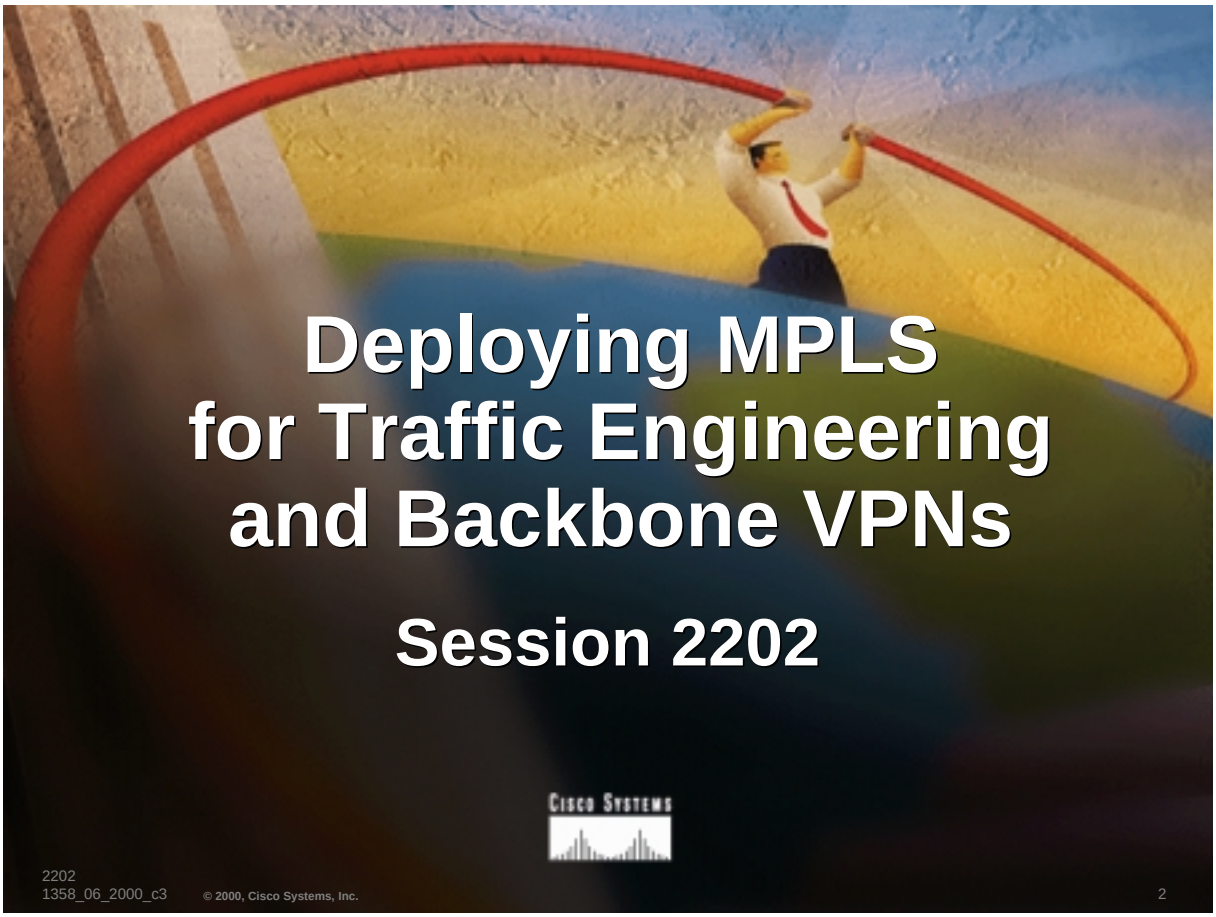

Agenda

- Things to Remember
- Why You Want MPLS TE
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Verifying Tunnel Setup
- Deployment Strategy

Things to Remember

- Tunnels are unidirectional
- Routing protocols are not run over TE tunnels

Agenda

- Things to Remember
- Why You Want MPLS TE
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Verifying Tunnel Setup
- Deployment Strategy

Why You Want MPLS TE

- Most backbones have underutilized capacity and over utilized capacity
- TE lets you use paths other than IGP
- Unequal-cost loadsharing!

Agenda

- Things to Remember
- Why You Want MPLS TE
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Verifying Tunnel Setup
- Deployment Strategy

Minimum TE Transit Config

- **What code to use?**
  - **12.0S, 12.1, 12.1T**
  - **Most networks are running 12.0S**

Minimum TE Transit Config

```
R1(config)#ip cef
```

- **CEF necessary for all MPLS features**
- **Can also be '`ip cef distributed`' on some platforms (GSR, 7500)**

Minimum TE Transit Config

- IGP **must** be link-state (OSPF or IS-IS)
- path determination could not work on DV without being extremely ugly
  - TE is all about calculating paths other than the shortest one
  - Distance-vector protocols don't have enough information in them to do this

Minimum TE Transit Config

- Link-state view of the network

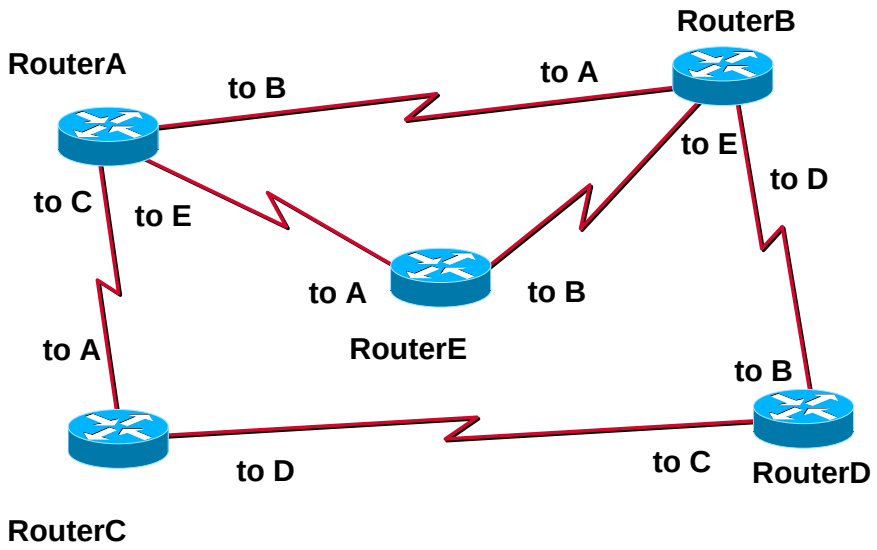

Minimum TE Transit Config

- Router A's DV view of the network

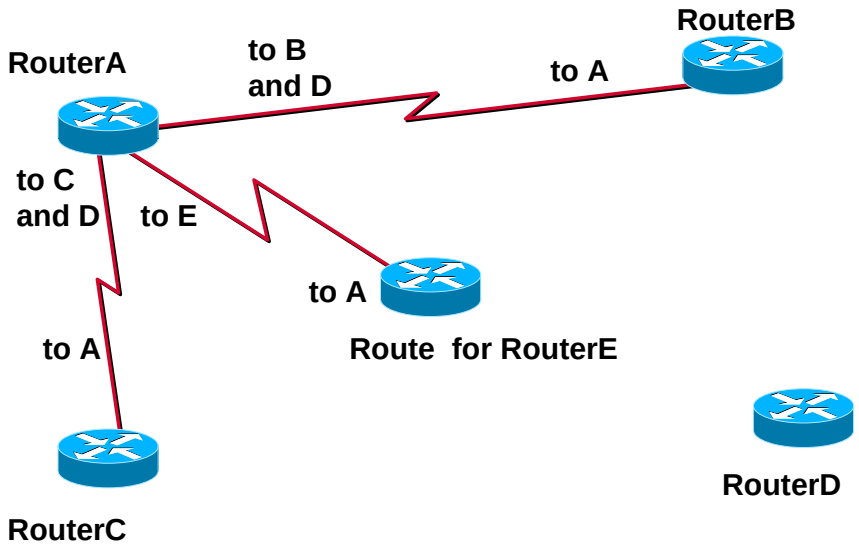

Minimum TE Transit Config

```
R1(config)#interface Loopback0
R1(config-if)#ip address
  192.168.1.1 255.255.255.255
```

