Enterprise Buyer's Guide to MPLS VPN Services

Dr. Peter J. Welcher,
Chesapeake Netcraftsmen

MPLScon 2006, New York City
May 23-25, 2006

About the Speaker

- Dr. Pete Welcher
  - Cisco CCIE #1773, CCSI #94014, CCIP
  - Specialties: Network Design, QoS, MPLS, Wireless, Large-Scale Routing & Switching, High Availability, Management of Networks
  - Customers include large enterprises, federal agencies, hospitals, universities
  - MPLS w/ major city government optical + MPLS deployment
  - Several large MPLS VPN customers
  - MPLS VPN Security Risk Analysis for major retailer (1700+ stores)
  - Taught many of the Cisco router/switch courses
  - Reviewer for many Cisco Press books, book proposals
  - Presenting (session + labs) on MPLS VPN Configuration at Networkers 2005
- Over 130 articles at http://www.netcraftsmen.net/welcher/

Rationale for This Presentation

- A buyer of MPLS services doesn't need to know a lot technically about how MPLS VPN works
  - You're buying a WAN service
  - All the MPLS is handled by the Service Provider
- There are some key questions to ask, to avoid surprises
  - Make sure you know what you're getting and NOT getting
- Experience shows that your choices will affect how easy or hard your design is
  - Routing, QoS, multicast, etc.
- Our focus here:
  - What's different about the customer side of MPLS VPNs?
  - Compared to consumer of FR and ATM services!

Why MPLS-Based VPN?

- Typical motivation:
  - Because MPLS VPN's are trendy
  - "My boss told me to"
  - Outsource routing, WAN headaches
  - Use internal staff for design and management, not day-to-day operations
  - Virtual full-mesh (within the SP cloud) is good for internal IP telephony, video-conferencing, etc.
  - More bandwidth for less

L2 MPLS VPN

- L2 VPN is basically private "WAN Ethernet"
  - Access: 10/100/1000 Mbps Ethernet, may be rate-limited
  - Service could be point-to-point or multi-point
    - Details later
  - May use Ethernet-switching, optical networking with QinQ, or MPLS L2 VPN "under the hood"
    - How it behaves is what matters
  - Network protocol neutral
  - "Should be" transparent

L3 MPLS VPN

- L3 MPLS uses MP-BGP to provide a virtual full mesh of private routed connections over the provider IP network
  - Access is via Leased line, FR, ATM, IPsec VPN, whatever
  - Big change from traditional WAN: routing is provided for you
    - If you don't want that, use L2 VPN, or use the work-around discussed later
  - You can have any network protocol you want
    - As long as it is IP

Making the Choice: L2 or L3 VPN?

- Do you want to outsource routing and WAN management?
  - Valuable for small organizations, retailers, etc.
  - Some government organizations
  - One indication: small network staff, retention or routing skills issues
- Wish / need to do own routing?
  - Are you a medium to large enterprise, have good net staff, want more control or don't trust SP's routing?
  - Do you have complex routing needs?
- Local availability of L2 VPN
  - Spotty availability, so what can you buy? L3 VPN is easier to deliver at a distance, e.g. via FR or ATM access
- Do you require Single or Dual SP?
  - We're seeing more and more organizations using two SPs, for diversity, avoiding lock-in, etc.
- What kind of management data and reports do you want / need?

L3 VPN: Managed Service, Single SP

- Managed CE router runs RIPv2, OSPF, EIGRP, iBGP to local routers, or uses static routing
- Generally you'll provide specs and SP will implement it

- CE and PE exchange IGP or eBGP routes

- PE adds route distinguisher (RD) to your prefixes, uses route targets (RTs) to control route selection

- Long prefixes (RD + IPv4) plus labels distributed using MP-BGP

- IGP or eBGP passes routes to CE at other sites. Or static routing.

