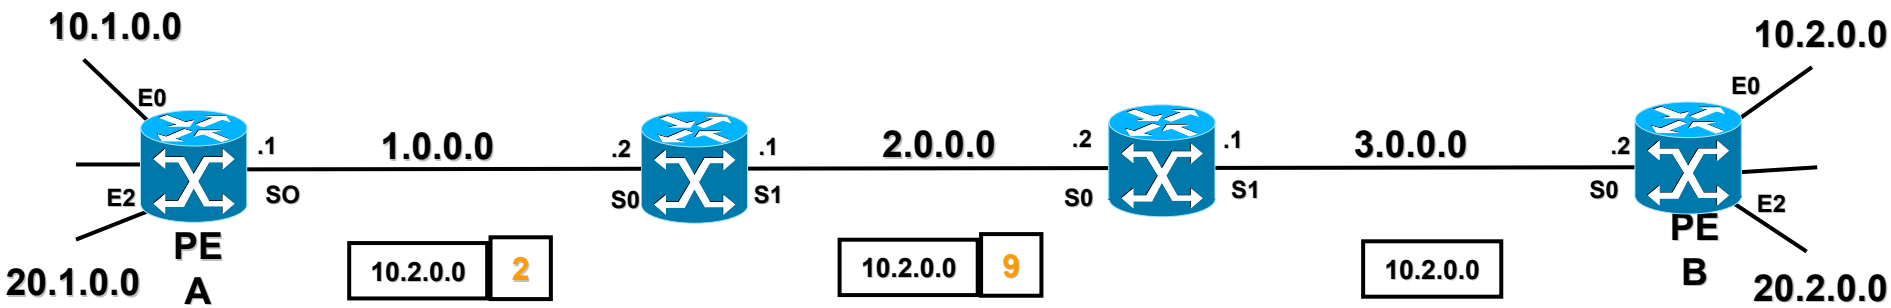

LFIB

Lbl IN	LblOUT	O/IF	MAC Hdr
2	2	S0	ABCD
4	4	S0	ABCD
6	POP	S0	ABCD
7	7	S0	ABCD

LFIB

Lbl IN	LblOUT	O/IF	MAC Hdr
8,10,15,17	POP	S0	AD8F
9,2	9	S1	DCBA
11,4	11	S1	DCBA
14,7	POP	S1	DCBA

LFIB

Lbl IN	LblOUT	O/IF	MAC Hdr
9,16,11,18	POP	S1	1A2B
15,22	15	S0	2B3C
17,24	17	S0	2B3C
19,26	POP	S0	2B3C

LFIB

Lbl IN	LblOUT	O/IF	MAC Hdr
22	22	S0	3C4D
24	24	S0	3C4D
26	26	S0	3C4D
27	POP	S0	3C4D

© 2001, Cisco Systems, Inc. All rights reserved.

Using 2nd Label for VPN service

Address	Prefix	Local Label
10.1.0.0	/16	1
10.2.0.0	/16	2
20.1.0.0	/16	3
20.2.0.0	/16	4
1.0.0.0	/16	5
2.0.0.0	/16	6
3.0.0.0	/16	7

Address	Prefix	Local Label
10.1.0.0	/16	22
10.2.0.0	/16	23
20.1.0.0	/16	24
20.2.0.0	/16	25
1.0.0.0	/16	26
2.0.0.0	/16	27
3.0.0.0	/16	28

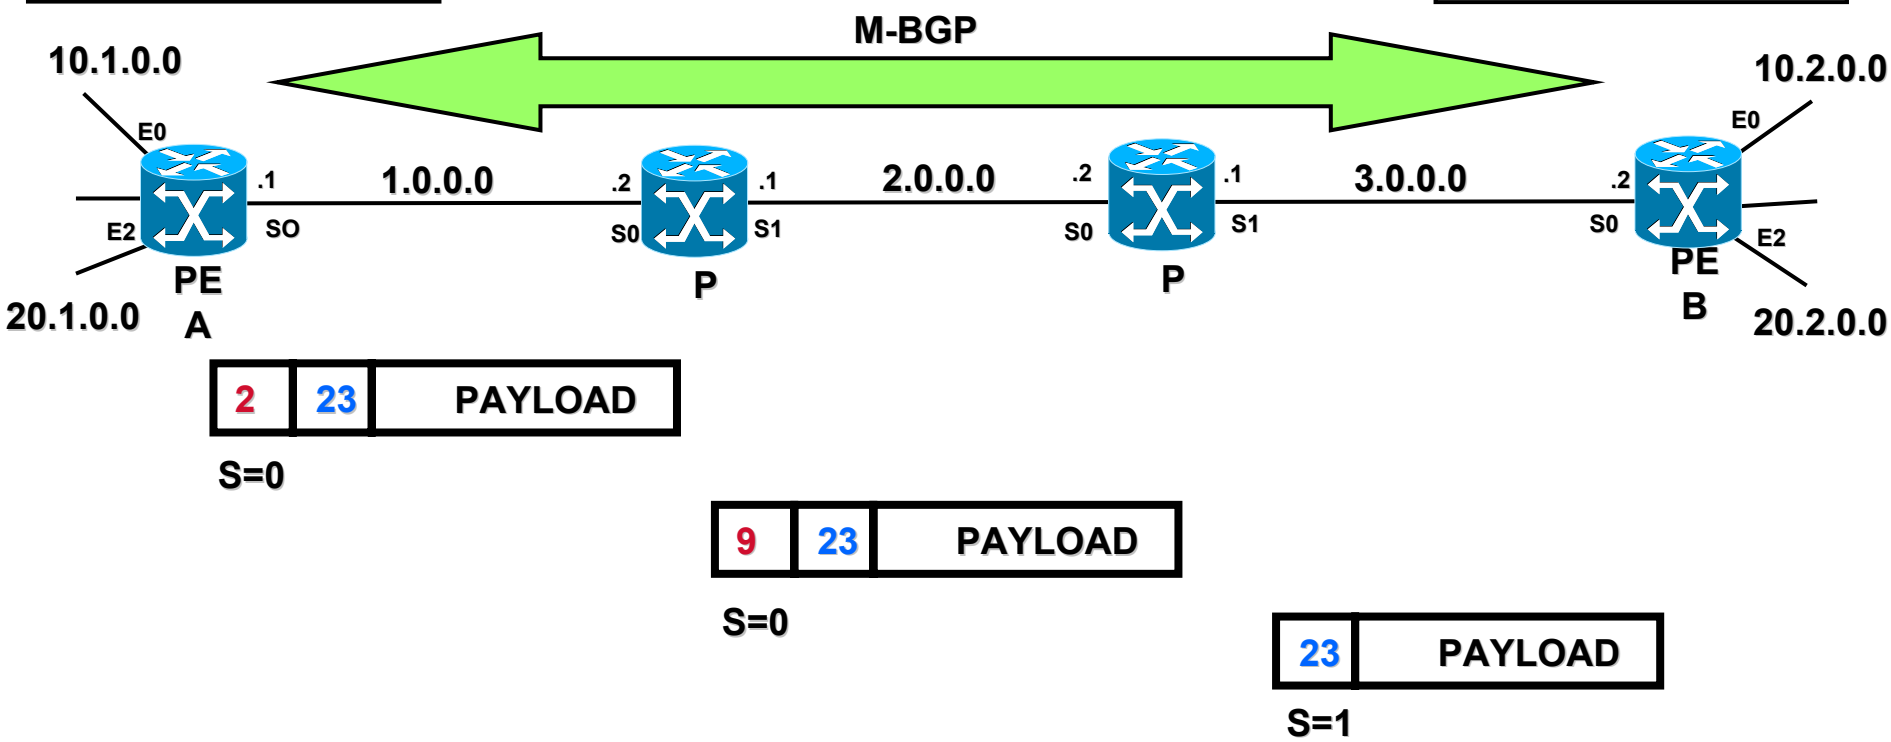

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

ATM LSRs

- ATM switches forward cells, not packets
- Label Dist is Downstream on-demand, Ordered
- IGP label is carried in the VPI/VCI field
- Merging LSR:
  - Ability to use the same label for different FECs if outgoing interface is the same
  - Save label space on ATM-LSRs
  - Cell interleave problem
- Non Merging LSR:
  - ATM-LSR requests one label per FEC and per incoming interface (upstream neighbors)
  - Downstream LSR may request itself new label to its downstream neighbors

© 2001, Cisco Systems, Inc. All rights reserved.

ATM LSRs Non-Merging Downstream on Demand

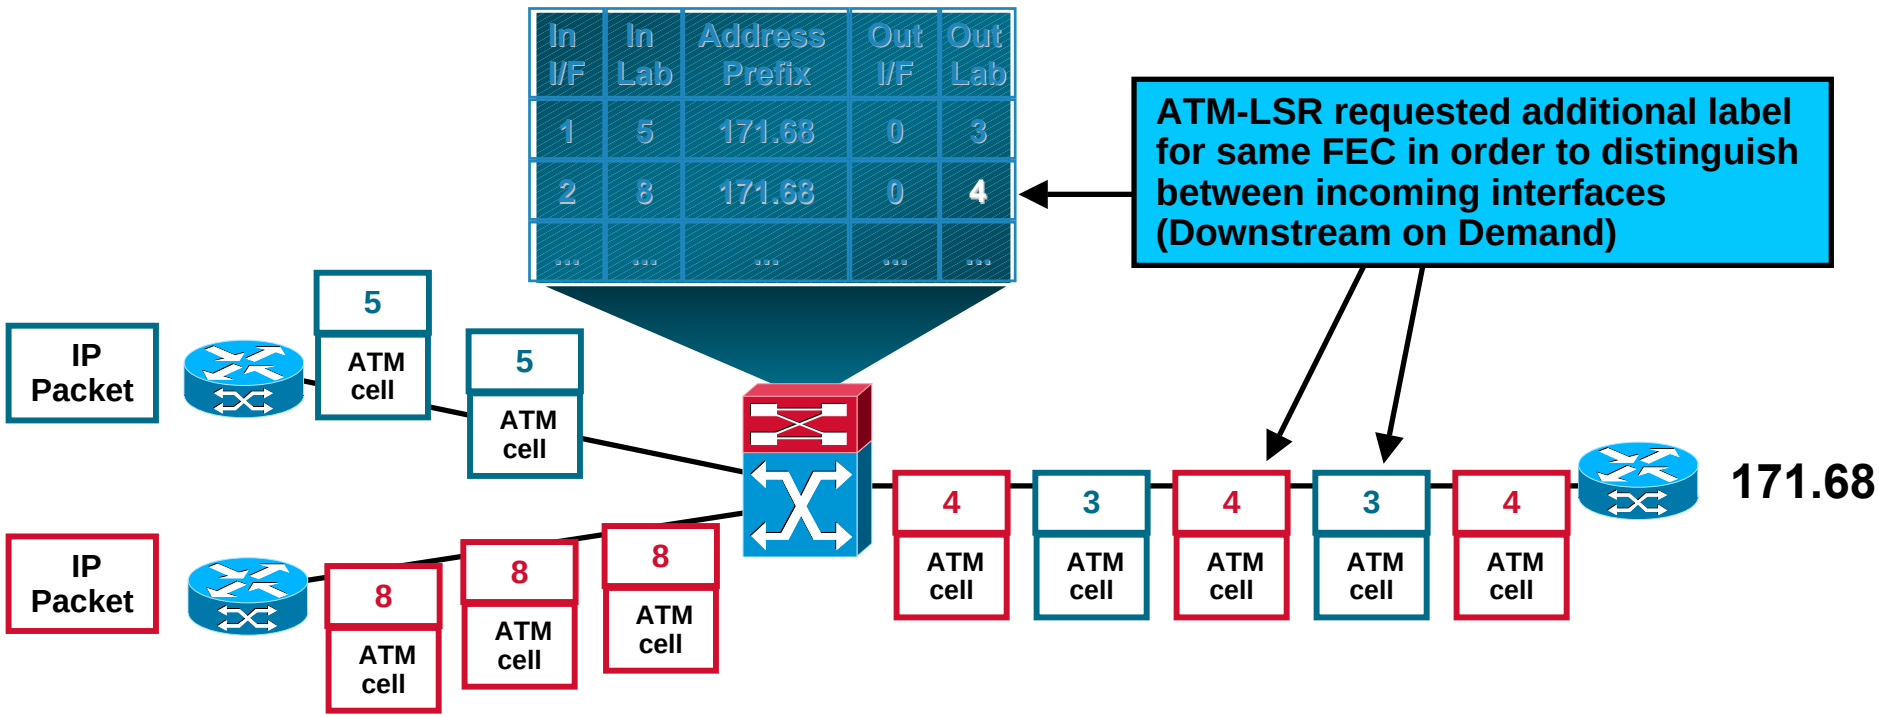

© 2001, Cisco Systems, Inc. All rights reserved.