- Configure loopbacks everywhere
- Loopbacks must be /32
- Lookback must be in your IGP

Minimum TE Transit Config

```
R1(config)#mpls traffic-eng tunnels
```

- Enables TE on the router
- Global command

Minimum TE Transit Config

```
R1(config-if)#
  mpls traffic-eng tunnels
```

- Enables RSVP signaling on an interface
- This is needed on both ends of any link an LSP could pass over

Minimum TE Transit Config

Configuring IS-IS for MPLS TE

```
router isis
 mpls traffic-eng level <1 or 2>
 mpls traffic-eng router-id loop0
 metric-style wide
```


Minimum TE Transit Config

Configuring IS-IS for MPLS TE

```
 mpls traffic-eng router-id loop0
```

- Must explicitly configure RID

Minimum TE Transit Config

Configuring IS-IS for MPLS TE

```
mpls traffic-eng level <1 or 2>
```

- TE support is currently single-area
- Most ISP backbones are all L1 or L2

Minimum TE Transit Config

Configuring IS-IS for MPLS TE

```
metric-style wide
```

- Wide metrics necessary for IS-IS TE
- Must transition IGP to wide metric
- See MPLS TE feature docs on CCO

Minimum TE Transit Config

Configuring OSPF for MPLS TE

```
router ospf <pid>
 mpls traffic-eng area <area>
 mpls traffic-eng router-id loop0
```


Minimum TE Transit Config

Configuring OSPF for MPLS TE

```
 mpls traffic-eng area <area>
```

- OSPF support currently single area
- Typically TE is in area 0 (backbone)

Minimum TE Transit Config

Configuring OSPF for MPLS TE

```
mpls traffic-eng router-id loop0
```

- Must explicitly configure RID

Agenda

- Things to Remember
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Other Useful Features
- Verifying Tunnel Setup
- Deployment Strategy

Minimum TE Headend Config

```
interface Tunnel0
  ip unnumbered Loopback0
  tunnel mode mpls traffic-eng
  tunnel destination <tail RID>
  tunnel mpls traffic-eng autoroute announce
  tunnel mpls traffic-eng path-option <priority>
                                      dynamic
```


Minimum TE Headend Config

`interface Tunnel0`

- **Creates LSP headend**
- **Remember, MPLS tunnels are unidirectional**

Minimum TE Headend Config

```
ip unnumbered Loopback0
```

- Gives the headend an IP address
- Tunnel inf. must have an IP address

  Cisco IOS® will not route IP across an interface with no IP address

Minimum TE Headend Config

```
tunnel mode mpls traffic-eng
```

- Tells IOS this is an MPLS TE tunnel
- Other choices would be GRE, etc...

Minimum TE Headend Config

```
tunnel destination <addr>
```

- Gives tail end of the tunnel
- Destination must be tailend RID

Minimum TE Headend Config

```
tunnel mpls traffic-eng
                autoroute announce
```

- Announces tunnel tail reachability to the RIB
- Without this, tunnel comes up but is not used

Minimum TE Headend Config

```
tunnel mpls traffic-eng path-option <priority>
                                    dynamic
```

- Need to tell the tunnel how to calculate its path
- '`path-option`' has other knobs; this is the minimum

Agenda

- Things to Remember
- Why You Want MPLS TE
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Verifying Tunnel Setup
- Deployment Strategy

Basic Options

Minimal Setup to Make LSP != IGP

- RSVP
- Dynamic vs. explicit path
- Tunnel bandwidth

Basic Options

```
R1(config-if)#ip rsvp
     bandwidth <x> <y>
```

- Needed to allow non-zero BW res.
- x = max reservable bw
- y = max reservable bw for single LSP
- no fallback BW on a tunnel

Basic Options

Unequal-cost load balancing!

```
interface Tunnel0
  tunnel destination 1.2.3.4
  tunnel mpls traffic-eng bandwidth 3

interface Tunnel1
  tunnel destination 1.2.3.4
  tunnel mpls traffic-eng bandwidth 1
```


Basic Options

Unequal-cost load balancing!

```
VXR-6#
Routing entry for 192.168.1.4/32
  Known via "isis", distance 115, metric 30, type level-2
  Redistributing via isis
  Last update from 192.168.1.4 on Tunnel0, 00:00:12 ago
  Routing Descriptor Blocks:
  * 192.168.1.4, from 192.168.1.4, via Tunnel1
      Route metric is 30, traffic share count is 1
    192.168.1.4, from 192.168.1.4, via Tunnel0
      Route metric is 30, traffic share count is 3
```

Basic Options

- Load-balancing is not perfect
- CEF has 16 buckets to store adjacency info in
- 1:9 or 1:10:100 won't work the way you think
  - 1:9 ends up as 1:15
  - 1:10:100 ends up as 1:4:9

Basic Options

```
tunnel mpls traffic-eng path-option
<prio> <dynamic|explicit <id|name>
[ID|NAME]>
```

- dynamic = router calculates path
- Router will take best IGP path that meets BW requirements
- BW req. of 0 == normal IGP path

Basic Options

```
tunnel mpls traffic-eng path-option
<prio> <dynamic|explicit <id|name>
[ID|NAME]>
```

- explicit = take specified path
- Path is a list of RIDs
- Strict source-routing of IP traffic

Basic Options

```
ip explicit-path <id|name> [ID|NAME]
    next-address 192.168.1.1
    next-address 192.168.2.1
    ...
```

- explicit = take specified path
- Router sets up path you specify
- Strict source-routing of IP traffic
- Each hop is a RID

Basic Options

```
VXR-6(cfg-ip-expl-path)#?
Explicit-Path configuration commands:
  append-after  Append additional entry after specified index
  exit          Exit from explicit-path configuration mode
  index         Specify the next entry index to add, edit (or
delete)
  list          Re-list all or part of the explicit path entries
  next-address  Specify the next (adjacent) address in the path
  no            Delete a specific explicit-path entry index
```