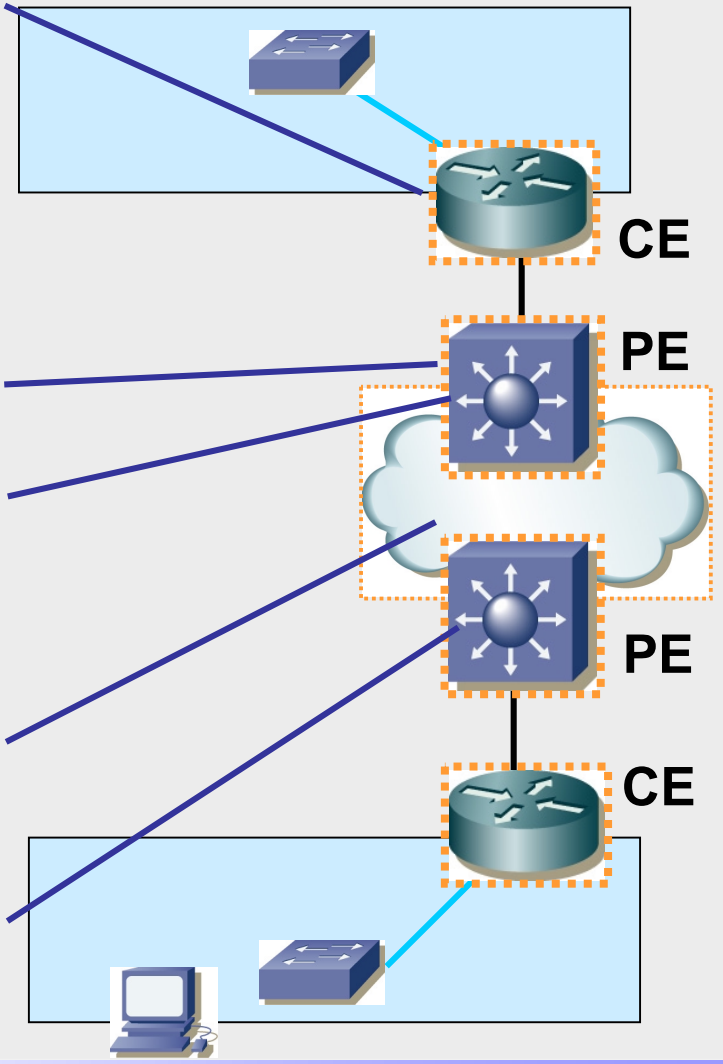

Consideration: Backdoor Routes

- **Suppose you wish to use the old WAN link as a backup path**
- **The problem: external vs. internal IGP routes**
  - **May be ok if eBGP used for routes from SP**
  - **Not so good if redistribution into IGP used and results in external routes**
  - **Route redistribution at many locations does not generally lead to stability…**

