

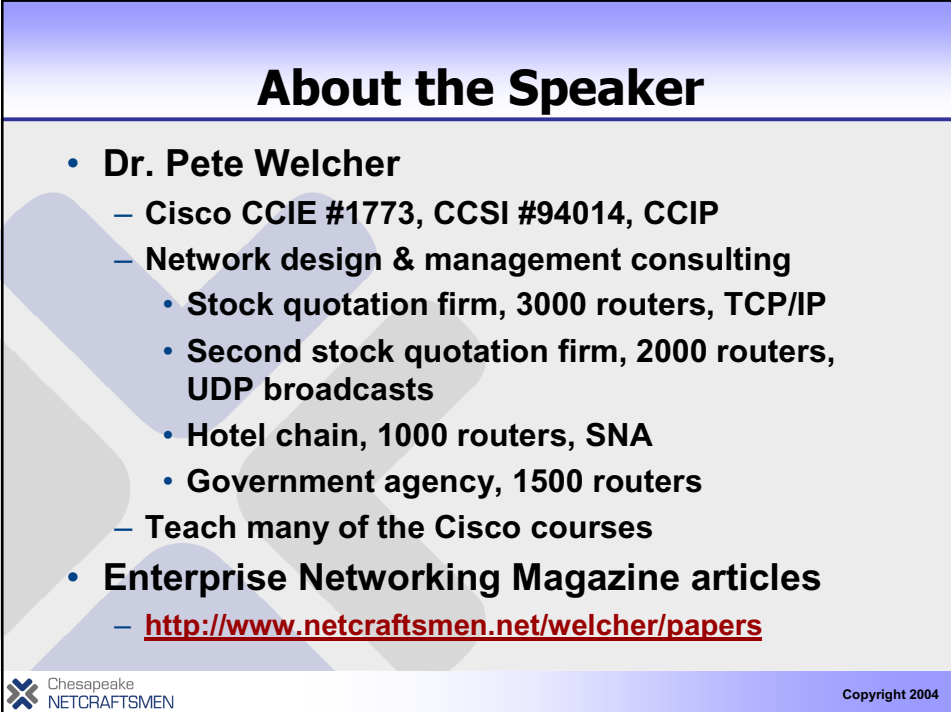
MPLS VPN's

Presented by:
Dr. Peter J. Welcher

Chesapeake
NETCRAFTSMEN

1

Copyright 2004



About the Speaker

- **Dr. Pete Welcher**
 - Cisco CCIE #1773, CCSI #94014, CCIP
 - Network design & management consulting
 - Stock quotation firm, 3000 routers, TCP/IP
 - Second stock quotation firm, 2000 routers, UDP broadcasts
 - Hotel chain, 1000 routers, SNA
 - Government agency, 1500 routers
 - Teach many of the Cisco courses
- **Enterprise Networking Magazine articles**
 - <http://www.netcraftsmen.net/welcher/papers>

Chesapeake
NETCRAFTSMEN

Copyright 2004

Objectives

- **Upon completion of this presentation, you should be able to:**
 - Explain the business case for MPLS VPN's
 - Discuss the basics of BGP, MBGP, and MPLS relevant to MPLS VPN's
 - Understand the fundamentals of how MPLS VPN's work, including
 - VRF's, Route distinguishers, Extended communities
 - MBGP next hop (egress PE router)
 - Use of LSP between ingress PE and egress PE
 - Explain some of the new MPLS VPN features, including **ATOM**

Topics

- **Why MPLS VPN's?**
- **Quick BGP Review**
- **Quick MPLS Review**
- **MPLS VPN Technical Overview**
- **Sample MPLS VPN Configurations**

What is a VPN?

- **Virtual Private Network**
- **Connectivity deployed on a shared infrastructure with the same policies and performance as a private network, with lower total cost of ownership**
- **Key policy: security!**

Existing VPN Technologies?

- **(Legacy) Frame Relay, ATM Virtual Circuits**
- **IPsec**
 - Key management (PKI)
 - Literal hardware cost of encryption at speed
 - Good for access via local ISP
 - SLA's???
- **IP-based**
 - GRE tunnel mesh from ISP
 - IP routing with route filtering by ISP
- **Other**
 - MAN: VLAN's
- **MPLS VPN's**
 - Good for intranet, extranet
 - All your routing is provided by the ISP (?!??)