ATM LSRs VC-Merging Downstream on Demand

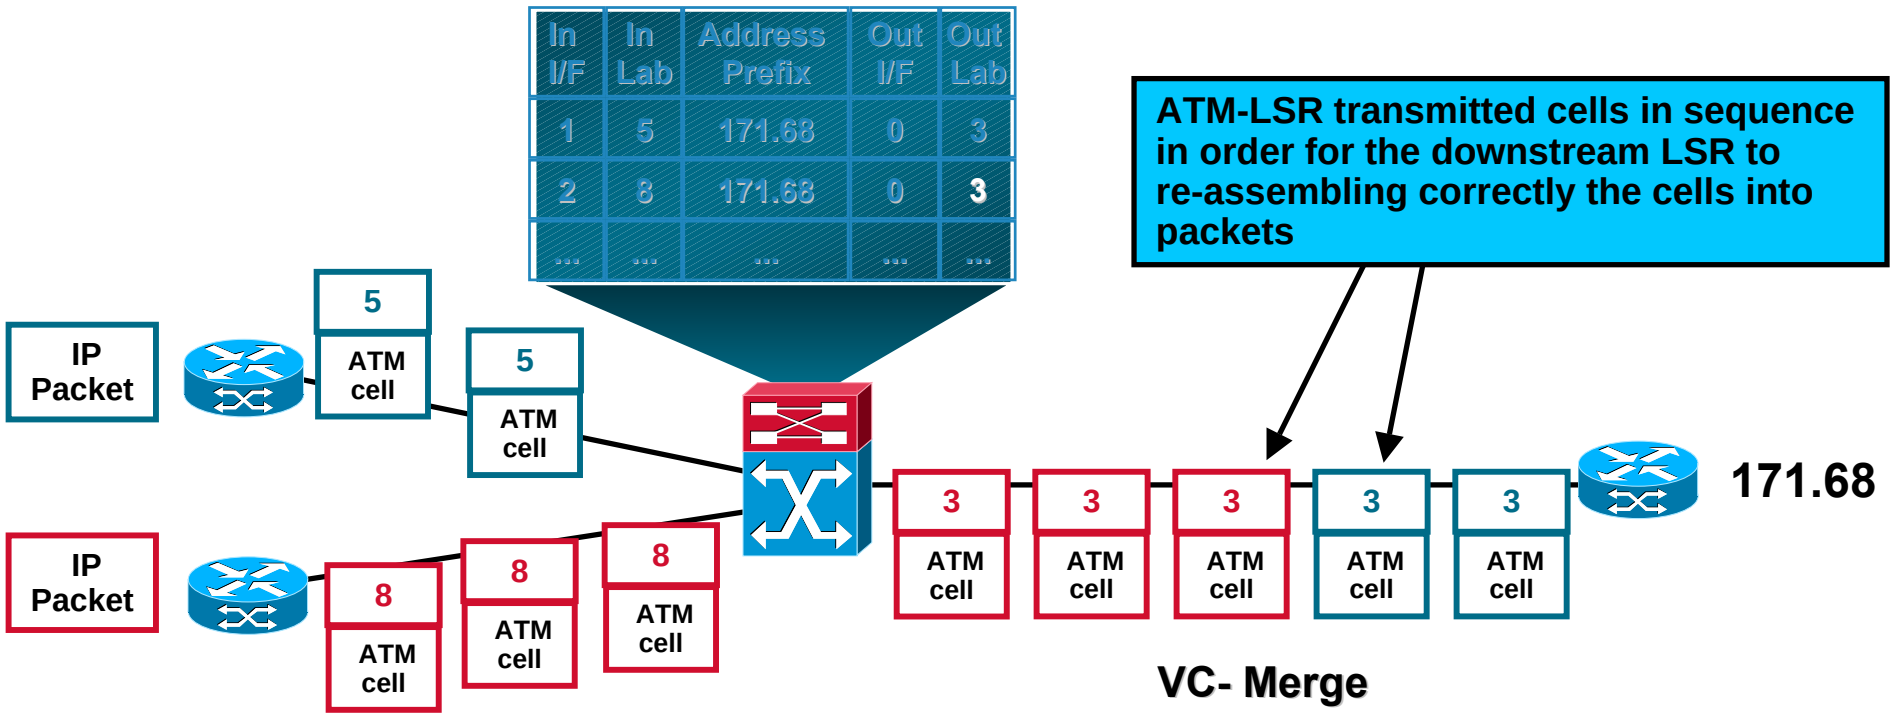

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

Loop prevention and TTL

- In IP networks TTL is used to prevent packets to travel indefinitely in the network
- MPLS ***may*** use same mechanism as IP, but not on all encapsulations
  - TTL is present in the label header for PPP and LAN headers (shim headers)
  - ATM cell header does not have TTL. It relies on IGP, and LDP mechanisms to do the job (Hop Count TLV,Path Vector max # hops), TDP (Hop Count)

draft-ietf-mpls-ldp

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

Route Distinguisher

- New address family: VPN-IP addresses

  VPN-IP address = Route Distinguisher (RD) + IP address

  RDs are assigned by a service provider to each PE

  RDs should be unique AS-wide (by virtue of assignment)

  convert non-unique IP addresses into unique VPN-IP addresses

- Route Distinguisher:

  Used to create a unique IP address, e.g. 2586:10:192.168.23.0

  An RD is needed per VRF signifying the route originator - therefore a single router may have multiple RDs

  The RD is NOT USED for filtering of routing information

© 2001, Cisco Systems, Inc. All rights reserved.

Route Target

- **BGP Extended Community:**

  Used to filter routing information

  Identifies VRFs that may receive set of routes tagged with given Route Target

  Example - 2856:101

  Same format as RDs - NOT the same function!

© 2001, Cisco Systems, Inc. All rights reserved.

MP-iBGP Update

- **VPN-IPV4 address**

  Route Distinguisher (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 (32bits)

- **Extended Community attribute (64 bits)**

  Site of Origin (SOO): identifies the originating site

  Route-target (RT): identifies the set of sites the route has to be advertised to

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Connection Model

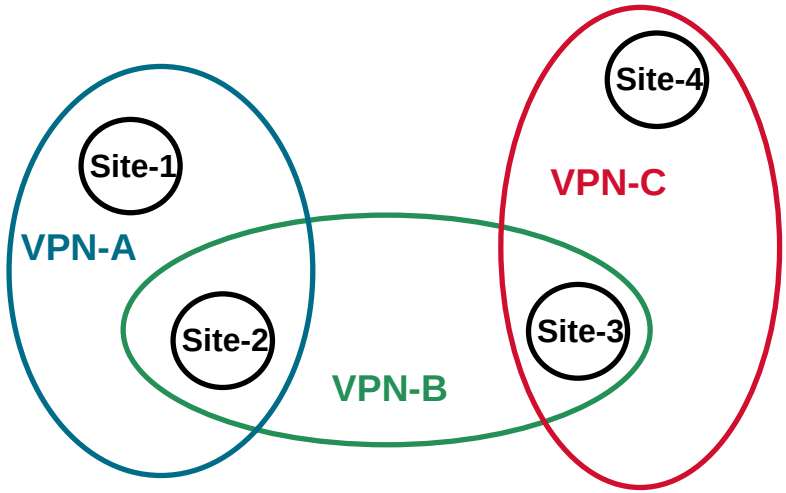

- A site belonging to different VPNs may or MAY NOT be used as a transit point between VPNs
- If two or more VPNs have a common site, address space must be unique among these VPNs

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Connection Model

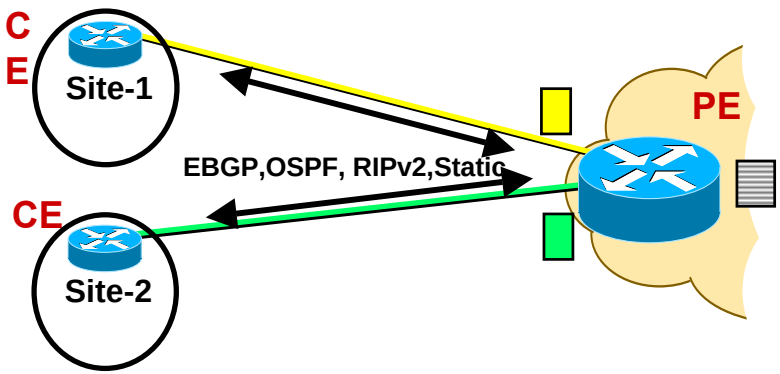

- PE and CE routers exchange routing information through:

  EBGP, OSPF, RIPv2, Static routing

- CE router run standard routing software

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Connection Model

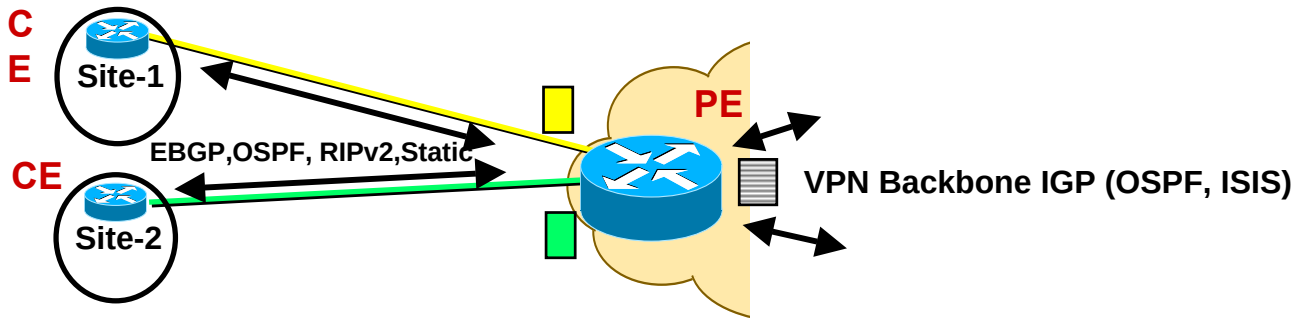

- PE routers maintain separate routing tables
  - The global routing table
    - With all PE and P routes
      Populated by the VPN backbone IGP (ISIS or OSPF)
  - VRF (VPN Routing and Forwarding)
    - Routing and Forwarding table associated with one or more directly connected sites (CEs)
    - VRF are associated to (sub/virtual/tunnel)interfaces
    - Interfaces may share the same VRF if the connected sites may share the same routing information

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Connection Model

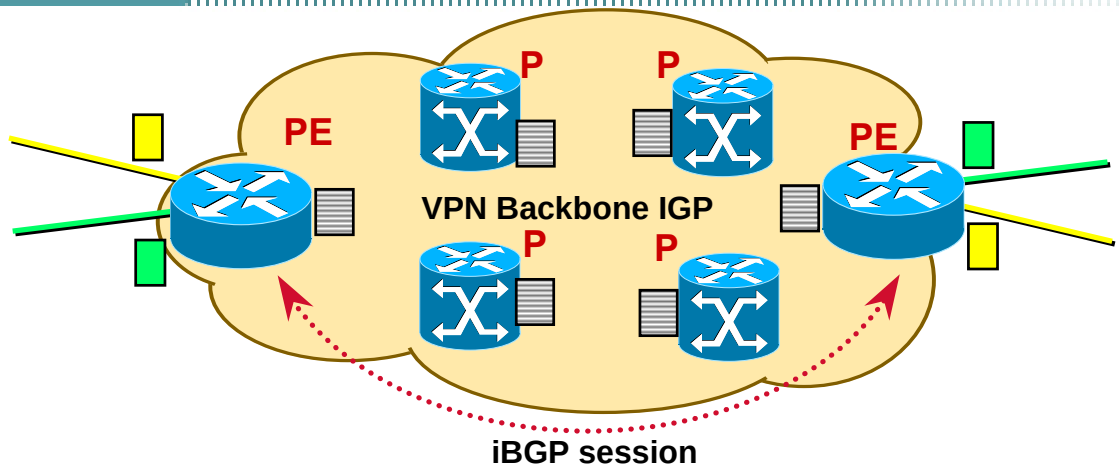

- PE and P routers share a common IGP (ISIS or OSPF)
- PEs establish MP-iBGP sessions between them
- PEs use MP-BGP to exchange routing information related to the connected sites and VPNs

  VPN-IPv4 addresses, Extended Community, Label

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Connection Model

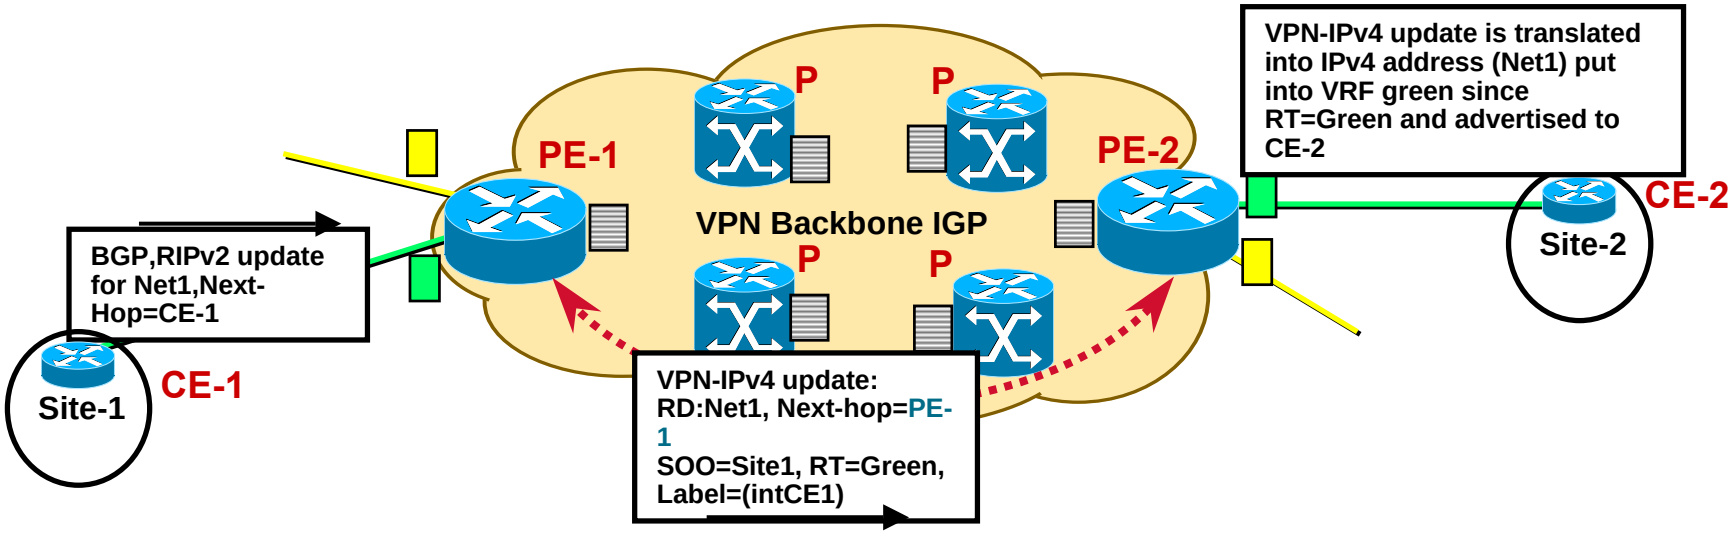

- PE routers receive IPv4 updates (EBGP, RIPv2, Static)
- PE routers translate into VPN-IPv4
  - Assign a SOO and RT based on configuration
  - Re-write Next-Hop attribute
  - Assign a label based on VRF and/or interface
  - Send MP-iBGP update to all PE neighbors