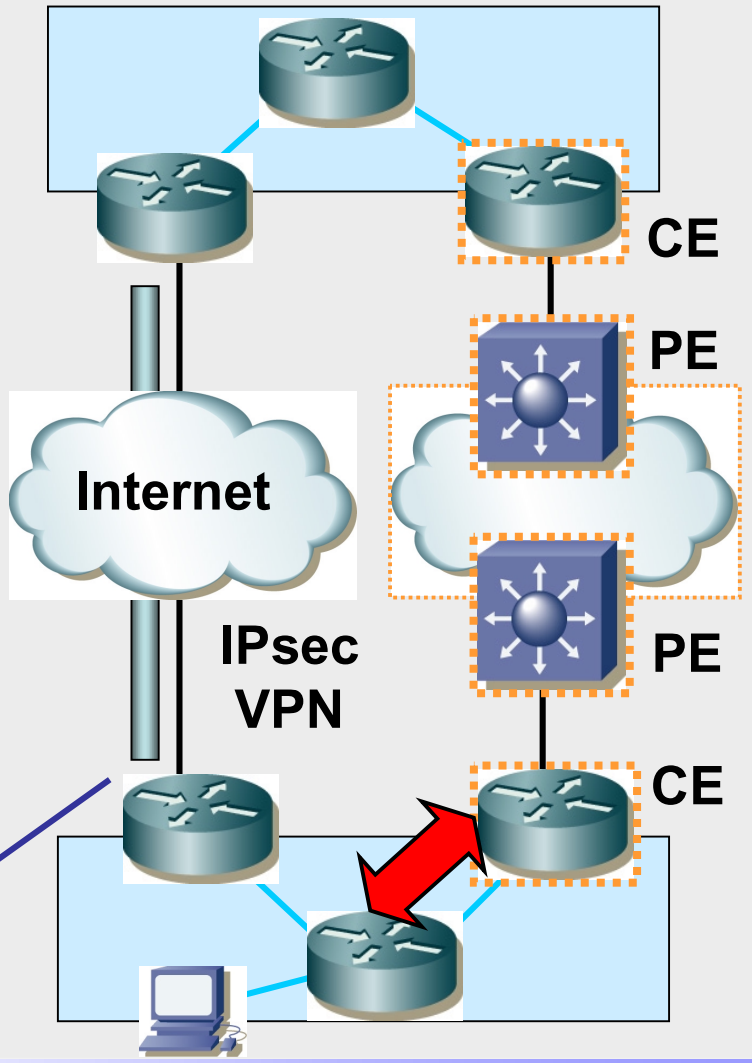

Might also be prior WAN, FR, ATM, etc.

L3 VPN: Unmanaged CE, Single SP, Routing

- **What protocol(s) does your SP offer for CE-PE connectivity?**
- **OSPF and EIGRP features possible with Cisco PE:**
  - **If the SP supports this**
  - **MPLS net can look like OSPF "super-area 0" to you!! (PE acts as ABR)**
  - **OSPF routes *can* remain internal**
  - **EIGRP routes *can* remain internal (with care)**
  - **Need distribute lists with EIGRP when multiple PE's connect to one site**
- **Backdoor route works OK unless MPLS routes become IGP external**