Why Use MPLS for VPN's?

- Use existing ISP BGP and router infrastructure
- Build Layer 2 or Layer 3 MPLS VPN's on top of the infrastructure
- Layer 2 VPN's: Anything Over MPLS (ATOM)
 - Currently: MAN/WAN point-to-point Ethernet, FR, ATM VC's
 - Future: MAN/WAN multi-point Ethernet connections
- Layer 3 VPN's
 - Basically, private (per-interface) routing tables

Why MPLS VPN's?

- For the MPLS VPN Customer...
 - Service Provider can provide secure private IP full mesh connectivity with dynamic routing
 - Can also get good SLA's and QoS support and low pricing
 - Outsourced connectivity and routing table management
 - Reduce need for in-house skills
 - Reduce need for additional staff by offloading work
 - No need to learn MPLS! All the work is done by the Service Provider.

Why MPLS VPN's?

- **For the MPLS VPN Implementer...**
 - Service Provider
 - Large Enterprise, Government, College
- **MPLS VPN Benefits**
 - Simpler provisioning of VPN's than with hand-crafted route filters or route maps, less error-prone and labor-intense
 - Scalable large-scale IP routing
- **Can layer IPsec on top**
 - Cost of IPsec only where needed
- **CCIE job security** [J](#) [L](#)
 - The simplification means automated tools and simpler provisioning and troubleshooting
 - More complex technology is being used

Business Case for MPLS VPN's

- **Service Provider case:**
 - Build a solid IP networking core and offer services over it
 - Any configuration for a new MPLS VPN is at the edge router connecting to the customer
 - Much more scalable configuration and routing than GRE or IPsec
 - The routing is equivalent to full tunnel mesh

Business Case for MPLS VPN's

- **Enterprise case:**
 - (See above)
 - **Compartmentalize routing for isolation and security**
 - Can route separately by department or business unit
 - Control external routing connectivity per department
 - Provide controlled connectivity at selected firewall points with IDS's

Business Case for MPLS VPN's

- **If and when MPLS VPN's are generally available in the 6500's / MSFC's**
 - Gain the traffic separation of VLAN's, yet use high availability Layer 3 techniques instead of extended Layer 2 VLAN's
 - Provide routing isolation for distributed wireless access subnets

Enterprise MPLS VPN: For Whom?

- **Large enterprises**
 - Especially those with subsidiaries
- **Universities**
 - Isolate academic departments, etc.
 - Control external exposure per-department
- **City, county, state governments**
 - Lower costs due to shared network, yet isolation for police, fire, and other sensitive services
 - Much lower ongoing and support costs than individual departmental networks

MPLS VPN's: Which Layer?

- **RFC 2547: Layer 3**
 - SP or MPLS network participates in the customer routing
 - Good isolation of routing between customers
- **Layer 2 Point-to-point**
 - Ethernet
 - FR, ATM AAL5
 - PPP, HDLC: virtual leased lines w/o TDM
 - (draft-martini-l2circuit-trans-mpls-07.txt).
 - (draft-martini-l2circuit-enp-mpls-03.txt).

Cisco Universal VPN

- **Common architectural approach allowing mixing of VPN types:**
 - IPsec VPN
 - MPLS VPN
 - L3TP VPN
- **Use each where appropriate**

Topics

- **Why MPLS VPN's?**
- **Quick BGP Review**
- **Quick MPLS Review**
- **MPLS VPN Technical Overview**
- **Sample MPLS VPN Configurations**