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN mechanisms
VRF and Multiple Routing Instances

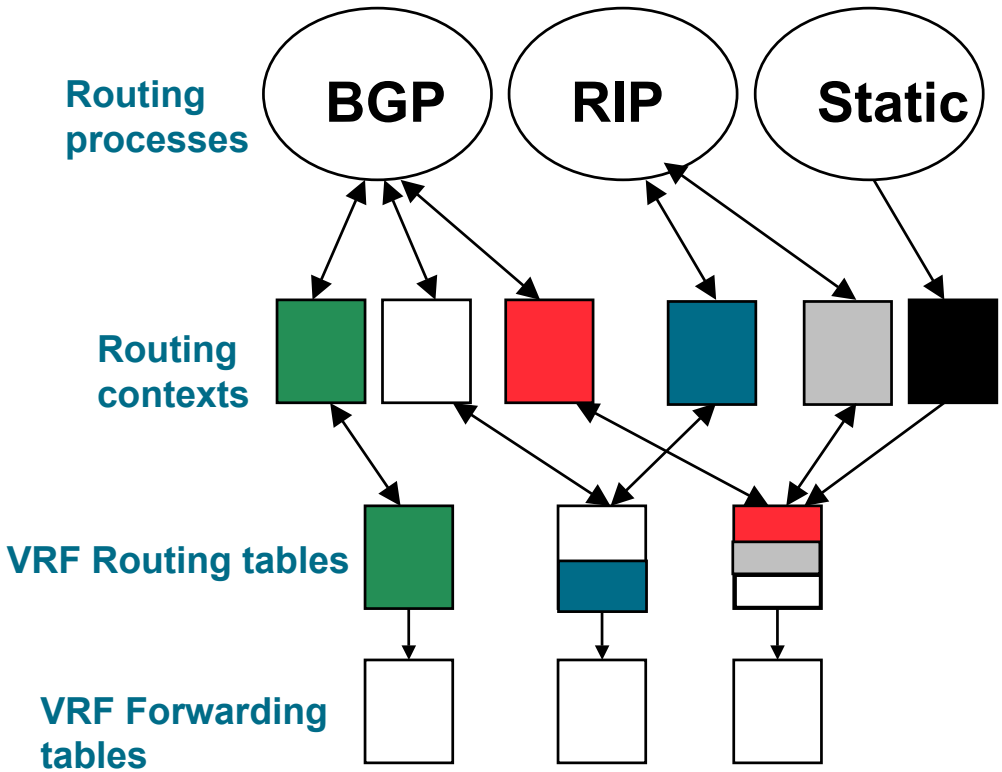

- Routing processes run within specific routing contexts
- Populate specific VPN routing table and FIBs (VRF)
- Interfaces are assigned to VRFs

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN mechanisms
VRF and Multiple Routing Instances

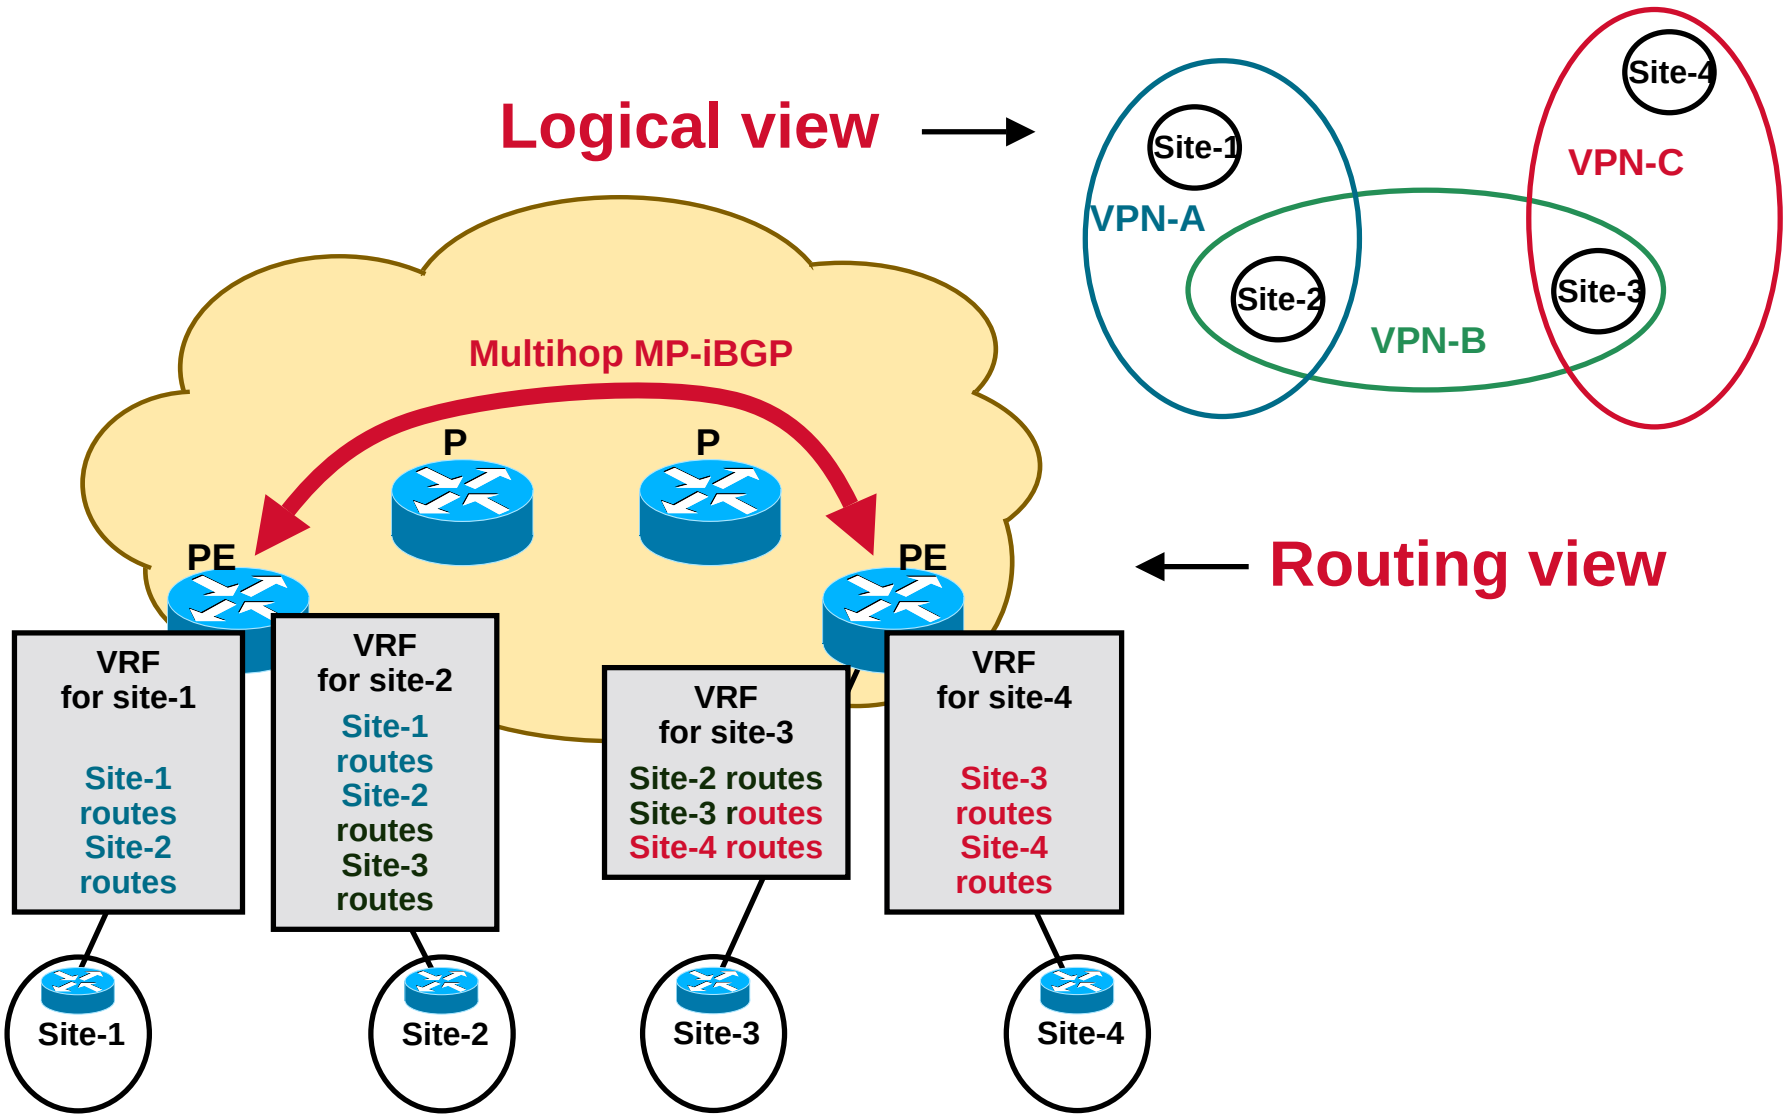

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

VPN sites with Hub & Spoke routing

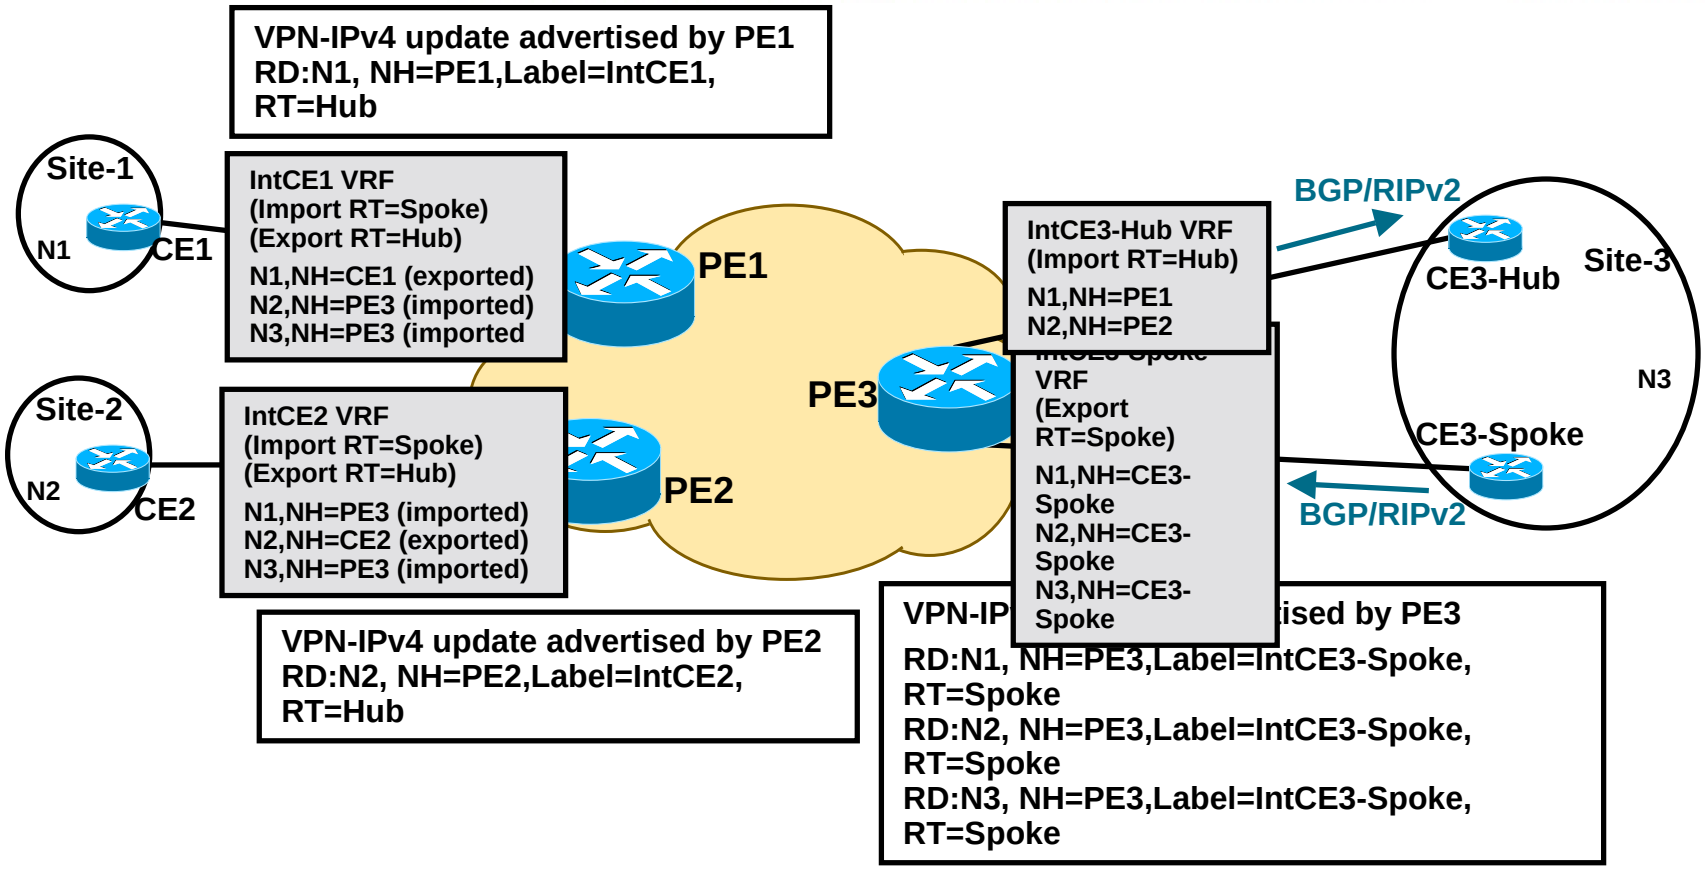

- Routes are imported/exported into VRFs based on RT value of the VPN-IPv4 updates
- PE3 uses 2 (sub) interfaces with two different VRFs

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Topologies
VPN sites with Hub & Spoke routing

