



Concepts and Operation of MPLS VPNs

Francisco Bolanos
fbolanos@cisco.com

Agenda

- **MPLS Concepts**
- **Label Structure**
- **Label assignment and distribution**
- **RD, RT and VRF instances**
- **Service Models**
- **MPLS/VPN Configuration**

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.**

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

- **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 or more PE/CE links

MPLS/VPN Model



Control Component

Protocols for label exchange:

Control Component is responsible for **binding between labels and routes**

LDP (646)/ TDP (711)

Label/Tag Distribution Protocol maps unicast IP destinations into labels

For Signalling and service control we have:

- **RSVP, CR-LDP**
Used in traffic engineering
- **BGP/M-BGP**
External labels (VPN)
- **PIM**

For multicast states label mapping

- **MPLS Switching defines traffic flows based on FECs**
- **FEC: Forwarding Equivalence Class**
- **A FEC can represent a: Destination address prefix, VPN, Traffic Engineering tunnel, Class of Service.**

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

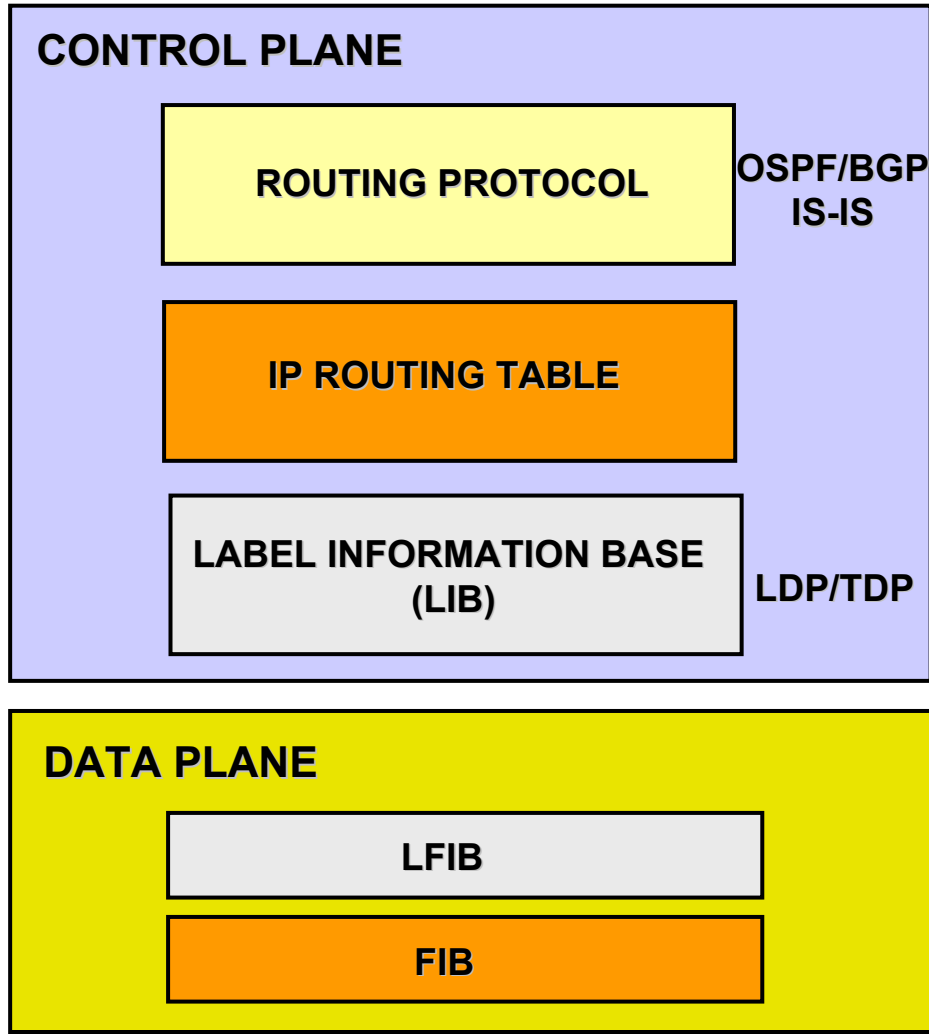
Contains only label/prefix mappings that are currently in use for label forwarding

MPLS Architecture

OSI



MPLS




Label Structure




MPLS Shim header structure




IP v4 HEADER



MPLS SHIM HEADER




MPLS LABEL STACKING




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

Label assignment and distribution



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.