BGP Review

- **iBGP and eBGP**
- **Next hop attribute**
- **BGP community attribute**

MBGP

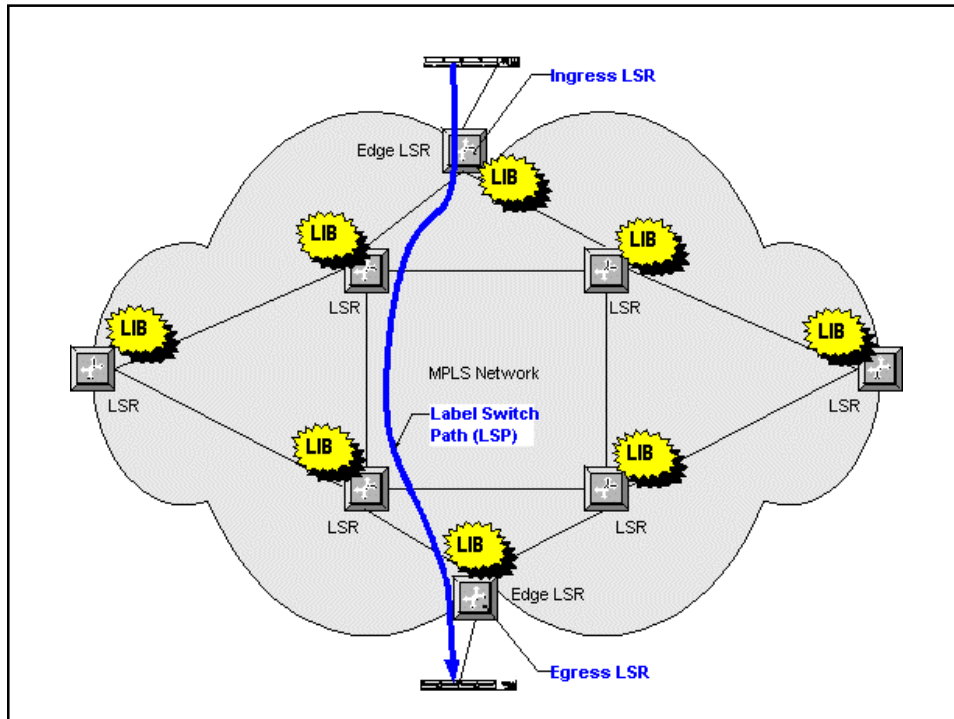
- **Extends BPG to carry information about other protocols**
- **Useful for:**
 - IPv6
 - IPX “internet” (who cares anymore?)
 - ISP-scale tracking of unicast routes for multicast purposes (RPF check “steering” of multicast traffic along different trunks)
 - MPLS VPN routing information (more detail later)

Topics

- Why MPLS VPN's?
- Quick BGP Review
- Quick MPLS Review
- MPLS VPN Technical Overview
- Sample MPLS VPN Configurations

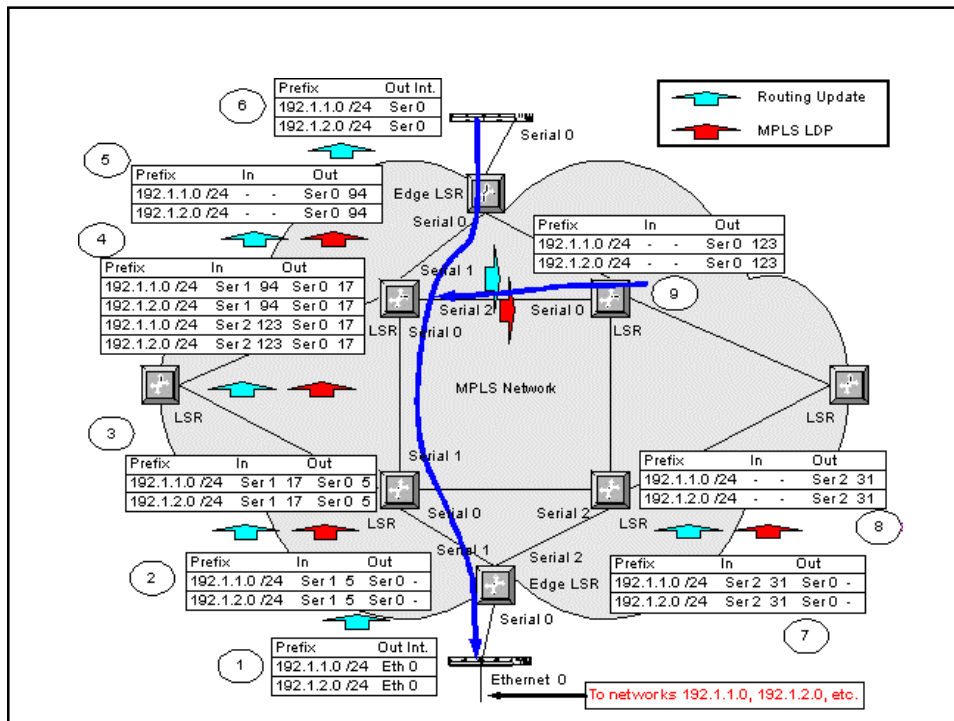
MPLS Terminology

- LSR – Label Switch Router
- LER – Label Edge Router
- Ingress, egress LSR
- LSP – Label Switch Path
- LIB – Label Information Base
- LFIB – Label Forwarding Information Base