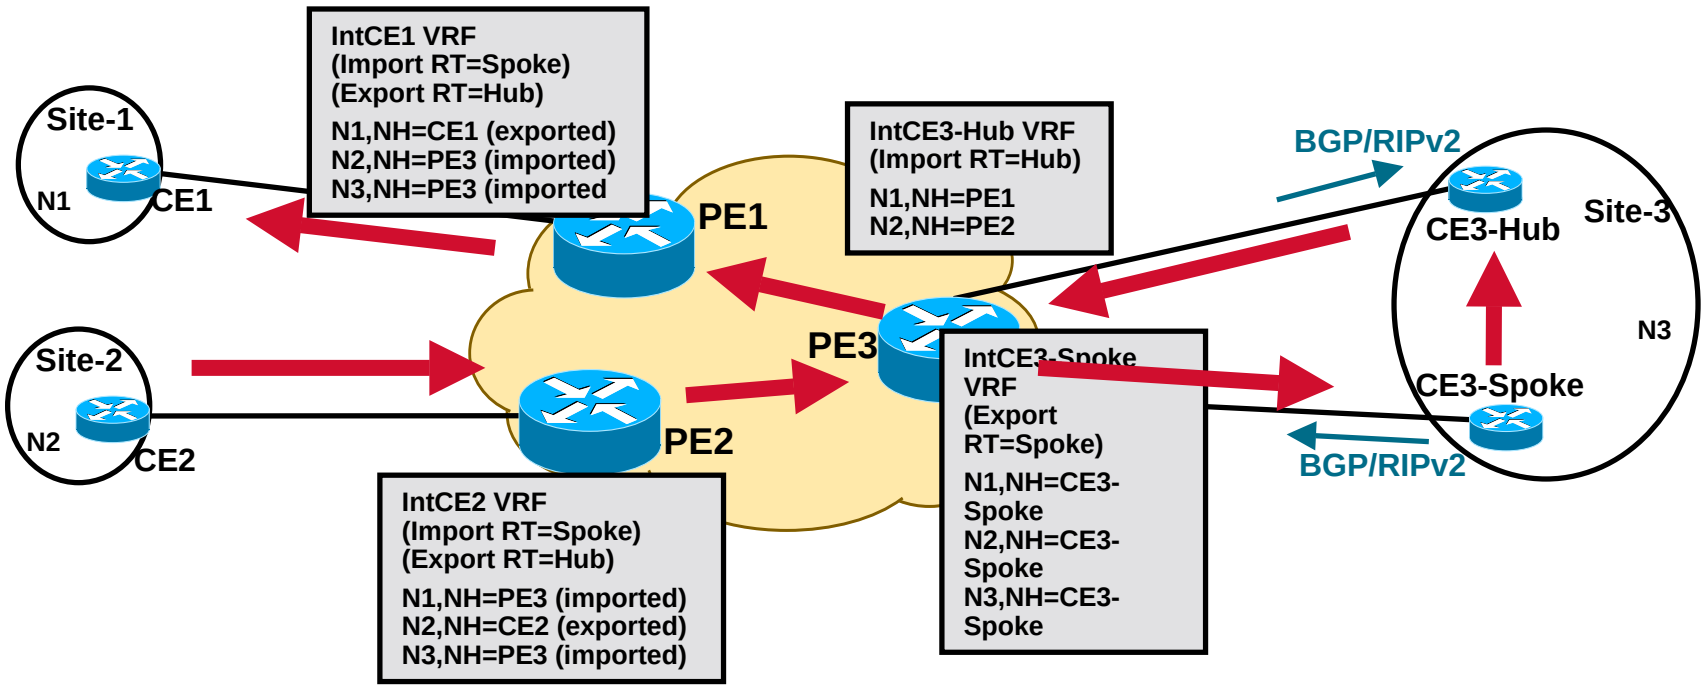

- Traffic from one spoke to another will travel across the hub site
- Hub site may host central services

  Security, NAT, centralised Internet access

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet Routing

- The Internet routing table is treated separately
- In the VPN backbone the Internet routes are in the Global routing table of PE routers
- Labels are not assigned to external (BGP) routes
- P routers need not (and will not) run BGP

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet routing
VRF specific default route

- A default route is installed into the site VRF and pointing to a Internet Gateway
- The default route is NOT part of any VPN

  A single label is used for packets forwarded according to the default route

  The label is the IGP label corresponding to the IP address of the Internet gateway

    Known in the IGP

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet routing
VRF specific default route

- PE router originates CE routes for the Internet

  Customer (site) routes are known in the site VRF

    Not in the global table

  The PE/CE interface is NOT known in the global table. However:

    A static route for customer routes and pointing to the PE/CE interface is installed in the global table

  This static route is redistributed into BGP-4 global table and advertised to the Internet Gateway

- The Internet gateway knows customer routes and with the PE address as next-hop

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet routing
VRF specific default route

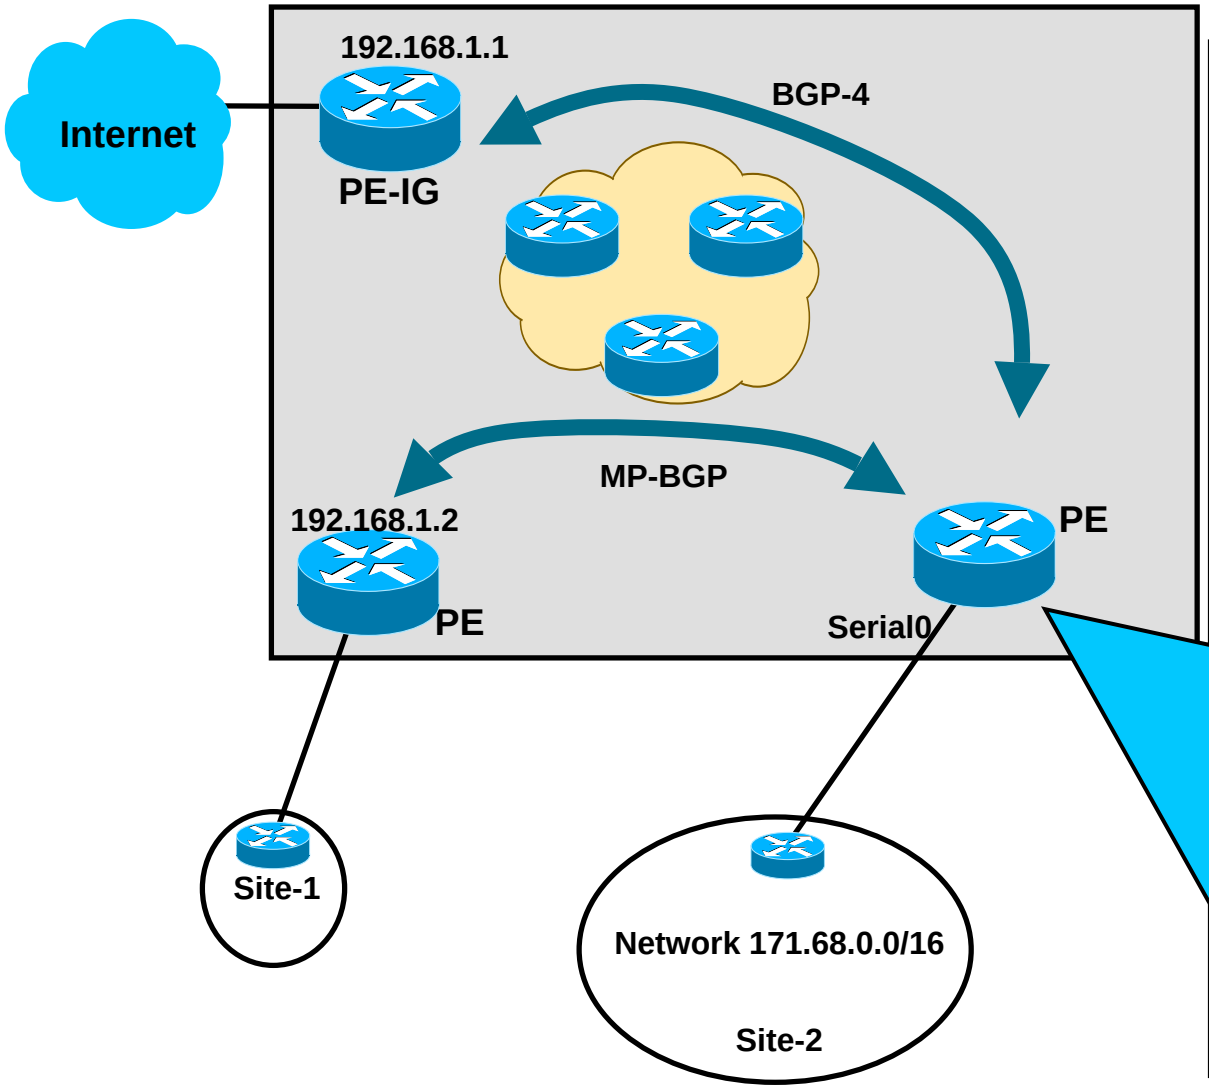


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
!
Interface Serial0
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
!
Router bgp 100
no bgp default ipv4-unicast
network 171.68.0.0 mask 255.255.0.0
neighbor 192.168.1.1 remote 100
neighbor 192.168.1.1 activate
neighbor 192.168.1.1 next-hop-self
neighbor 192.168.1.1 update-source loopback0
!
address-family ipv4 vrf VPN-A
 neighbor 192.168.10.2 remote-as 65502
 neighbor 192.168.10.2 activate
 exit-address-family
!
address-family vpnv4
 neighbor 192.168.1.2  activate
exit-address-family
!
ip route 171.68.0.0 255.255.0.0 Serial0
ip route vrf VPN-A 0.0.0.0 0.0.0.0 192.168.1.1 glob
```

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet Routing
Separated (sub) interfaces

- If CE wishes to receive and announce routes from/to the Internet
  - A dedicated BGP session is used over a separate (sub) interface
  - The PE imports CE routes into the global routing table and advertise them to the Internet
    - The interface is not part of any VPN and does not use any VRF
  - Default route or Internet routes are exported to the CE
    - PE needs to have Internet routing table

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet Routing
Separated (sub) interfaces

- The PE uses separate (sub) interfaces with the CE
  - One (sub) interface for VPN routing
    - associated to a VRF
    - Can be a tunnel interface
  - One (sub) interface for Internet routing
    - Associated to the global routing table

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN Internet Routing
Separated (sub) interfaces

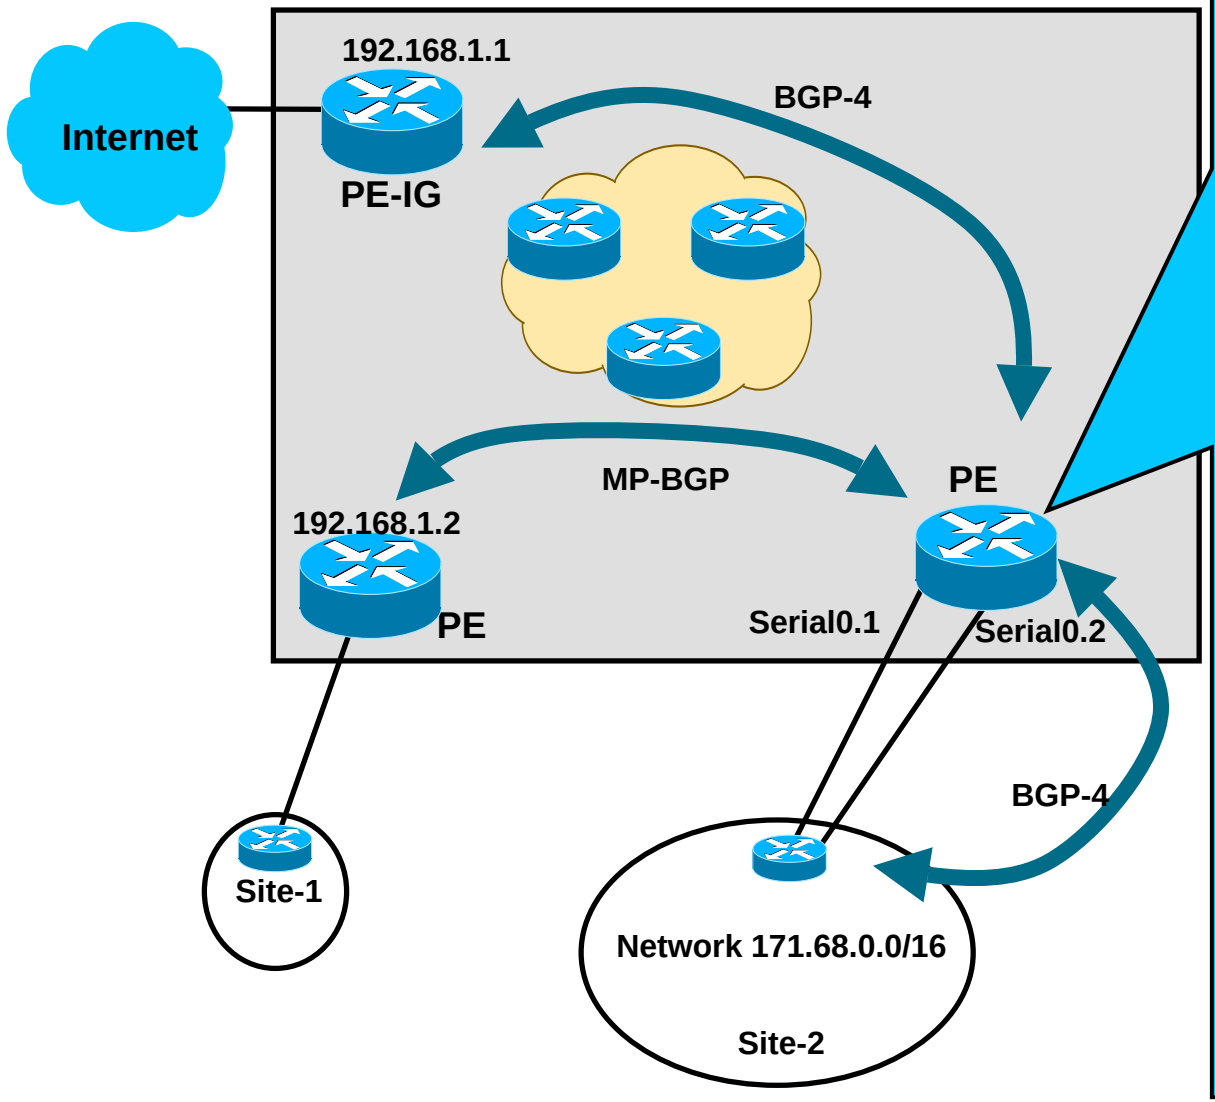


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
!
Interface Serial0
no ip address
!
Interface Serial0.1
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
!
Interface Serial0.2
ip address 171.68.10.1 255.255.255.0
!
Router bgp 100
no bgp default ipv4-unicast
neighbor 192.168.1.1 remote 100
neighbor 192.168.1.1 activate
neighbor 192.168.1.1 next-hop-self
neighbor 192.168.1.1 update-source loopback0
neighbor 171.68.10.2 remote 502
!
address-family ipv4 vrf VPN-A
 neighbor 192.168.10.2 remote-as 502
 neighbor 192.168.10.2 activate
 exit-address-family
!
address-family vpnv4
 neighbor 192.168.1.2  activate
exit-address-family
```

© 2001, Cisco Systems, Inc. All rights reserved.