Basic Options

- **How RSVP sets up a path**

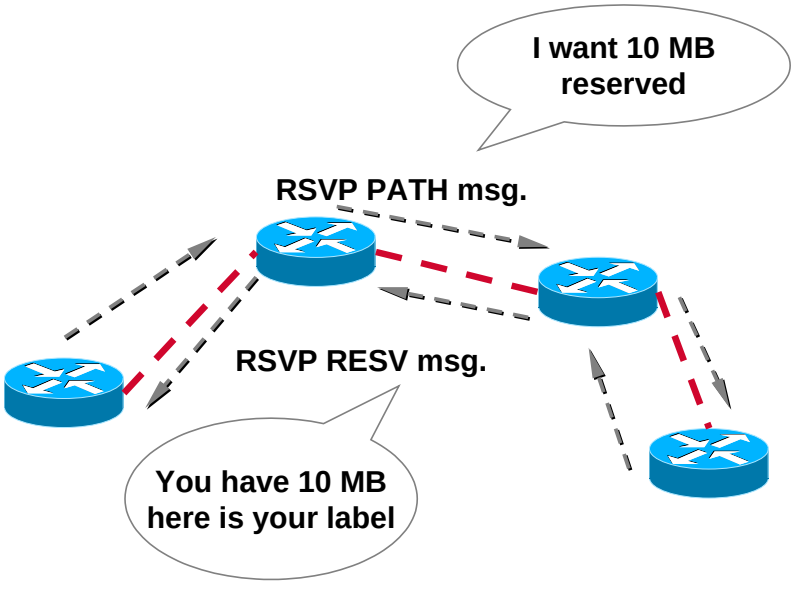

Basic Options

- Use tunnel priorities for fallback paths
- Good idea to have last fallback a dynamic tunnel

Basic Options

- Multiple tunnel priorities

```
interface Tunnel0

   tunnel mpls traffic-eng path-
option 10 explicit name first-path

   tunnl mpls traffic-eng path-
option 20 explicit name second-path

   tunnel mpls traffic-eng path-
option 30 dynamic
```

Agenda

- Things to Remember
- Minimum TE Transit Config
- Minimum TE headend Config
- Basic Options
- Other Useful Features
- Verifying Tunnel Setup
- Deployment Strategy

Verifying Tunnel Setup

- sh int tun<x>

  up/up == tunnel is up/LSP is successfully built
- sh ip route <tunnel dest>

  should be via the configured tunnel(s)
- sh mpls traffic-eng tunnels

  lots of options under here

Verifying Tunnel Setup

- **sh int tun<x>**

  up/up == tunnel is up/LSP is successfully built

```
VXR-6#sh int tun0
Tunnel0 is up, line protocol is up
  Hardware is Tunnel0
  ...
  Encapsulation TUNNEL, loopback not set
  Keepalive set (10 sec)
  Tunnel source 192.168.1.6, destination 192.168.1.4
```


Verifying Tunnel Setup

- **sh ip route <tunnel dest>**

  should be directly connected via tunnel

```
VXR-6#sh ip rou 192.168.1.4
Routing entry for 192.168.1.4/32
  Known via "isis", distance 115, metric 30, type level-2
  ...
  Routing Descriptor Blocks:
  * 192.168.1.4, from 192.168.1.4, via Tunnel0
      Route metric is 30, traffic share count is 1
```

Verifying Tunnel Setup

- **sh mpls traffic-eng tunnels**

  lots of options here

```
VXR-6#sh mpls traffic-eng tunnels ?
  Tunnel       Tunnel interface
  brief        Brief summary of tunnel status and configuration
  destination  Restrict display to tunnels with this destination
  name         Restrict display to tunnels with this name
  role         Restrict display to tunnels with specified role
  source-id    Tunnel identifier address/id
  summary      Show summary information
  up           Restrict display to tunnels in up state
  <cr>
```


Verifying Tunnel Setup

- **sh mpls traffic-eng tunnels <intf>**

```
VXR-6#sh mpls traffic-eng tunnels Tunnel0
Name: VXR-6_t0                           (Tunnel0) Destination: 192.168.1.4
  Status:
    Admin: up         Oper: up     Path: valid       Signalling: connected

    path option 10, type dynamic (Basis for Setup, path weight 20)

  Config Paramters:
    Bandwidth:  6         Priority: 7  7   Affinity: 0x0/0xFFFF
    AutoRoute:  enabled   LockDown: disabled
...
    RSVP Path Info:
      My Address: 192.168.4.6
      Explicit Route: 192.168.4.5 192.168.6.5 192.168.6.4 192.168.1.4
      Record   Route:  NONE
      Tspec: ave rate=6 kbits, burst=1000 bytes, peak rate=6 kbits
```

Verifying Tunnel Setup

- **sh mpls traffic-eng tunnels brief**

```
VXR-6#sh mpls traffic-eng tunnels brief
Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Periodic reoptimization:        every 10 seconds, next in 5 seconds
TUNNEL NAME                          DESTINATION      STATUS      STATE
VXR-6_t0                             192.168.1.4      up          up
VXR-6_t1                             192.168.1.4      up          up
VXR-4_t0                             192.168.1.6      signalled   up
Displayed 2 (of 2) heads, 0 (of 0) midpoints, 1 (of 1) tails
```


Verifying Tunnel Setup

- **sh mpls traffic-eng tunnels summary**

```
VXR-6#sh mpls tr tu summary
Signalling Summary:
    LSP Tunnels Process:            running
    RSVP Process:                   running
    Forwarding:                     enabled
    Head: 2 interfaces, 2 active signalling attempts, 2 established
          24 activations, 22 deactivations
    Midpoints: 0, Tails: 1
    Periodic reoptimization:        every 10 seconds, next in 8 seconds
```

Agenda

- Things to Remember
- Minimum TE Transit Config
- Minimum TE Headend Config
- Basic Options
- Other Useful Features
- Verifying Tunnel Setup
- Deployment Strategy

Deployment Strategy

Two strategies to TE deployment

- Full mesh of TE tunnels

  "Proper"—more complex procedure, but more benefits overall

- As needed to clear up congestion

  "Hack"—may be harder to track over time

Deployment Strategy

- TE tunnel bandwidth is not policed
- Ratio is more important than actual bandwidth for small hacks
- Actual BW use is more important for large-scale deployment

Deployment Strategy

Full Mesh of TE Tunnels—Steps

1) Enable CEF everywhere
2) Config your IGP, RSVP
3) Configure full mesh of zero-bandwidth tunnels
4) Turn up tunnels via ‘autoroute announce’

Deployment Strategy

Full Mesh of TE Tunnels—Steps

5) Let tunnels run for 1–2 weeks

6) Add BW requirements to tunnels

7) Use off-line tools if necessary

   WANDL, Cisco

Deployment Strategy

As Needed to Clear Up Congestion

1) Enable CEF everywhere

2) Config your IGP, RSVP

3) Configure TE tunnels around congested links

   One IGP tunnel, one or more explicit-path tunnels

Deployment Strategy

As Needed to Clear Up Congestion

4) Turn up tunnels one at a time via ‘autoroute announce’