Label Switching Process

- Router forwards packet to LER
- LER CEF lookup determines a label is available for the destination prefix
- Ingress LSR applies labels
- Labels used to forward packet through MPLS network
- Egress LSR forwards normal packet



Label Path Establishment

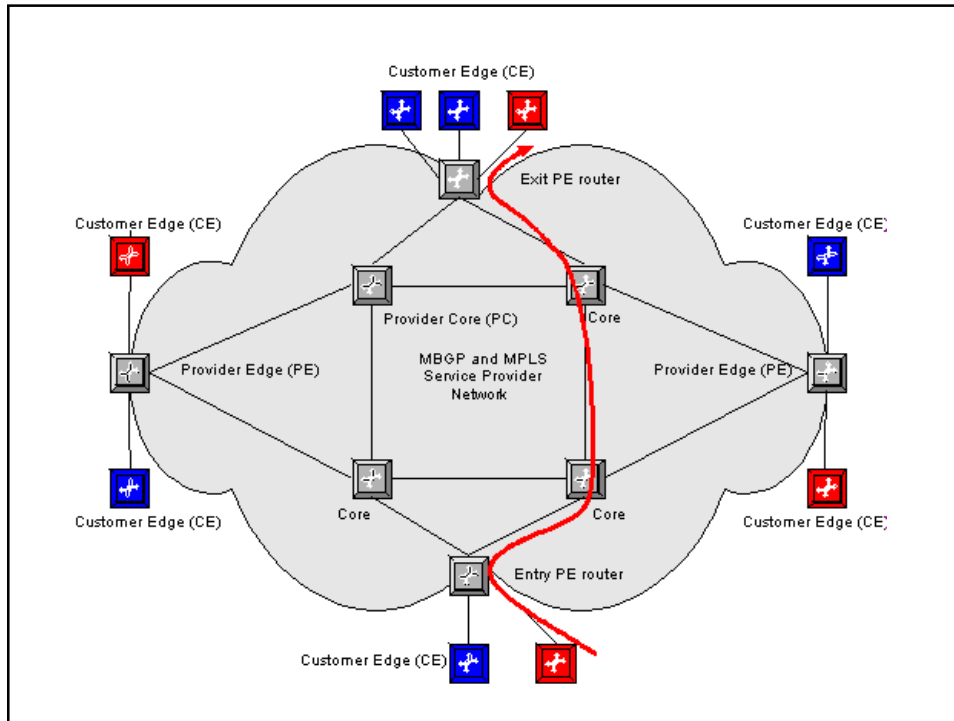
- Routing updates propagate routing information
- TDP or LDP establishes control connection between adjacent LSR's
- LSR advertises a label to upstream neighbor for each new prefix
- Upstream neighbor picks outbound label(s) from next hop router(s) determined by routing protocol, matches inbound label to the outbound label in the LFIB

Topics

- Why MPLS VPN's?
- Quick BGP Review
- Quick MPLS Review
- **MPLS VPN Technical Overview**
- Sample MPLS VPN Configurations

Service Provider Terminology

- Customer Edge router (CE)
- Provider Edge router (PE)
- Provider Core router (P)
- Intranet, extranet
- Entry, exit PE routers