- MPLS Concepts
- Label Structure
- Label assignment and distribution
- ATM LSRs
- Loop prevention
- RD, RT and VRF instances
- Service Models
- MPLS/VPN Configuration

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration

- VPN knowledge is on PE routers
- PE router have to be configured for
  - VRF and Route Distinguisher
  - VRF import/export policies (based on Route-target)
  - Routing protocol used with CEs
  - MP-BGP between PE routers
  - BGP for Internet routers
    - With other PE routers
    - With CE routers

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
VRF and Route Distinguisher

- RD is configured on PE routers (for each VRF)
- VRFs are associated to RDs in each PE
- Common (good) practice is to use the same RD for the same VPN in all PEs

  But not mandatory
- VRF configuration command

```
ip vrf <vrf-symbolic-name>
 rd <route-distinguisher-value>
 route-target import <community>
 route-target export <community>
```

© 2001, Cisco Systems, Inc. All rights reserved.

CLI - VRF configuration


```
ip vrf site1
    rd 100:1
    route-target export
100:1
    route-target import
100:1
ip vrf site2
    rd 100:2
    route-target export
100:2
    route-target import
100:2
    route-target import
100:1
    route-target export
100:1
```

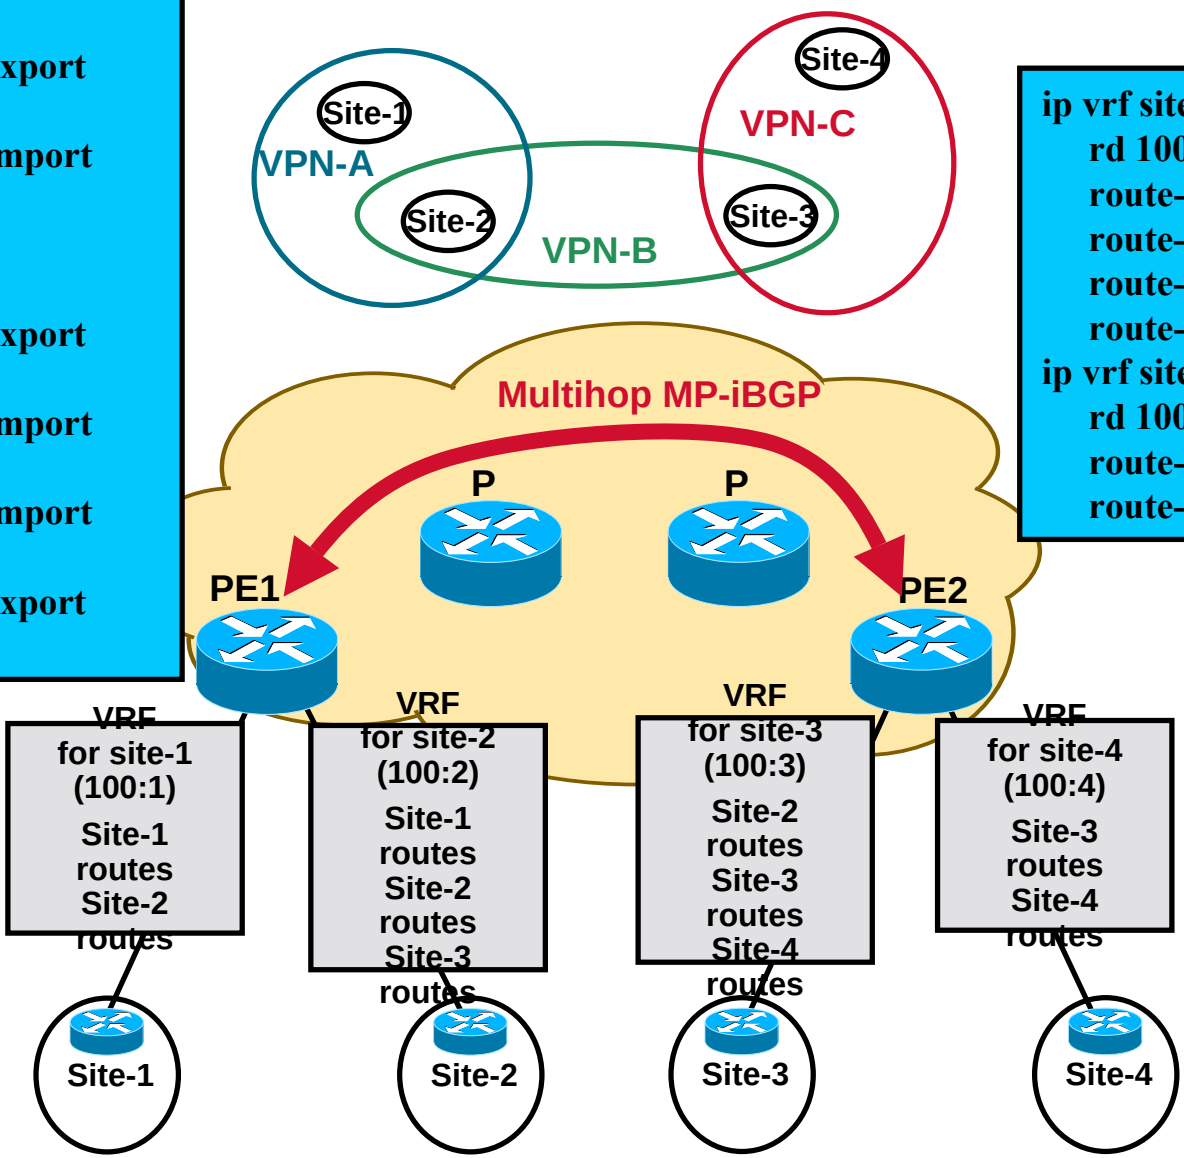


```
ip vrf site3
    rd 100:3
    route-target export 100:2
    route-target import 100:2
    route-target import 100:3
    route-target export  100:3
ip vrf site-4
    rd 100:4
    route-target export 100:3
    route-target import 100:3
```

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
PE/CE routing protocols

- PE/CE may use BGP, RIPv2 or Static routes
- A routing context is used for each VRF
- Routing contexts are defined within the routing protocol instance

Address-family router sub-command

```
Router rip
  version 2
  address-family ipv4 vrf <vrf-symbolic-name>
  …
    any common router sub-command
  …
```

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
PE/CE routing protocols

- BGP uses same “address-family” command

  Router BGP <asn>

  …

  address-family ipv4 vrf <vrf-symbolic-name>

  …

  any common router BGP sub-command

  …

- Static routes are configured per VRF

  ip route vrf <vrf-symbolic-name> …

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
PE router commands

- All show commands are VRF based

  Show ip route vrf <vrf-symbolic-name> ...

  Show ip protocol vrf <vrf-symbolic-name>

  Show ip cef <vrf-symbolic-name> …

  …

- PING and Telnet commands are VRF based

  telnet /vrf <vrf-symbolic-name>

  ping vrf <vrf-symbolic-name>

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
PE/CE routing protocols

```
ip vrf site1
    rd 100:1
    route-target export 100:12
    route-target import 100:12
ip vrf site2
    rd 100:2
    route-target export 100:12
    route-target import 100:12
    route-target import 100:23
    route-target export 100:23
!
interface Serial3/6
 ip vrf forwarding site1
 ip address 192.168.61.6
255.255.255.0
 encapsulation ppp
!
interface Serial3/7
 ip vrf forwarding site2
 ip address 192.168.62.6
255.255.255.0
 encapsulation ppp
```

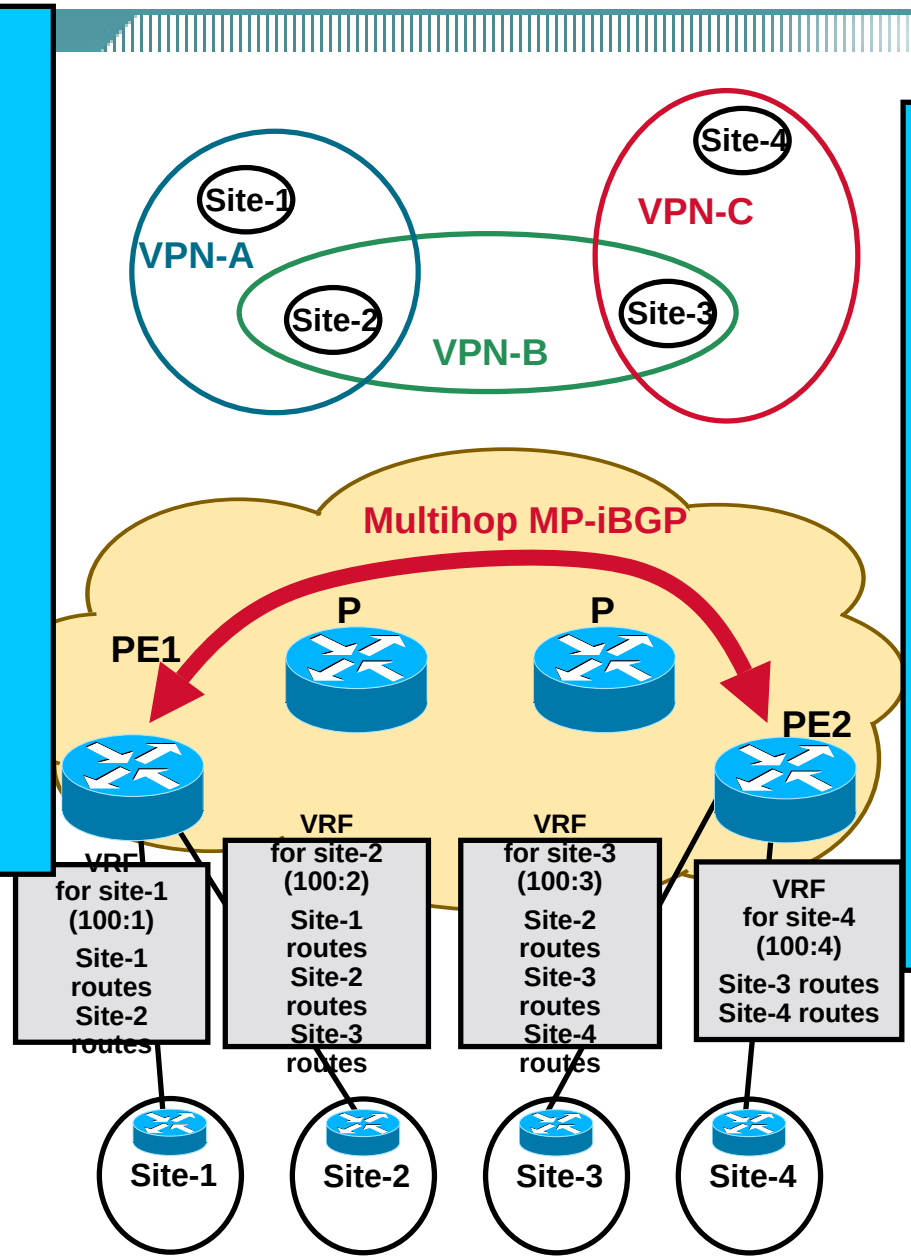


```
ip vrf site3
    rd 100:3
    route-target export 100:23
    route-target import 100:23
    route-target import 100:34
    route-target export 100:34
ip vrf site-4
    rd 100:4
    route-target export 100:34
    route-target import 100:34
!
interface Serial4/6
 ip vrf forwarding site3
 ip address 192.168.73.7
255.255.255.0
 encapsulation ppp
!
interface Serial4/7
 ip vrf forwarding site4
 ip address 192.168.74.7
255.255.255.0
 encapsulation ppp
```

© 2001, Cisco Systems, Inc. All rights reserved.