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 |

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

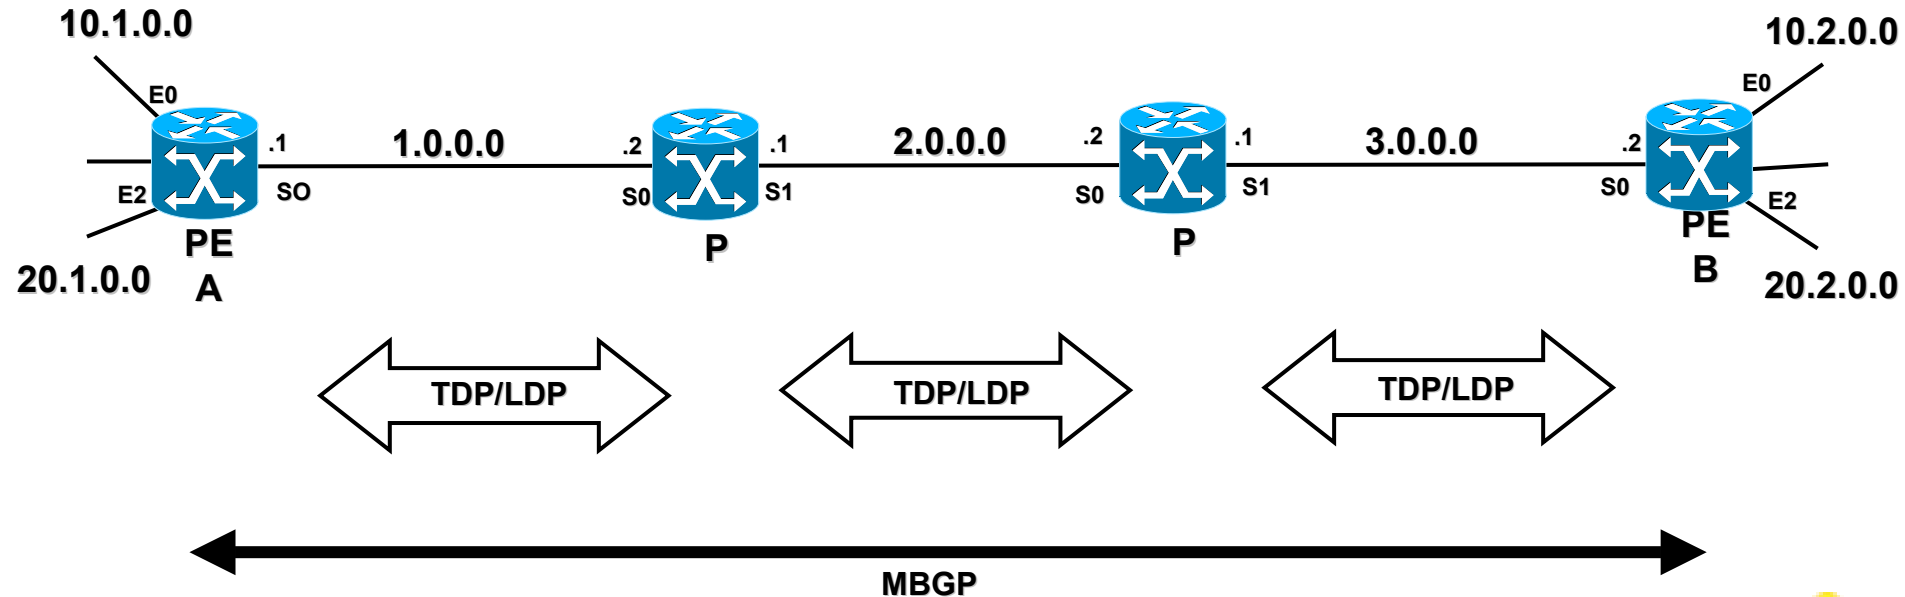
| Address | Prefix | LocLbl | NHLbl |
|----------|--------|--------|-------|
| 10.1.0.0 | /16 | 8,15 | POP |
| 10.2.0.0 | /16 | 9,2 | 9 |
| 20.1.0.0 | /16 | 10,17 | POP |
| 20.2.0.0 | /16 | 11,4 | 11 |
| 1.0.0.0 | /16 | 12,19 | Null |
| 2.0.0.0 | /16 | 13,6 | Null |
| 3.0.0.0 | /16 | 14,7 | POP |