Route Distinguisher

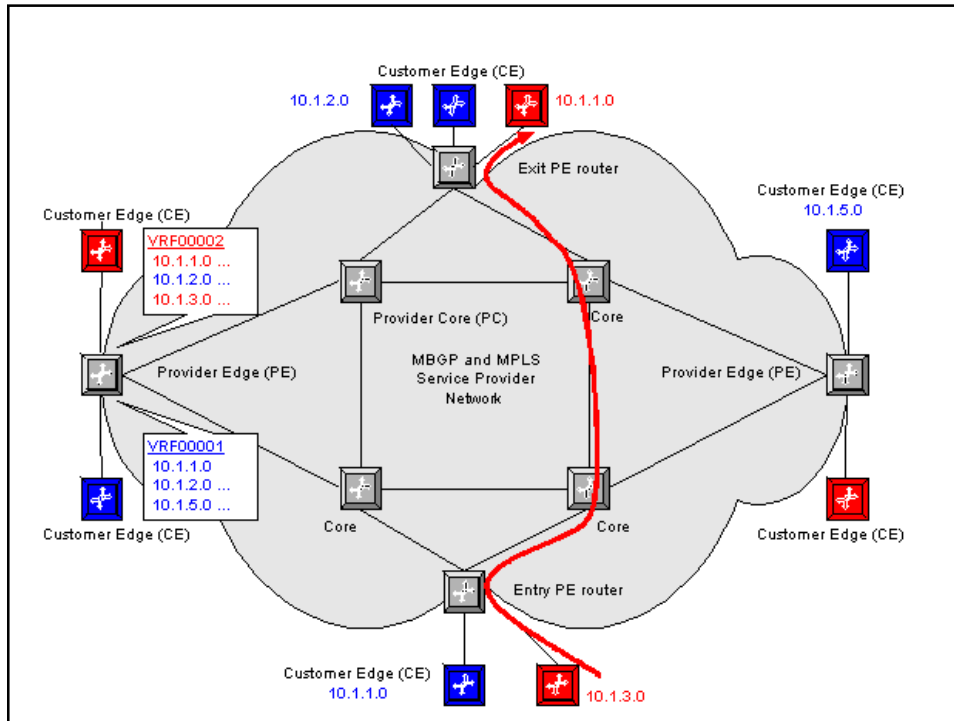
- **Route distinguisher (RD) – 8 byte prefix used by MBGP**
 - **Not** VPN identifier
 - Does allow private addresses from different customers to be distinguished, kept separate
 - As long as those sites don't need to communicate
 - Think of MBGP as distributing routing information for 12 byte prefixes (8 bytes of RD, 4 of normal IPv4 prefix)

Key New Ideas for MPLS VPN's

- VRF – Virtual Routing Facility
 - Per-interface IPv4 routing table on PE router
- MBGP extended community
 - **Not** VPN identifier
 - One of possibly several communities of routes to import to VRF
 - Local VRF routes generally exported to at least one MBGP extended community
- What is reachable from the VRF is destinations whose extended communities were imported
 - Route maps for fine-tuning

An Example Using Extended Communities

- Customer A sites use extended community 100:1
- Their credit card clearinghouse bank's routes are associated with (exported to) community 100:1234
- Their suppliers routes to bidding/quoting servers are exported to community 100:5678
- VRF imports routes from 100:1, 100:1234, 100:5678
- VRF exports local routes to 100:1, and ...
 - **What else do you need?**



Role of MBGP in MPLS VPN's

- MBGP uses address family vpnv4 tracks RD + IPv4 prefixes: MPLS routing info
- MBGP also transmits two critical attributes for MPLS VPN's:
 - Extended communities associated with each vpnv4 routing prefix
 - MPLS label indirectly identifying the appropriate VRF is associated with each advertised prefix

BGP and MPLS VPN's

```
address-family ipv4 unicast vrf name-of-vrf
```

- We generally configure Cisco BGP with VRF-specific instances
- eBGP to a CE router exchanges information with the VRF instance
- The BGP VRF instance can redistribute to/from an OSPF process (one per CE) or a global RIPv2 process
- The VRF BGP instance automatically exchanges information with MBGP vpnv4 routing
 - Adds/removes the RD