MPLS VPN - Configuration
PE/CE routing protocols

```
router bgp 100
no bgp default ipv4-unicast
 neighbor 7.7.7.7 remote-as 100
 neighbor 7.7.7.7 update-source
Loop0
!
 address-family ipv4 vrf site2
   neighbor 192.168.62.2 remote-as
                  65502
 neighbor 192.168.62.2 activate
 exit-address-family
 !
 address-family ipv4 vrf site1
   neighbor 192.168.61.1 remote-as
                  65501
 neighbor 192.168.61.1 activate
 exit-address-family
 !
 address-family vpnv4
 neighbor 7.7.7.7 activate
 neighbor 7.7.7.7 next-hop-self
exit-address-family
```

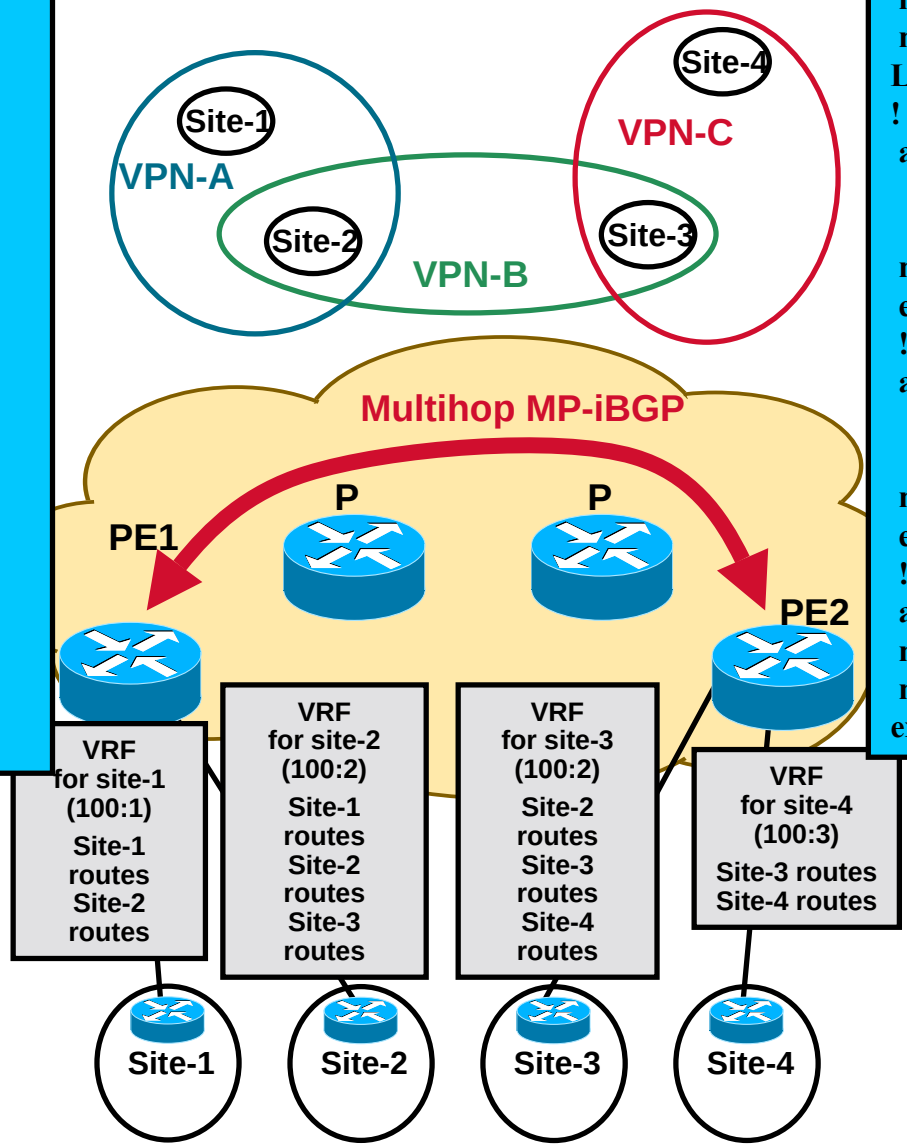


```
router bgp 100
no bgp default ipv4-unicast
 neighbor 6.6.6.6 remote-as 100
 neighbor 6.6.6.6 update-source
Loop0
!
 address-family ipv4 vrf site4
    neighbor 192.168.74.4 remote-as
                  65504
 neighbor 192.168.74.4 activate
 exit-address-family
 !
 address-family ipv4 vrf site3
    neighbor 192.168.73.3 remote-as
                  65503
 neighbor 192.168.73.3 activate
 exit-address-family
 !
 address-family vpnv4
 neighbor 6.6.6.6 activate
 neighbor 6.6.6.6 next-hop-self
exit-address-family
```

© 2001, Cisco Systems, Inc. All rights reserved.

Allow AS

- If PE and Hub-site use BGP the PE should not check the received AS_PATH

  The update the Hub-site advertise contains the VPN backbone AS number

  By configuration the AS_PATH check is disabled

  Allow AS

  Routing loops are suppressed by the limit of occurrence of provider ASN in the AS_PATH

  Therefore, PE will **REJECT** the update if its ASN appears more than ***3 times*** in the AS_PATH 3 is the default and can be overwritten with <opt>

© 2001, Cisco Systems, Inc. All rights reserved.

Allow AS

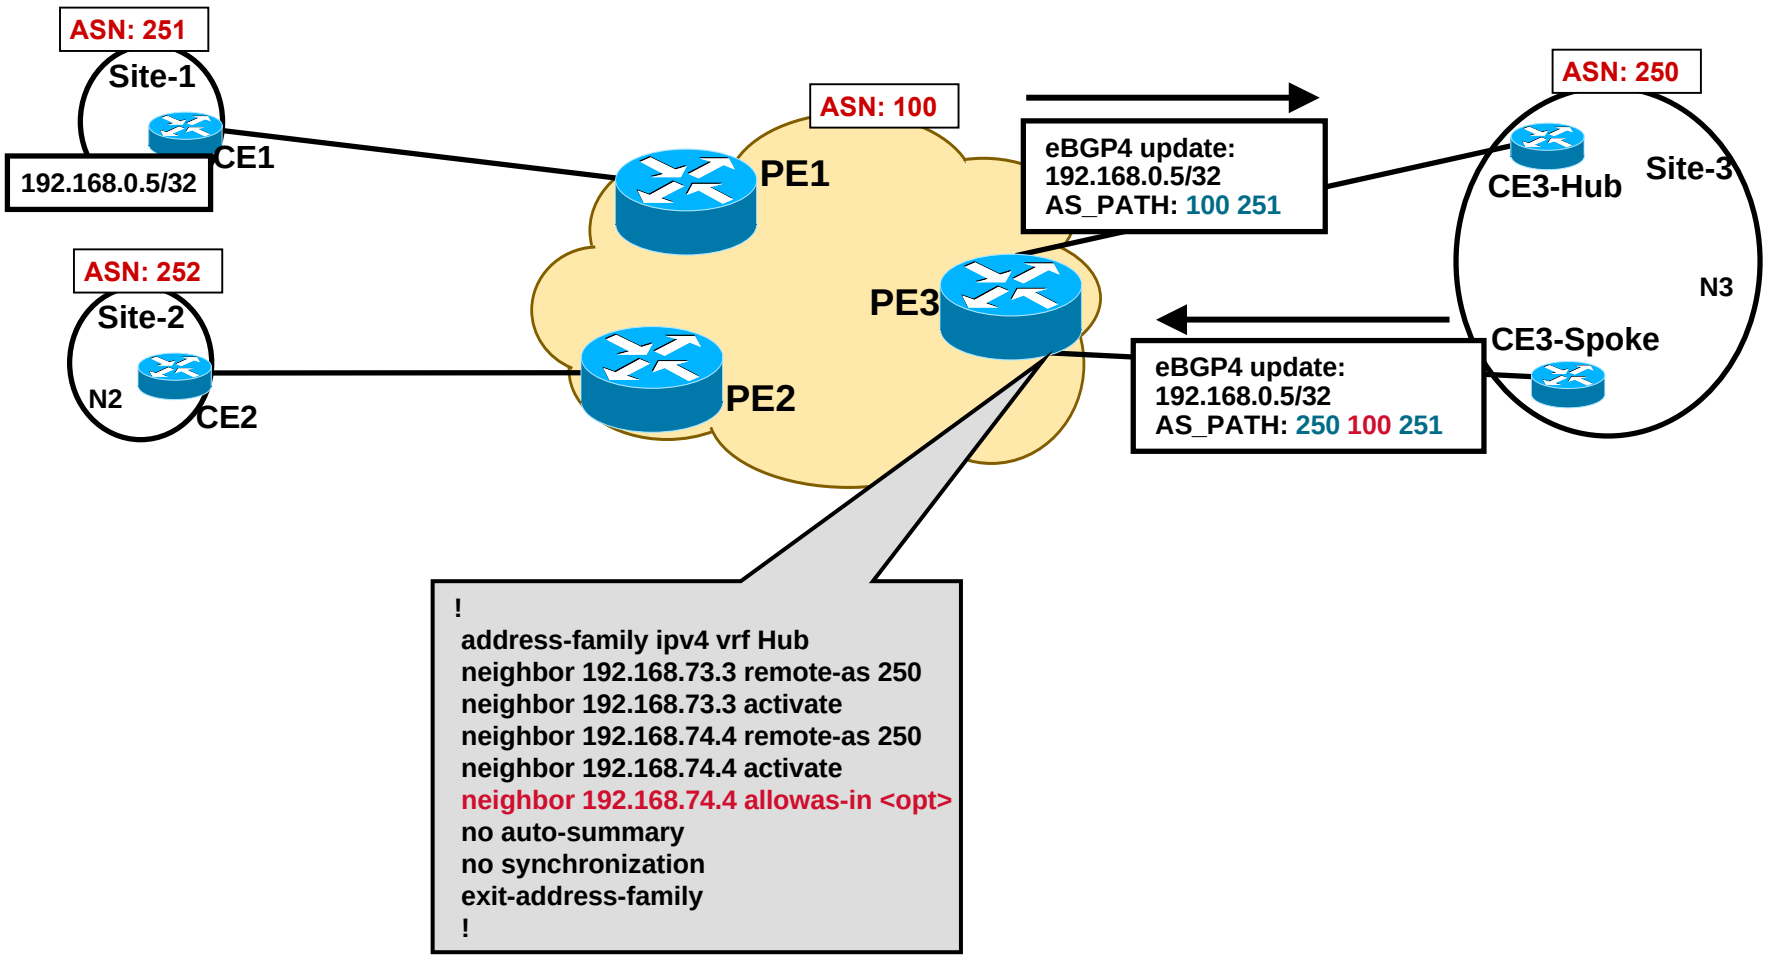

© 2001, Cisco Systems, Inc. All rights reserved.

Allow AS with ASN override

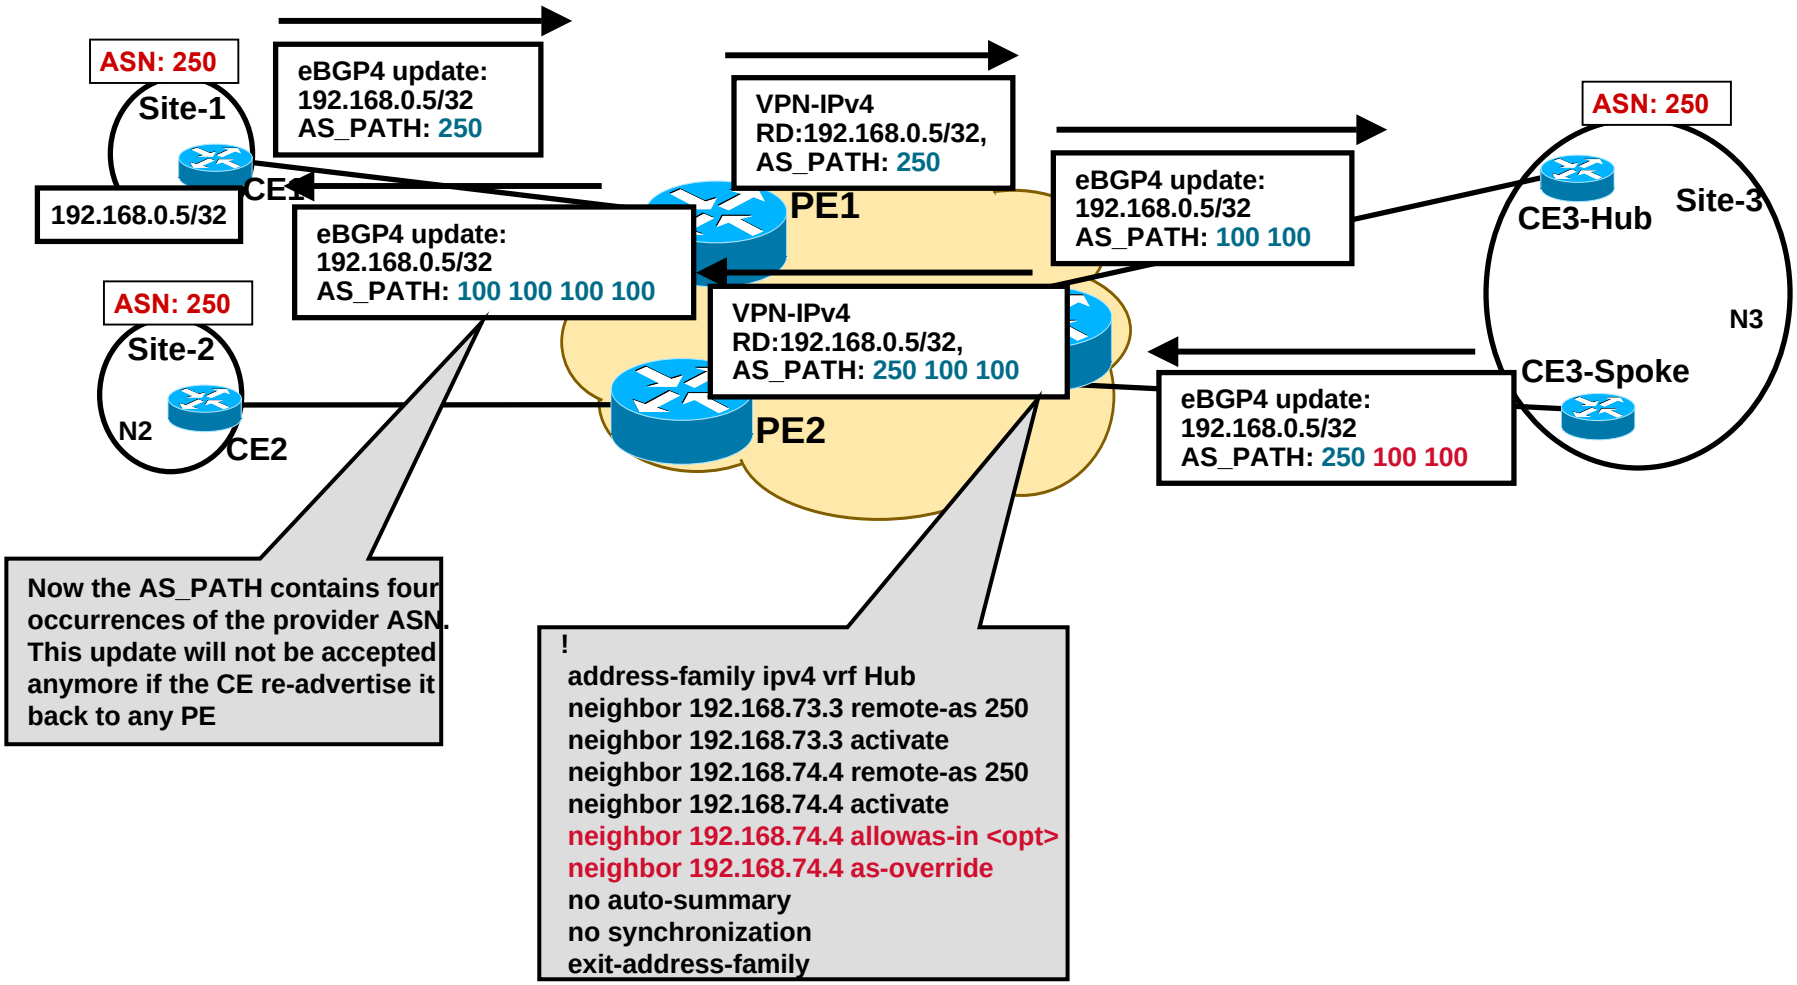

© 2001, Cisco Systems, Inc. All rights reserved.