5) Add BW requirements to tunnels

Tunnel BW ratio is important

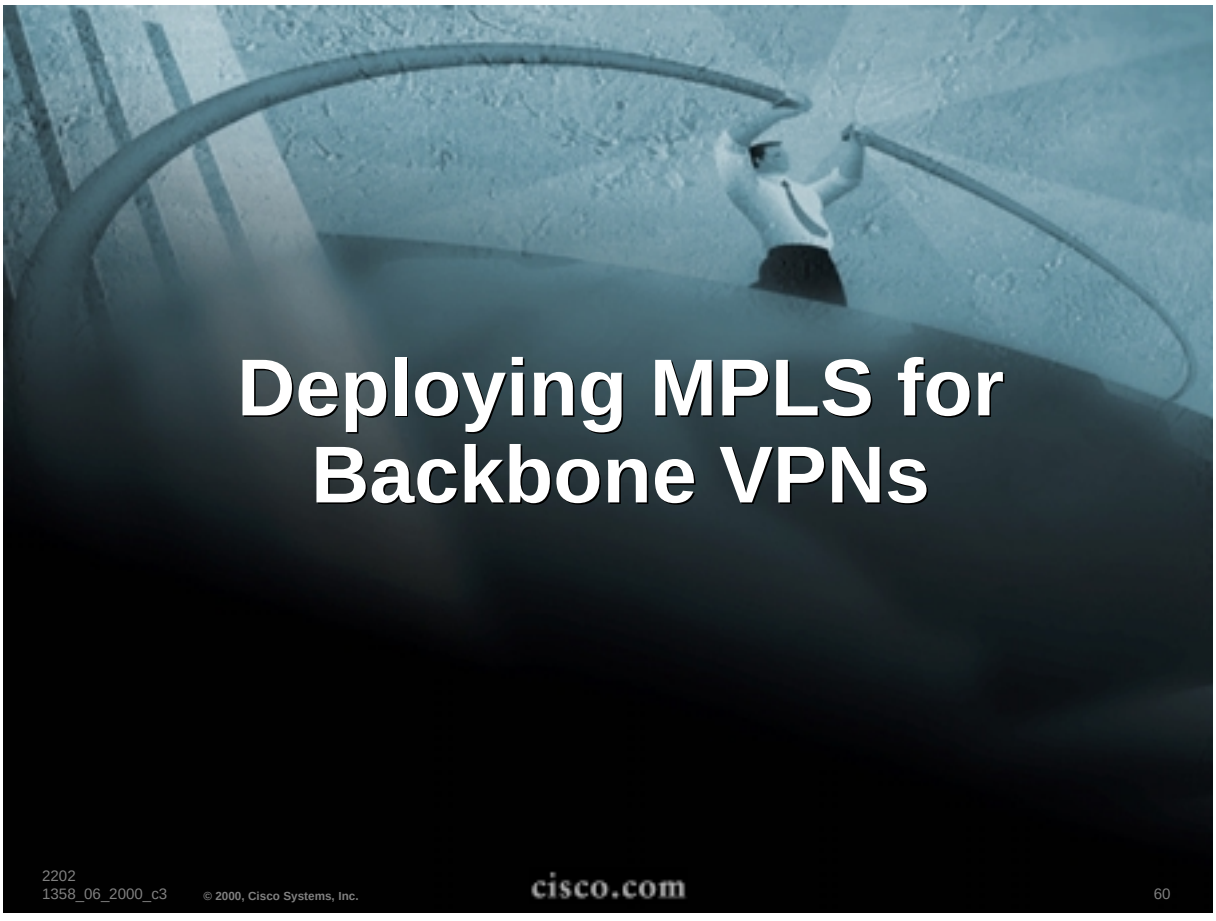

Agenda

- **MPLS-VPN Terminology**
- **Network Topology**
- **Step By Step Configuration Guideline**
  - **Sample configuration (VRF, interfaces)**
  - **PE-CE Routing Protocol configuration (Static, RIPv2, BGP, OSPF)**
  - **MP-BGP configuration**
- **MPLS VPN Scaling**
  - **Route-Reflectors, Filters**
- **MPLS VPN Internet Routing**

MPLS VPN Network

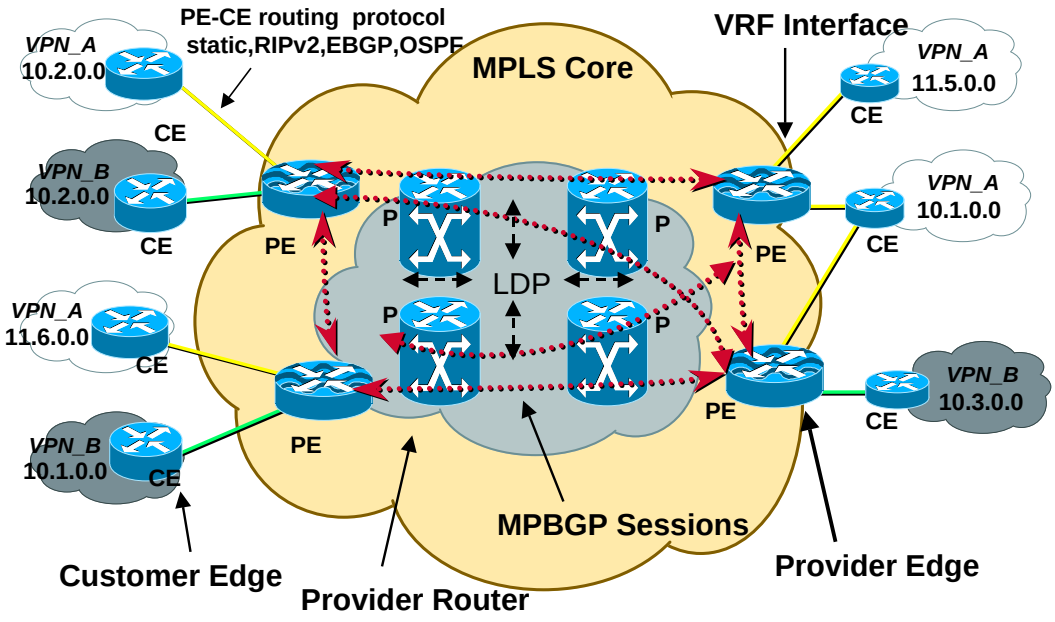

Working of MPLS VPN Network

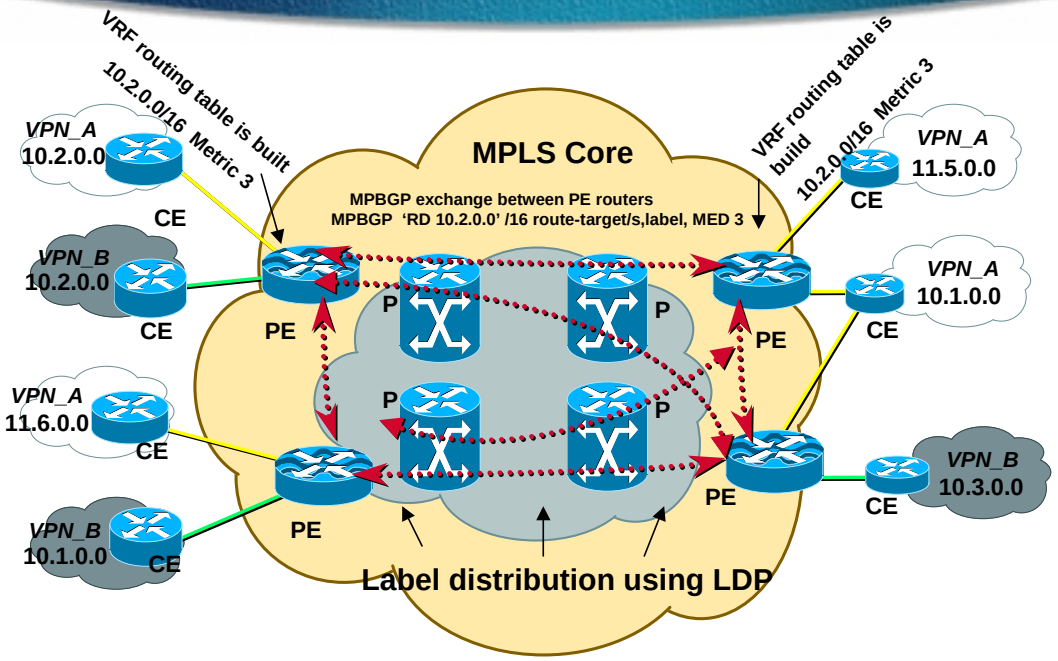

MPLS-VPN Terminology

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

  The backbone under control of a Service Provider

- **Customer Network (C-Network)**

  Network under customer control

- **CE router**

  Customer Edge router. Part of the C-network and interfaces to a PE router

MPLS-VPN Terminology

- **Site**

  Set of (sub)networks part of the C-network and co-located

  A site is connected to the VPN backbone through one or more PE/CE links

- **PE router**

  Provider Edge router. Part of the P-Network and interfaces to CE routers

- **P router**

  Provider (core) router, without knowledge of VPN

MPLS-VPN Terminology

- **Route-Target**

  64 bits identifying routers that should receive the route

- **Route Distinguisher**

  Attributes of each route used to uniquely identify prefixes among VPNs (64 bits)