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 |



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 |

10.1.0.0

10.2.0.0



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 |





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 |

M-BGP



S=0



S=0



S=1



RD, RT and VRF instances



Route Distinguisher (RD)

- New address family: VPN-IP addresses

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

RDs are assigned by a service provider to each PE

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

- Route Distinguisher:

Used to create a unique IP address within MPLS Domain,
VPNv4 address format would look like:

2586:10:192.168.23.0
└──┬──┘ └──────────────────┘
RD **IP Add**

Route Target (RT)

- RT controls the import/export
- 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!

- **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

Service Models




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**

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)

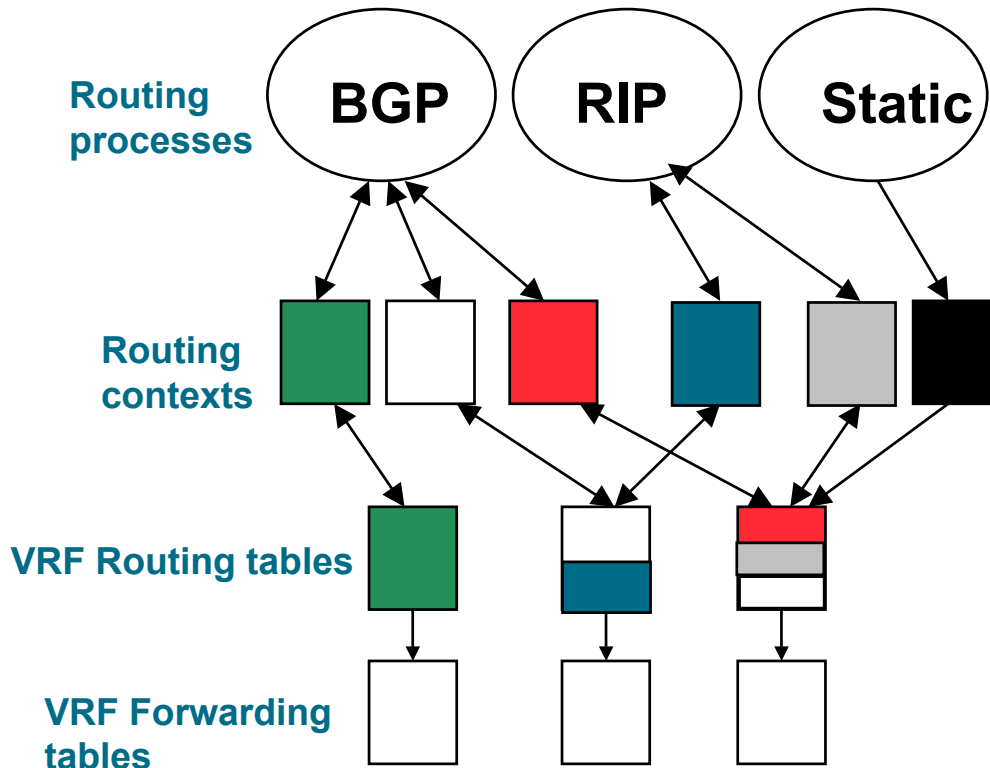
VRF are associated to (sub/virtual/tunnel)interfaces