ASN Override

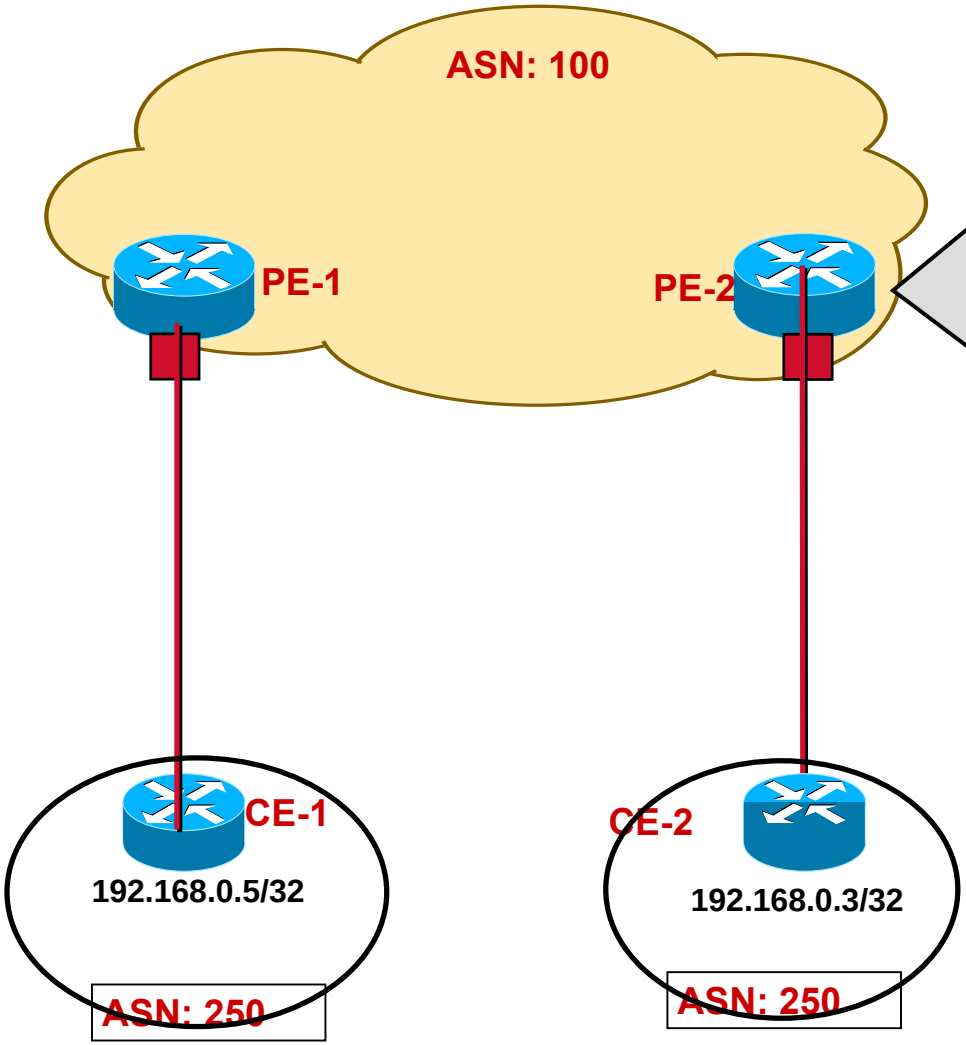


```
ip vrf odd
 rd 100:1
 route-target export 100:3
 route-target import 100:3
!
interface Serial1
 ip vrf forwarding odd
 ip address 192.168.73.7 255.255.255.0
!
router bgp 100
 no synchronization
 no bgp default ipv4-unicast
 neighbor 192.168.0.6 remote-as 100
 neighbor 192.168.0.6 update-source Loop0
 neighbor 192.168.0.6 activate
 neighbor 192.168.0.6 next-hop-self
 no auto-summary
 !
 address-family ipv4 vrf odd
 neighbor 192.168.73.3 remote-as 250
 neighbor 192.168.73.3 activate
 neighbor 192.168.73.3 as-override
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 192.168.0.6 activate
 neighbor 192.168.0.6 send-community extended
 no auto-summary
 exit-address-family
!
```

© 2001, Cisco Systems, Inc. All rights reserved.

ASN Override

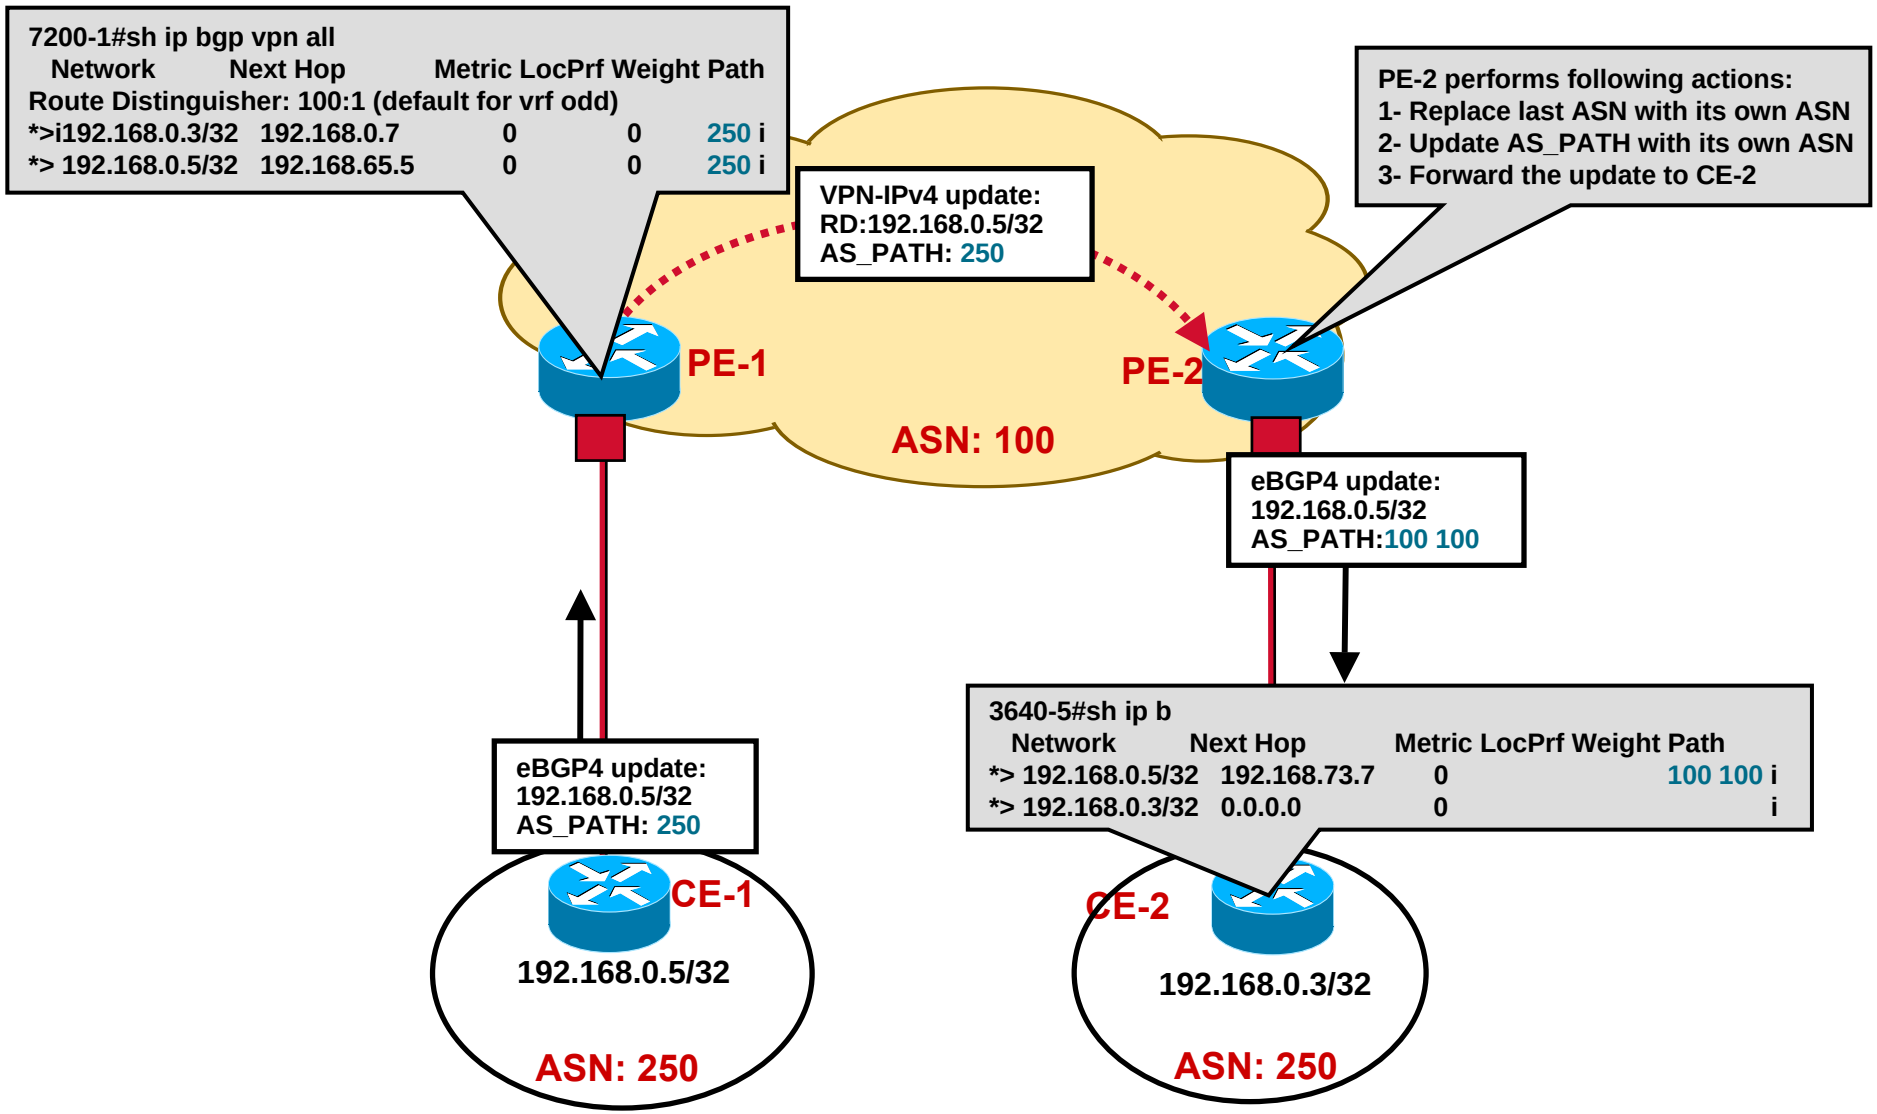

© 2001, Cisco Systems, Inc. All rights reserved.

Site of Origin

- Used to identify the site
- Extended Community type
- Used to prevent loops when AS_PATH cannot be used
  - When BGP is used between PE and multihomed sites
  - A BGP route is NOT advertised back to the same site
    - Even through different PE/CE connections

© 2001, Cisco Systems, Inc. All rights reserved.

Site of Origin

- SOO for eBGP learned routes

  SOO is configured through a route-map command

- SOO can be applied to routes learned through a particular VRF interface (without the use of BGP between PE and CE)

  SOO is then configured on the interface

  SOO is propagated into BGP during redistribution

© 2001, Cisco Systems, Inc. All rights reserved.

Site of Origin


```
7200-1#sh ip route vrf odd
C    192.168.65.0/24 is directly connected, Serial2
B       192.168.0.5 [20/0] via 192.168.65.5, 00:08:44, Serial2
7200-1#
7200-1#sh ip bgp vpn all
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf odd)
*> 192.168.0.5/32    192.168.65.5             0             0       250 i
7200-1#sh ip bgp vpn all 192.168.0.5
BGP routing table entry for 100:1:192.168.0.5/32, version 17
Paths: (1 available, best #1)
  Advertised to non peer-group peers:
  192.168.0.7
  250
    192.168.65.5 from 192.168.65.5 (192.168.0.5)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Extended Community: SoO:100:65 RT:100:3
7200-1#
```

```
ip vrf odd
 rd 100:1
 route-target export 100:3
 route-target import 100:3
!
interface Serial1
 ip vrf forwarding odd
 ip address 192.168.65.6 255.255.255.0
!
router bgp 100
 no synchronization
 no bgp default ipv4-unicast
 neighbor 192.168.0.7 remote-as 100
 neighbor 192.168.0.7 update-source Loop0
 neighbor 192.168.0.7 activate
 neighbor 192.168.0.7 next-hop-self
 no auto-summary
 !
 address-family ipv4 vrf odd
 neighbor 192.168.65.5 remote-as 250
 neighbor 192.168.65.5 activate
 neighbor 192.168.65.5 route-map setsoo in
 no auto-summary
 no synchronization
 exit-address-family
 !
 address-family vpnv4
 neighbor 192.168.0.7 activate
 neighbor 192.168.0.7 send-community extended
 no auto-summary
 exit-address-family
!
route-map setsoo permit 10
 set extcommunity soo 100:65
```

© 2001, Cisco Systems, Inc. All rights reserved.