  VRF based (not VPN based)

- **VPN-IPv4 addresses**

  Address including the 64 bits Route Distinguisher and the 32 bits IP address

MPLS-VPN Terminology

- **VRF**
  - **VPN Routing and Forwarding Instance**
  - **Routing table and FIB table**
  - **Populated by routing protocol contexts**
- **VPN-Aware network**
  - **A provider backbone where MPLS-VPN is deployed**

Agenda

- **MPLS-VPN Terminology**
- **Network Topology**
- **Step By Step Configuration Guideline**
  - **Sample configuration (VRF, interfaces)**
  - **PE-CE Routing Protocol configuration (Static, RIPv2, BGP, OSPF)**
  - **MP-BGP configuration**
- **MPLS VPN Scaling**
  - **Route-Reflectors, Filters**
- **MPLS VPN Internet Routing**

Step By Step Configuration Guideline

- LDP configuration in MPLS backbone
- 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
  - (Static, RIPv2, BGP)
  - MP-BGP between PE routers

Step By Step Configuration Guideline

- Configure LDP in MPLS backbone
  - On all interfaces in MPLS backbone configure
  - tag-switching ip
- Verify the configuration
  - Show tag forwarding
  - Show tag tdp neighbour
  - Labels for Non BGP routes are created
  - Note: Configuration will be slightly different in future IOS

Step By Step Configuration Guideline

Configure VRF instance on PE router

ip vrf <vrf-symbolic-name> (Name of the VPN site)
rd <route-distinguisher-value>
route-target import <Import route-target community>
route-target export <Import route-target community>

example:

```
ip vrf VPN-A
rd 100:1
route-target export 100:1
route-target import 100:1
```


Step By Step Configuration Guideline

Configure VRF instance on PE router

rd <route-distinguisher-value>

You may pick a unique **rd** for each **vrf** interface.

Ease in manageability

Can achieve load sharing using unique **rd** every where

route-target import <Import route-target community>
route-target export <Import route-target community>

dictates who will receive the routes

Inter VPN connectivity is very simple by importing **route-target** form other VPN

It can be same value as **rd**

Step By Step Configuration Guideline

Identify VRF Interface on PE Router

```
Interface Serial0
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
```

Verify the configuration

Show ip vrf interface

```
Interface              IP-Address      VRF                            Protocol
Serial0                192.168.10.1    VPN-A                          up
```


Step By Step Configuration Guideline

Configure Routing Protocol Between Provider and Customer (PE and CE)

- **Supported protocols**

  Static, RIPv2, eBGP, OSPF

- **Learned routes will populate the VRF associated to the interface where the protocol is configured**
- **Same routing protocol instance with different routing contexts**

Step By Step Configuration Guideline PE-CE Routing Protocol (Static)

Static routes into VRFs need outgoing interface

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

ip route vrf VPN-A 10.0.0.0 255.0.0.0 Serial0 192.168.10.2

Verify the configuration

```
sh ip route vrf VPN-A
C    192.168.10.0/24 is directly connected, Serial0
S    10.0.0.0/8 [1/0] via 192.168.10.2,  Serial0
```


Step By Step Configuration Guideline PE-CE Routing Protocol (RIPv2)

- **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
…

Step By Step Configuration Guideline PE-CE Routing Protocol (RIPv2)

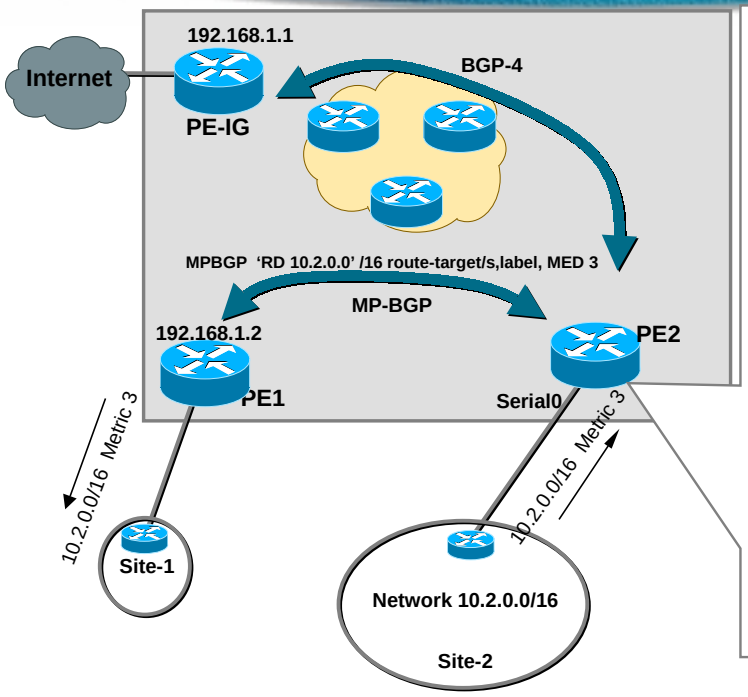


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
!
Interface Serial0
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
!
router rip
 version 2
 !
 address-family ipv4 vrf VPN-A
 version 2
 redistribute connected metric 2
 redistribute bgp 100 metric 2
 network 192.168.10.0
 no auto-summary
 exit-address-family

[snip]
```

```
PE2#sh ip route vrf VPN-A rip
```


Step By Step Configuration Guideline PE-CE Routing Protocol (RIPv2)

- **RIPv2/BGP interaction**

  During redistribution,

  RIP metric propagated BGP MED

  BGP MED automatically propagated into RIP metric

  Allows transparency of RIP metric through the MPLS backbone

Step By Step Configuration Guideline PE-CE Routing Protocol (RIPv2)

- **Use RIP with stub sites and when convergence is not an issue**

Step By Step Configuration Guideline PE-CE Routing Protocol (BGP)

BGP uses same “address-family” command

```
Router BGP <asn>
...
address-family ipv4 vrf <vrf-symbolic-name>
…
  any common router BGP sub-command
…
```

Step By Step Configuration Guideline PE-CE Routing Protocol (BGP)

```
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
 !
```

```
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
 !