MPLS VPN Packet Flow

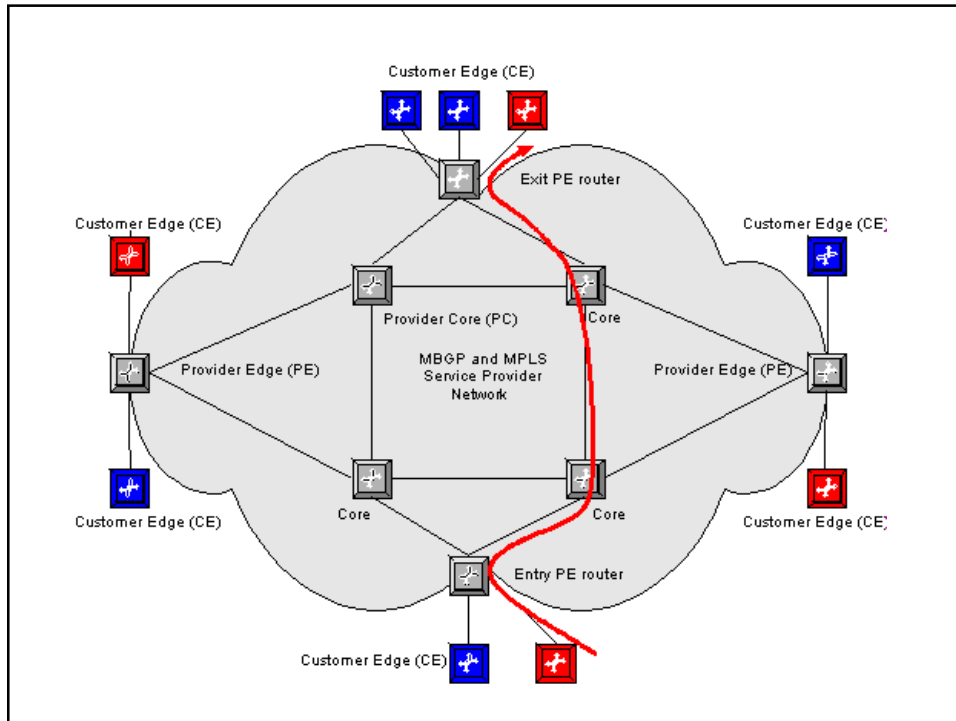
- CE router forwards packet to PE, based on static route or dynamic route to destination.
 - Dynamic route is advertisement of VRF to CE router
 - Usually requires redistribution (unless eBGP used)
- PE router looks at VRF forwarding table associated with inbound interface, applies appropriate RD, determines MBGP next hop
- BGP next hop is reachable via MPLS label path

MPLS VPN Packet Flow - 2

- PE router applies label stack: label for the MBGP next hop, along with a VRF identifying label learned via MBGP
- Labels allow packet to be sent across P routers without examining IPv4 header
 - No need for new routing based on “long” RD + IPv4 addresses
 - No confusion if private address for destination

MPLS VPN Packet Flow - 3

- Penultimate hop LSR pops outer label and forwards packet to exit router (iMBGP next hop)
- Egress PE uses inner label to ensure correct VRF and interface used to continue forwarding packet to CE router
 - Security safety mechanism!



Some Details

- **No MPLS, no MBGP needed on CE router?**
 - No need to upgrade CE routers!
- **CE – PE routing choices:**
 - Static
 - OSPF
 - RIPv2
 - eBGP
 - EIGRP
- **Service Provider does provide the routing for the customer: CE router might well just use static default!**

Scaling Benefit - 1

- **P routers do not need to speak MBGP**
 - P routers only need to run an IGP to provide reachability and MPLS labels to get to all iMBGP next hop addresses
 - Next hops are usually loopbacks on PE routers
- This results in much better scalability / stability in SP core: no customer routes, much smaller routing table

Scaling Benefit - 2

- **PE routers automatically filter routes with communities that are not imported by any local VRF**
- **SP can therefore partition their customers and route reflectors**
 - Keeps size of aggregate routing database smaller
 - Still does require some design and planning, but nothing very complex

MPLS VPN's and IP Multicast

- **Multicast Over MPLS VPN**
 - MPLS can support multicast directly
 - Does not scale to very large sizes well
 - Not appropriate for multicast within a VPN
 - Cisco now has scalable techniques for supporting multicast within an MPLS VPN
 - MPLS labeling can be directly used to support distribution of high-volume multicasts

ATOM Futures

- **Future possibilities**
 - Any-to-any service interworking
 - General ATM cell transport
 - POS over MPLS
 - Multipoint support (MPLS L2 VPN acts like LAN interconnecting sites)
 - Crypto support with ATOM