Interfaces may share the same VRF if the connected sites may share the same routing information

MPLS VPN mechanisms


VRF and Multiple Routing Instances

Cisco.com



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

MPLS VPN Connection Model



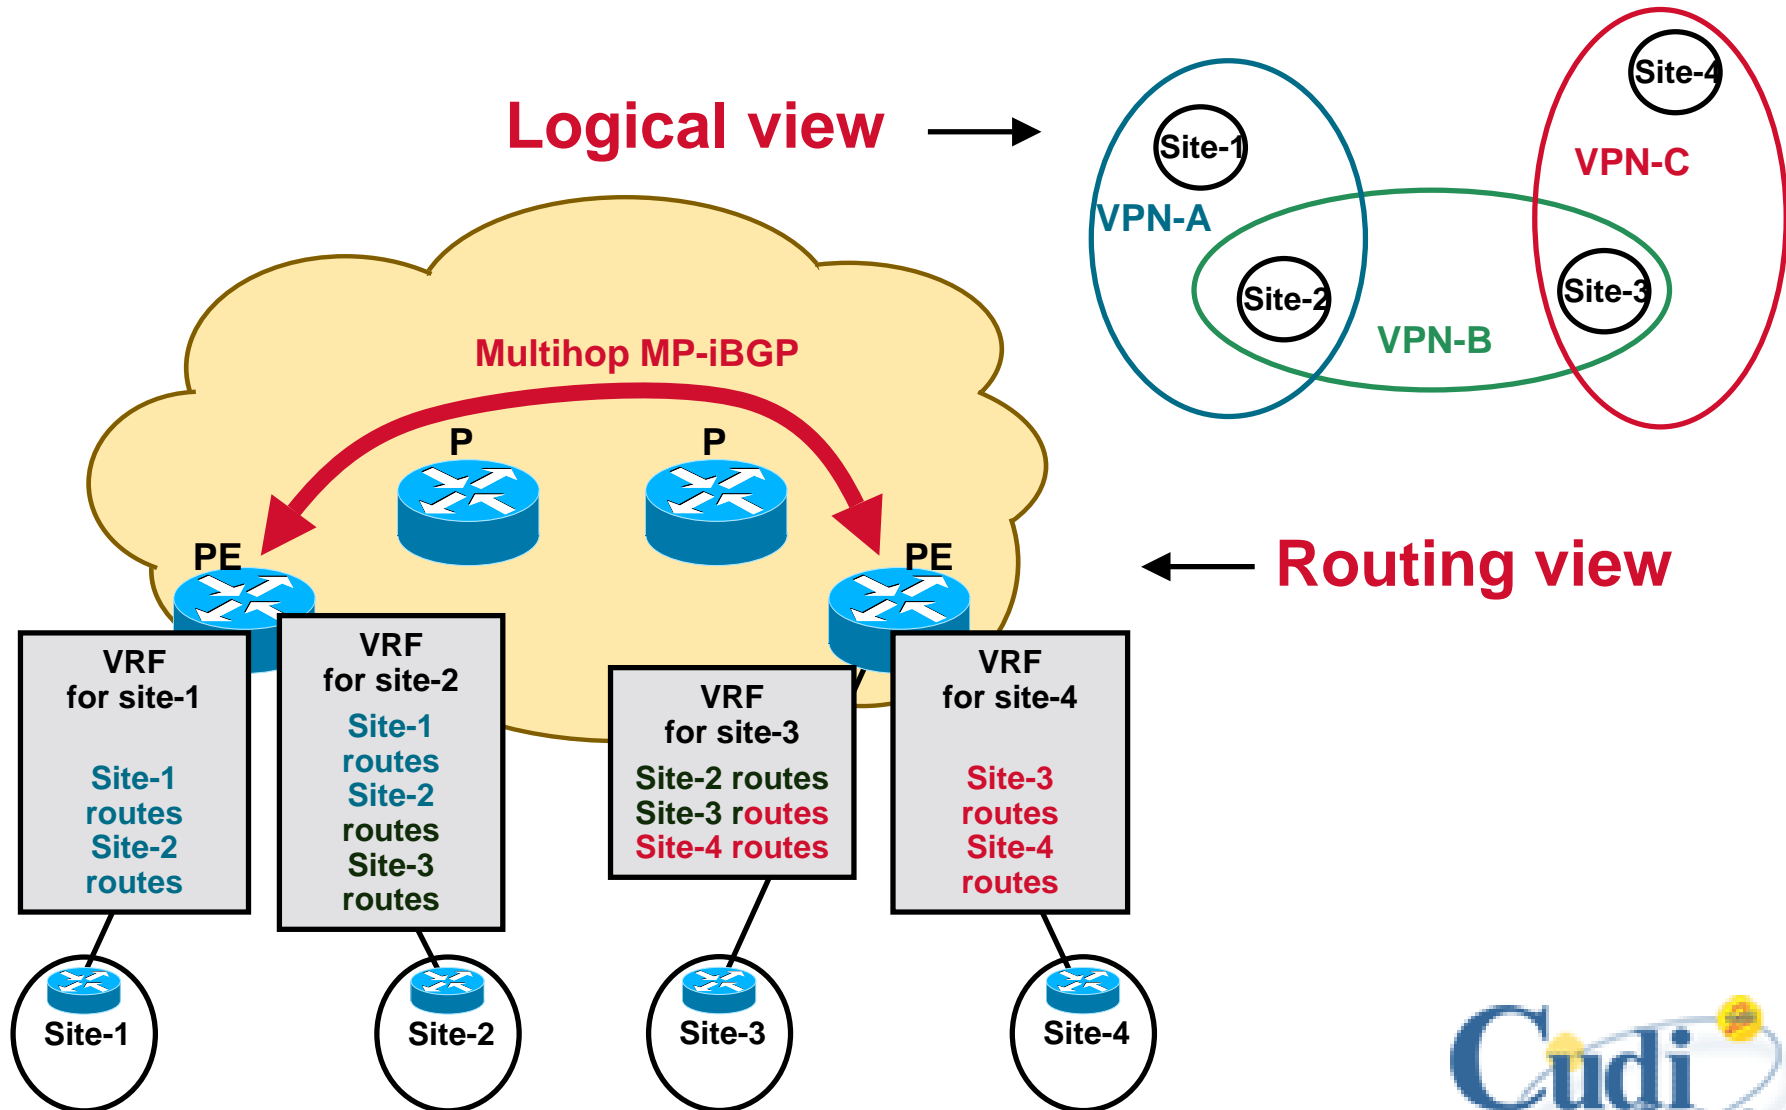
- **PE and P routers share a common IGP (ISIS or OSPF)**
- **PEs establish MP-iBGP sessions between them**
- **PEs use MP-BGP to exchange routing information related to the connected sites and VPNs**

**VPN-IPv4 addresses, Extended Community,
Label**

MPLS VPN mechanisms

VRF and Multiple Routing Instances

Cisco.com



MPLS/VPN Configuration



MPLS VPN - Configuration

- The following global commands are mandatory **BEFORE** configuring MPLS

Ip cef

Mpls ip

No tag-switching ip propagate-ttl (optional if no ttl propagation is needed)

- **VRF configuration command**

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


MPLS VPN - Configuration

PE/CE routing protocols

Cisco.com

- PE/CE may use BGP, RIPv2, OSPF 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

...



- 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> ...
```

MPLS VPN – Configuration Verification

- 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>

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

www.cisco.com/go/mpls

CISCO SYSTEMS



EMPOWERING THE
INTERNET GENERATIONSM