Site of Origin

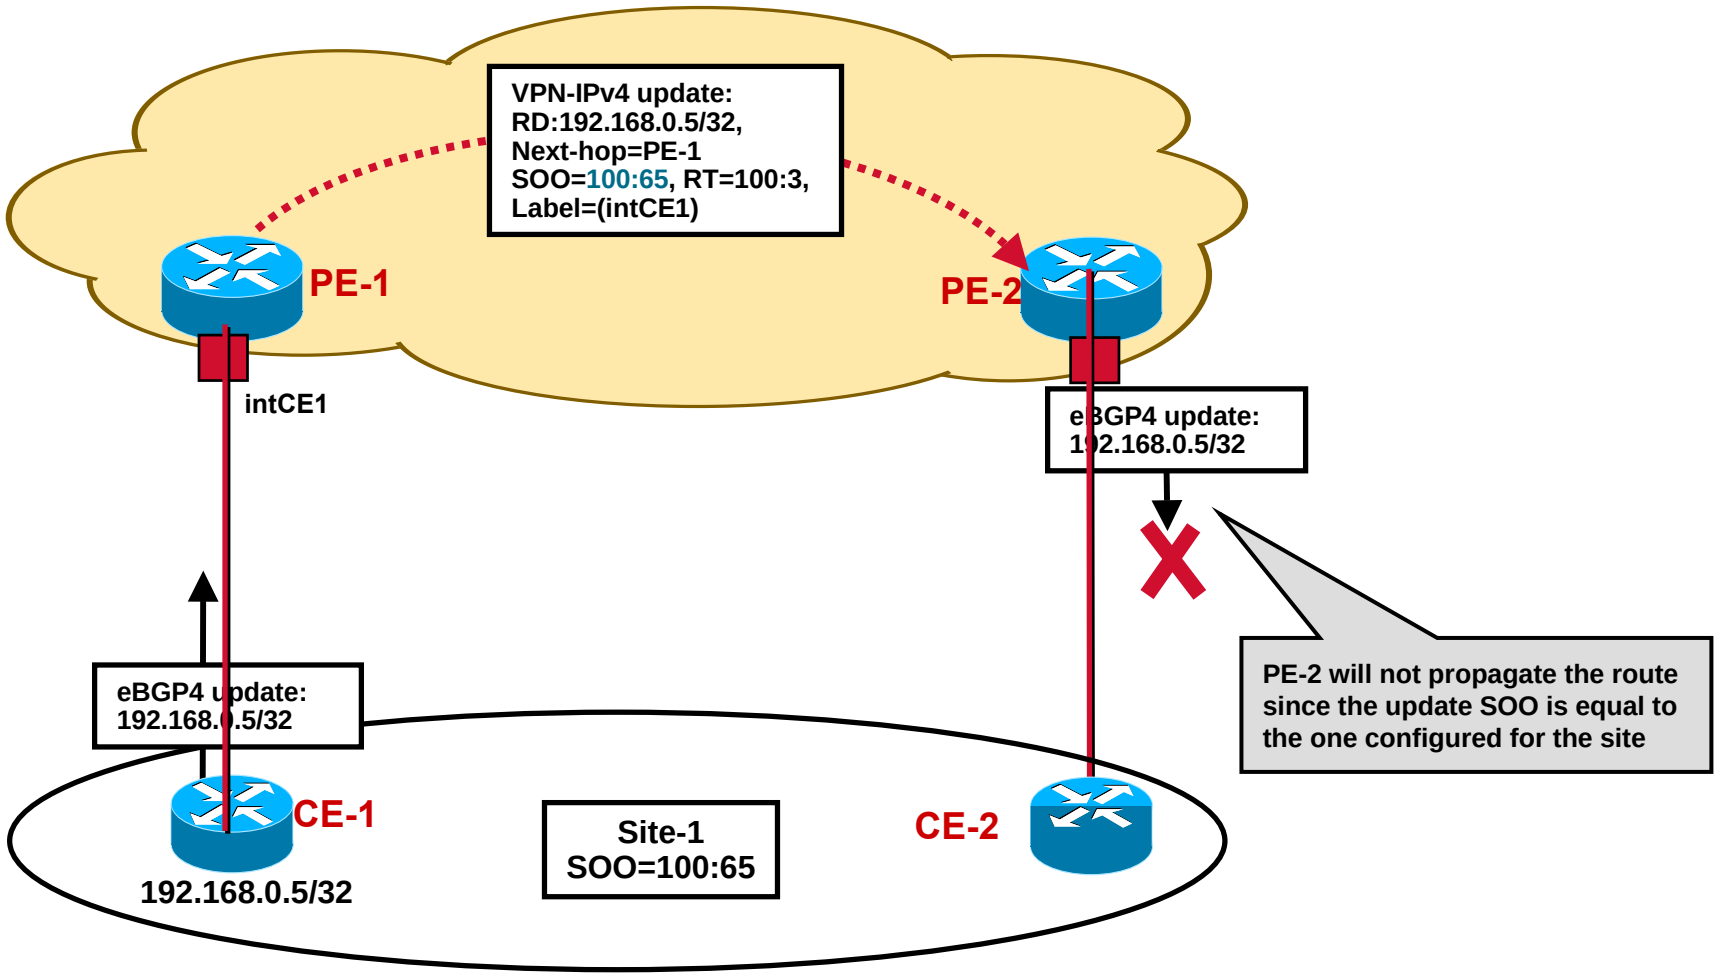

© 2001, Cisco Systems, Inc. All rights reserved.

Selective Export

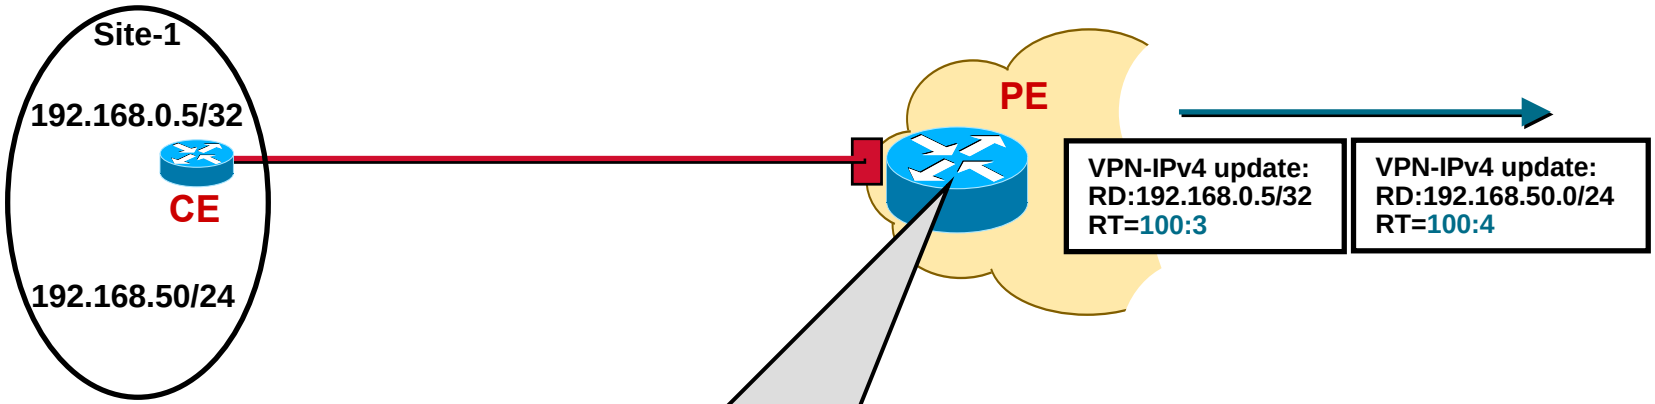


```
ip vrf odd
 rd 100:1
 export map RTMAP
 route-target import 100:3
!
…
…
!
access-list 10 permit 192.168.0.5 0.0.0.0
access-list 11 permit any
!
route-map RTMAP permit 10
  match ip address 10
  set extcommunity rt 100:3
!
route-map RTMAP permit 20
  match ip address 11
  set extcommunity rt 100:4
!
```

© 2001, Cisco Systems, Inc. All rights reserved.

Selective Import

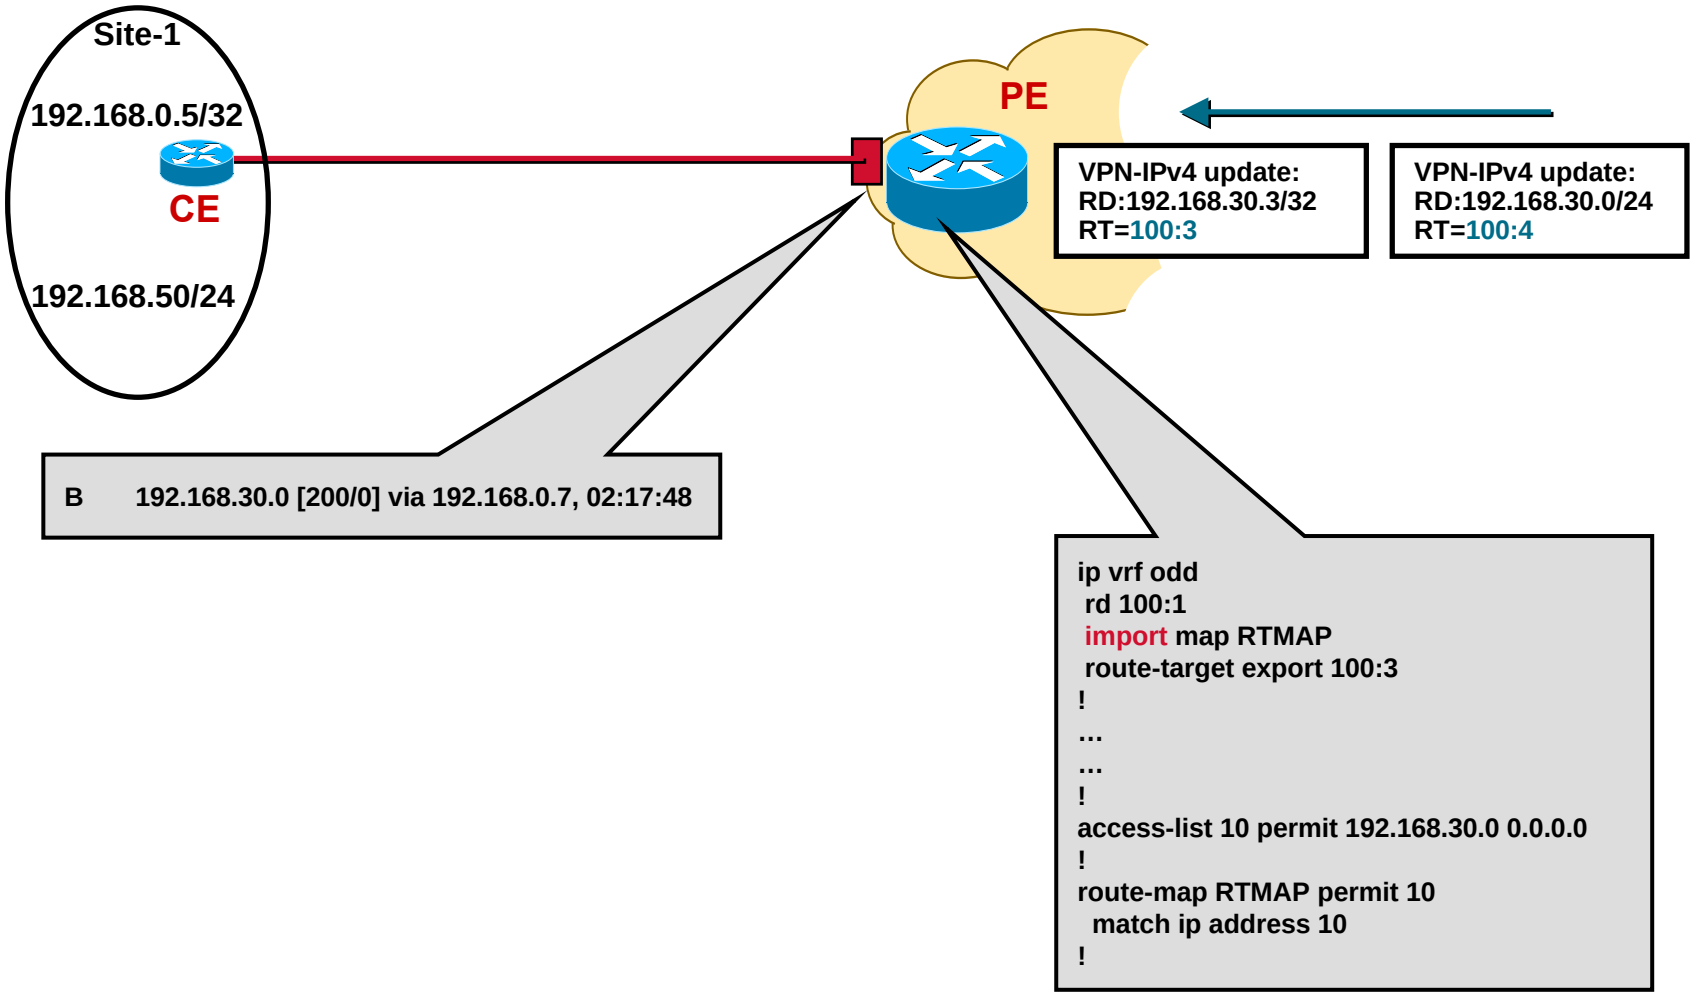

© 2001, Cisco Systems, Inc. All rights reserved.

Suggested reading

draft-ietf-mpls-arch-05.txt

draft-ietf-mpls-label-encaps-04.txt

draft-ietf-mpls-atm-02.txt

draft-ietf-mpls-ldp-03.txt

© 2001, Cisco Systems, Inc. All rights reserved.

© 2001, Cisco Systems, Inc. All rights reserved.