Topics

- Why MPLS VPN's?
- Quick BGP Review
- Quick MPLS Review
- MPLS VPN Technical Overview
- **Sample MPLS VPN Configurations**

Step 1: Create VRF's

```
ip vrf vrf00001
  rd 888:1
  route-target both 888:1
ip vrf vrf00002
  rd 888:2
  route-target both 888:2
  route-target import 888:1
  import map vrf00002-import-map
```

Step 2: Apply VRF's to Interfaces

```
interface Fastethernet 0/2
  ip vrf forwarding vrf00001
  ip address ...
interface Fastethernet 0/3
  ip vrf forwarding vrf00002
  ip address ...
```

Step 3: MBGP (One Time Only)

- *(Configure Core IGP, probably OSPF or IS-IS)*
- *Set up MBGP between PE's*
 - *You only have to do this ONCE, not per-VRF*

```
router bgp 888
  no auto-summary
  no bgp default ipv4-activate
  neighbor 10.60.0.5 remote-as 888
  neighbor 10.61.0.1 remote-as 888
```

Step 3: MBGP (Continued)

```
...
address-family vpnv4 unicast
  neighbor 10.60.0.5 activate
  neighbor 10.60.0.5 send-community extended
  neighbor 10.61.0.1 activate
  neighbor 10.61.0.1 send-community extended
  no synchronization
  no auto-summary
exit-address-family
```

Step 4: PE-CE Routing

```
address-family ipv4 unicast vrf vrf00001
  neighbor 10.20.1.1 remote-as 65535
  neighbor 10.20.1.1 activate
  no synchronization
  no auto-summary
address-family ipv4 unicast vrf vrf00002
  neighbor 10.20.2.2 remote-as 65535
  neighbor 10.20.2.2 activate
  no auto-summary
  no synchronization
exit-address-family
```

PE-CE Routing Alternative: Static

```
ip route vrf vrf00001 15.0.0.0  
                255.0.0.0 e0/2 10.20.1.1
```

PE-CE Routing Alternative: RIP

```
router rip  
  version 2  
  address-family ipv4 vrf a  
    version 2  
    network 150.1.0.0  
    default-information originate  
    default-metric 1  
  no auto-summary
```

PE-CE Routing Alternative: OSPF

```
router ospf 1 vrf vpna
  log-adjacency-changes
  redistribute bgp 3 subnets
  default-metric 10000
  network 150.1.0.0 0.0.255.255 area 0
```

PE-CE Routing Alternative: EIGRP

```
router eigrp 1
  address-family ipv4 vrf VRF-RED
  autonomous-system 1234
  network 172.16.0.0
  redistribute BGP 1234 metric 10 ...
  exit-address-family
```

CE Configuration

- **STANDARD** routing from the CE perspective

```
router bgp 65001
  network 150.1.0.0
  network 203.1.0.0
  network 203.1.1.0
  neighbor 150.1.31.1 remote-as 3
```

Summary

- **VRF's isolate VPN's**
- **MBGP extended communities (route targets) allow for simple scalable selective connectivity**
- **MBGP extends BGP to carry VRF label and extended addressing information needed for MPLS VPN's**
- **MPLS uses labels across provider core to reach egress MBGP next hop router without IP header examination – MPLS tunnels the VPN packets between PE routers**

Any Questions ?

For more info, see my CiscoWorld articles —
<http://www.netcraftsmen.net/welcher/papers>

 Chesapeake
NETCRAFTSMEN 55 Copyright 2004

A Word About Us ...

- **We can provide**
 - Network design review: how to make what you have work better
 - Periodic strategic advice: what's the next step for your network or staff
 - Network management tools & procedures advice: what's right for you
 - Implementation guidance (your staff does the details) or full implementation
- **We do**
 - Small- and Large-Scale Routing and Switching (design, health check, etc.)
 - IPsec VPN and V3PN (design and implementation)
 - QoS (strategy, design and implementation)
 - IP Telephony (preparedness survey, design, and implementation)
 - Call Manager deployment
 - Security
 - Network Management (design, installation, tuning, tech transfer, services, etc.)



 Chesapeake
NETCRAFTSMEN 56 Copyright 2004

A Word About Us ...



Certified by Cisco in:

- IP Telephony
- Network Management
- Wireless
- Security
- (Routing and Switching)