PE2#sh ip route vrf site3 bgp 100
```

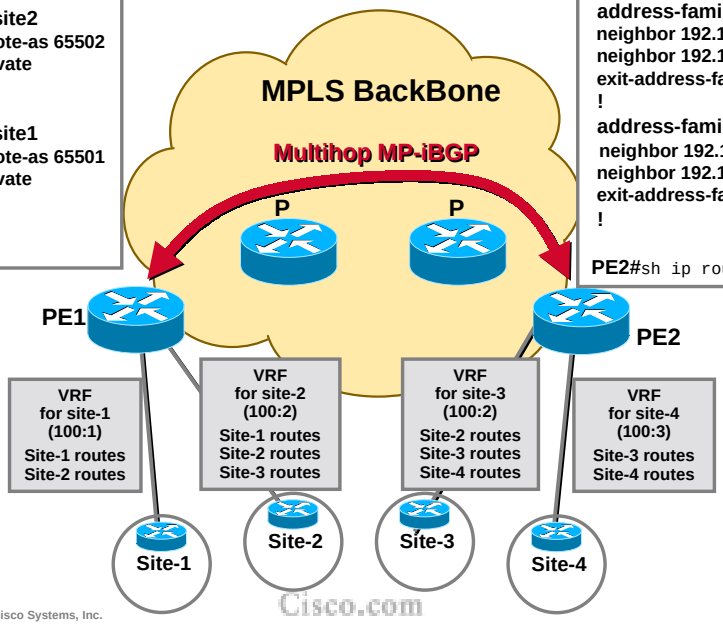

Step By Step Configuration Guideline MPBGP

```
router bgp 100
no bgp default ipv4-unicast
 neighbor 7.7.7.7 remote-as 100
 neighbor 7.7.7.7 update-source Loop0
!
[snip]

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

[snip]

address-family vpnv4
 neighbor 6.6.6.6 activate
 neighbor 6.6.6.6 next-hop-self
exit-address-family
```

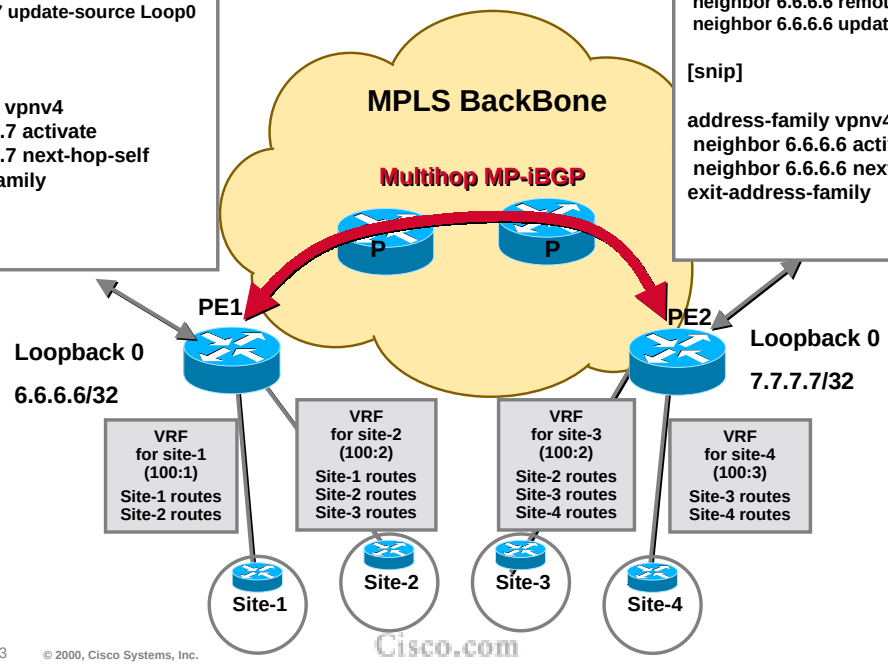

Step By Step Configuration Guideline MPBGP

Verify the working of MPBGP

```
sh ip bgp vpnv4 ?
  all  Display information about all VPNv4 NLRIs
  rd   Display information for a route distinguisher
  vrf  Display information for a VPN Routing/Forwarding instance

example:
```

```
sh ip bgp vpnv4 rd 100:1
```


Agenda

- MPLS-VPN Terminology
- Network Topology
- Step By Step Configuration Guideline
  - Sample configuration (VRF, interfaces)
  - PE-CE Routing Protocol configuration (Static, RIPv2, BGP, OSPF)
  - MP-BGP configuration
- MPLS VPN Scaling
  - Route-Reflectors, Filters
- MPLS VPN Internet Routing

MPLS-VPN Scaling BGP

- Existing BGP techniques can be used to scale the route distribution: route reflectors
- Each edge router needs only the information for the VPNs it supports

  Directly connected VPNs

- RRs are used to distribute VPN routing information

MPLS-VPN Scaling BGP

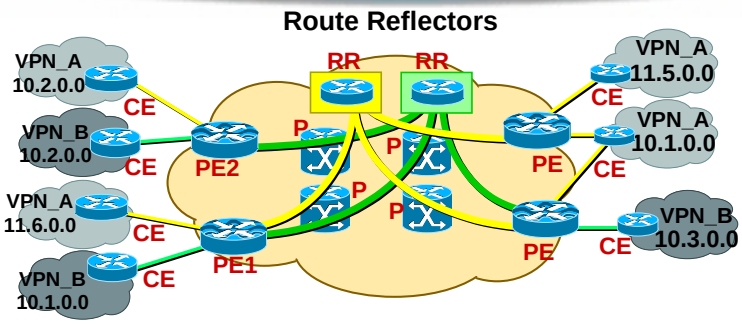

- Route reflectors may be partitioned

  Each RR store routes for a set of VPNs

- Thus, no BGP router needs to store all VPNs information
- PEs will peer to RRs according to the VPNs they directly connect

MPLS-VPN Scaling
BGP Updates Filtering

- iBGP full mesh between PEs results in flooding all VPNs routes to all PEs
- Scaling problems when large amount of routes. In addition PEs need only routes for attached VRFs
- Therefore each PE will discard any VPN-IPv4 route that hasn't a route-target configured to be imported in any of the attached VRFs
- This reduces significantly the amount of information each PE has to store
- Volume of BGP table is equivalent of volume of attached VRFs (nothing more)

MPLS-VPN Scaling
BGP Updates Filtering

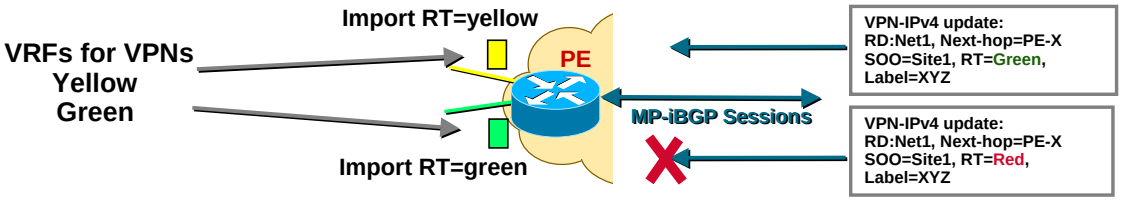

- Each VRF has an import and export policy configured
- Policies use route-target attribute (extended community)
- PE receives MP-iBGP updates for VPN-IPv4 routes
- If route-target is equal to any of the import values configured in the PE, the update is accepted
- Otherwise it is silently discarded

Agenda

- MPLS-VPN Terminology
- Network Topology
- Step By Step Configuration Guideline
  - Sample configuration (VRF, interfaces)
  - PE-CE Routing Protocol configuration (Static,RIPv2, BGP,OSPF)
  - MP-BGP configuration
- MPLS VPN Scaling
  - Route-Reflectors, Filters
- MPLS VPN Internet Routing

MPLS VPN Internet Routing

- In a VPN, sites may need to have Internet connectivity
- Connectivity to the Internet means:
  - Being able to reach Internet destinations
  - Being able to be reachable from any Internet source
- Security mechanism must be used as in any other kind of Internet connectivity

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

MPLS-VPN Internet Connectivity

1) VRF Specific default route

2) Separate sub-interface (non VRF)

May run BGP Between PE and CE

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

MPLS VPN Internet Routing VRF Specific Default Route

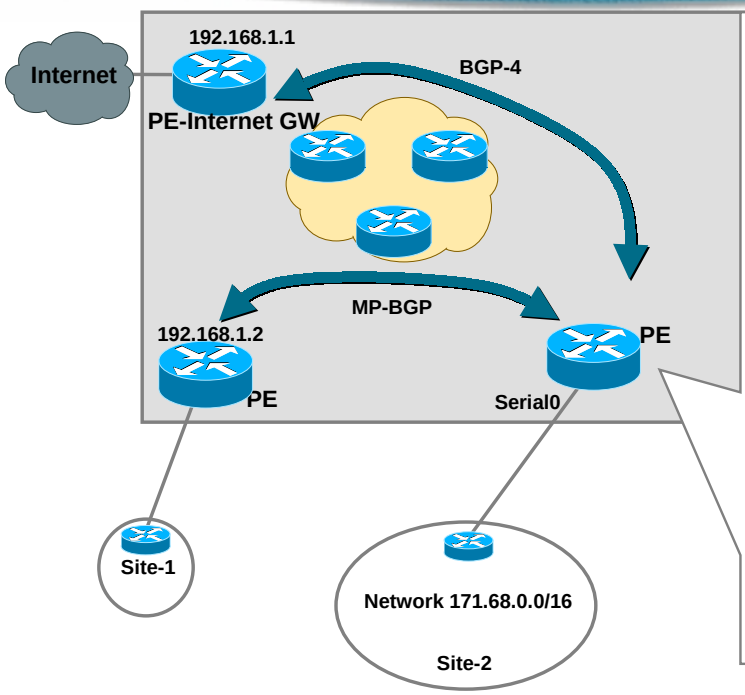


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
!
Interface Serial0
ip address 192.168.10.1 255.255.255.0
ip vrf forwarding VPN-A
!
Router bgp 100
no bgp default ipv4-unicast
network 171.68.0.0 mask 255.255.0.0
redistribute static
neighbor 192.168.1.1 remote 100
neighbor 192.168.1.1 activate
neighbor 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 global
```

2202
1358_06_2000_c3
© 2000, Cisco Systems, Inc.
Cisco.com

MPLS VPN Internet Routing VRF Specific Default Route

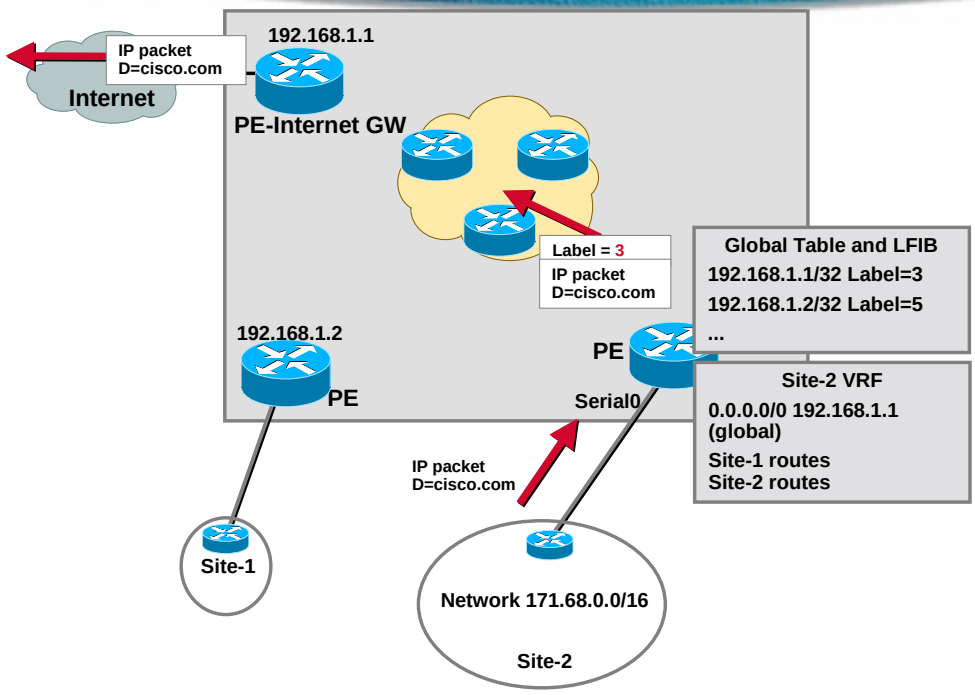

MPLS VPN Internet Routing VRF Specific Default Route

- **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**

MPLS VPN Internet Routing VRF Specific Default Route

- The Internet Gateway specified in the default route (into the VRF) need NOT to be directly connected
- Different Internet gateways can be used for different VRFs
- Using default route for Internet routing does NOT allow any other default route for intra-VPN routing

MPLS VPN Internet Routing VRF Specific Default Route

- PE routers need not to hold the Internet table
- PE routers will use BGP-4 sessions to originate customer routes
- Packet forwarding is done with a single label identifying the Internet Gateway IP address

MPLS VPN Internet Routing Separated (Sub)interfaces

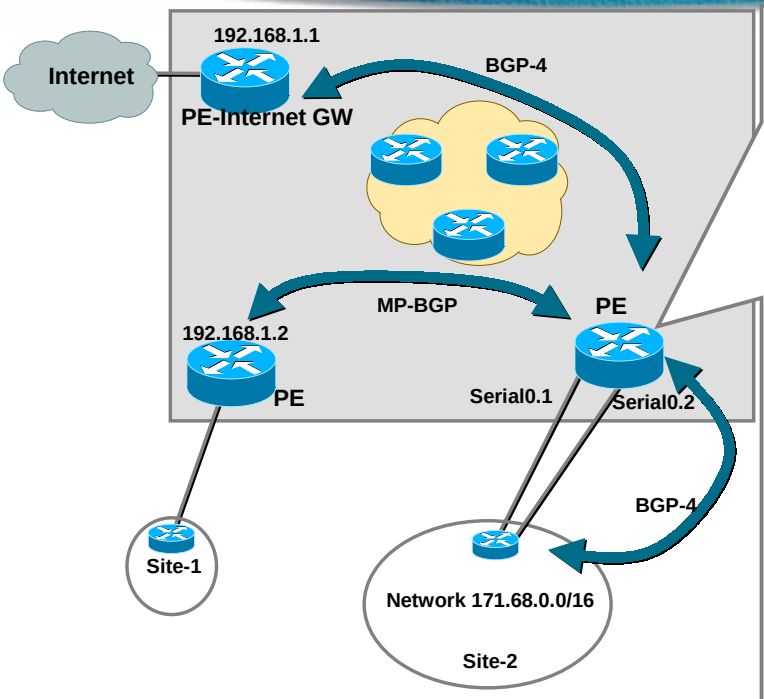


```
ip vrf VPN-A
rd 100:1
route-target both 100:1
!
Interface Serial0
no ip address
!
Interface Serial0.1
ip address 192.168.20.1 255.255.255.0
ip vrf forwarding VPN-A
!
Interface Serial0.2
ip address 171.68.10.1 255.255.255.0
!
Router bgp 100
no bgp default ipv4-unicast
neighbor 192.168.1.1 remote 100
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 192.168.20.2 remote-as 502
 neighbor 192.168.20.2 activate
 exit-address-family
!
address-family vpnv4
 neighbor 192.168.1.2  activate
exit-address-family
```


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**

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**
  - **One (sub)interface for Internet routing**
    - **Associated to the global routing table**

MPLS VPN Internet Routing Separated (Sub)interfaces

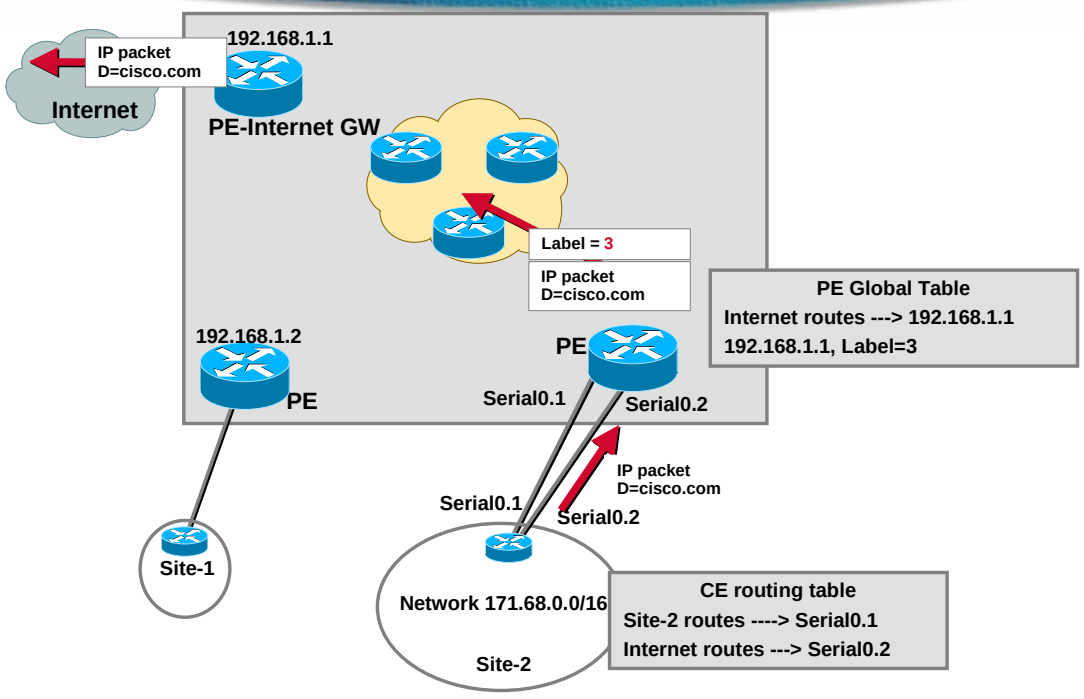

MPLS VPN—Configuration Cheat Sheet

- Identify PE routers
- Configure VRF instance for each VPN
  - Design rd and route-target assignments
- Configure MP-BGP between PEs (Route-Reflector or full-mesh)
- Identify Interfaces going to VPN sites
  - Assign interfaces to the VRF instance
- Choose the routing protocol between PE and CE
  - Configure necessary redistribution between MP-BGP and VRF routing protocol

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, Telnet, Traceroute commands are VRF based
  - ping vrf <vrf-symbolic-name>
  - telnet /vrf <vrf-symbolic-name>
  - traceroute vrf <vrf-symbolic-name>

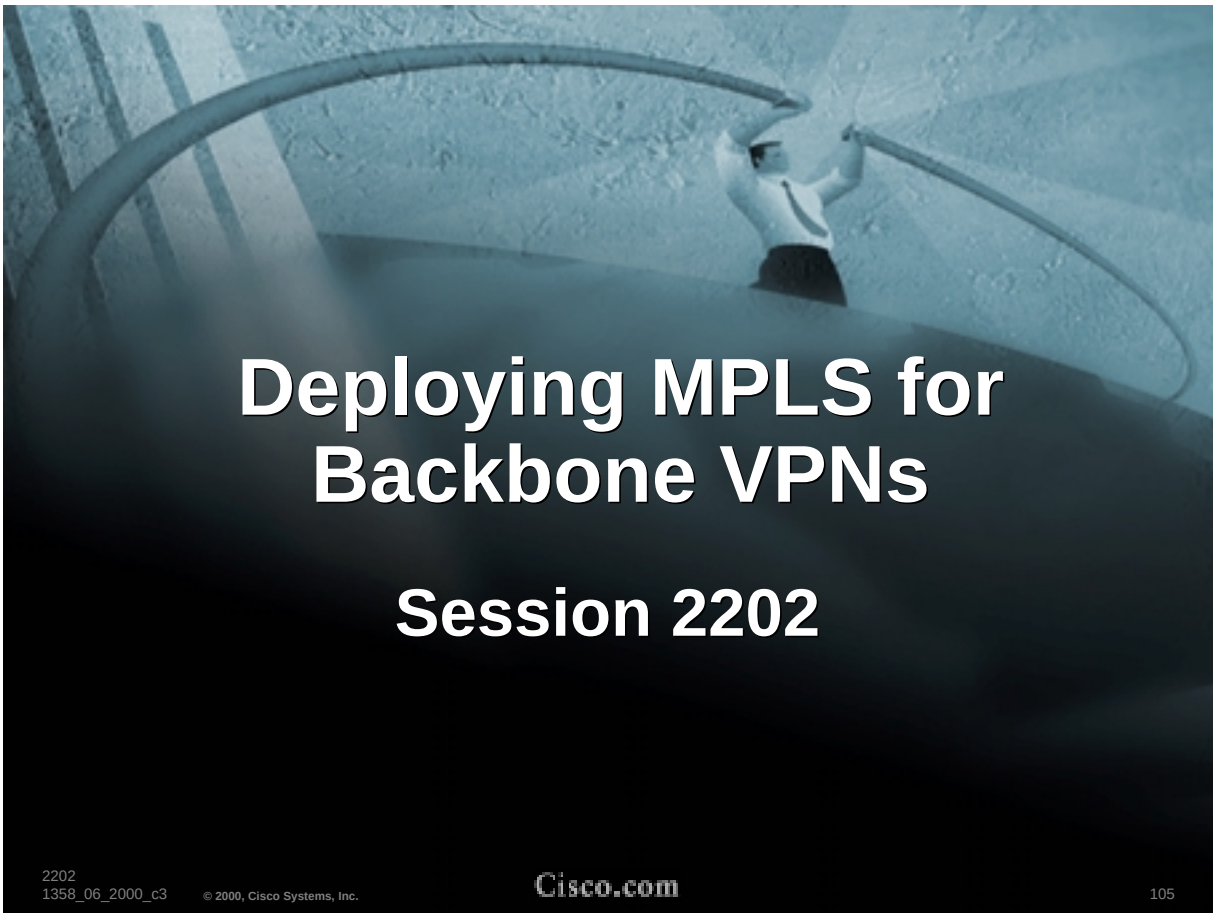

Empowering the Internet Generation℠