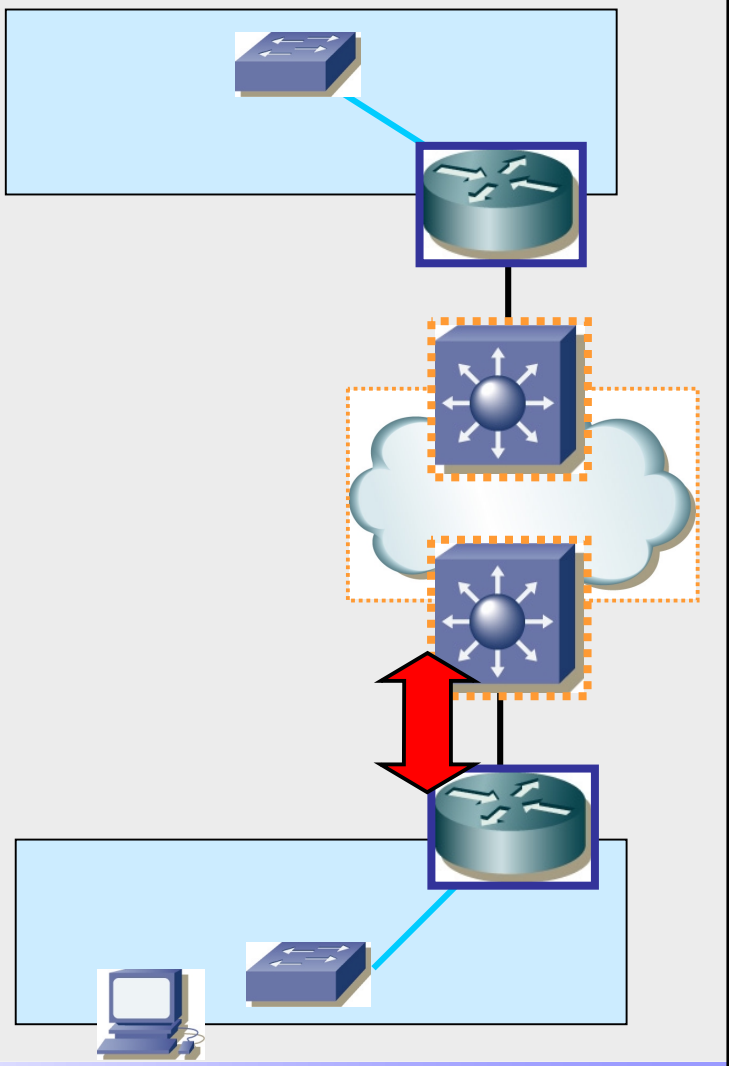

Managed Router + Your Own Routing

- If all you can get is a managed router, but you want your own routing…
- Regard the SP routing as connecting your site routers
  - Cost of a router at each site!
- Can set up GRE tunnels using that connectivity
  - Do your own routing, multicast, etc. over GRE
  - Doesn't clash with backdoor routes
- Possible issues:
  - Encapsulation performance hit
  - MTU: fragmentation performance hit could be major

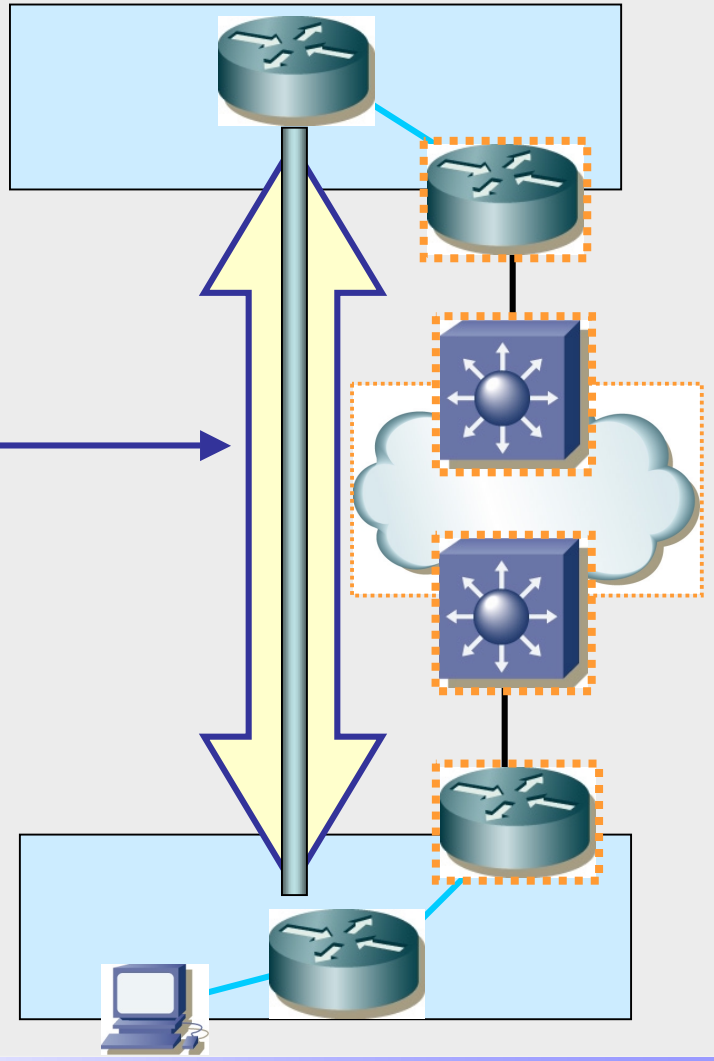

L3 VPN: Managed Service, Dual SP

- First Hop Routing Protocol
- With dual SP's, how do you handle failover?
  - HSRP? VRRP? GLBP?
    - First Hop Routing Protocol (FHRP)?
    - Offered by both SP's?
  - Proxy ARP? Slow!
  - Router at site
    - Defeats purpose of managed service
    - L3 switch ditto
  - Other?
- General issue: need path deadness detection when have 2 SP's
  - iBGP across SP's?
  - Route sync issues?

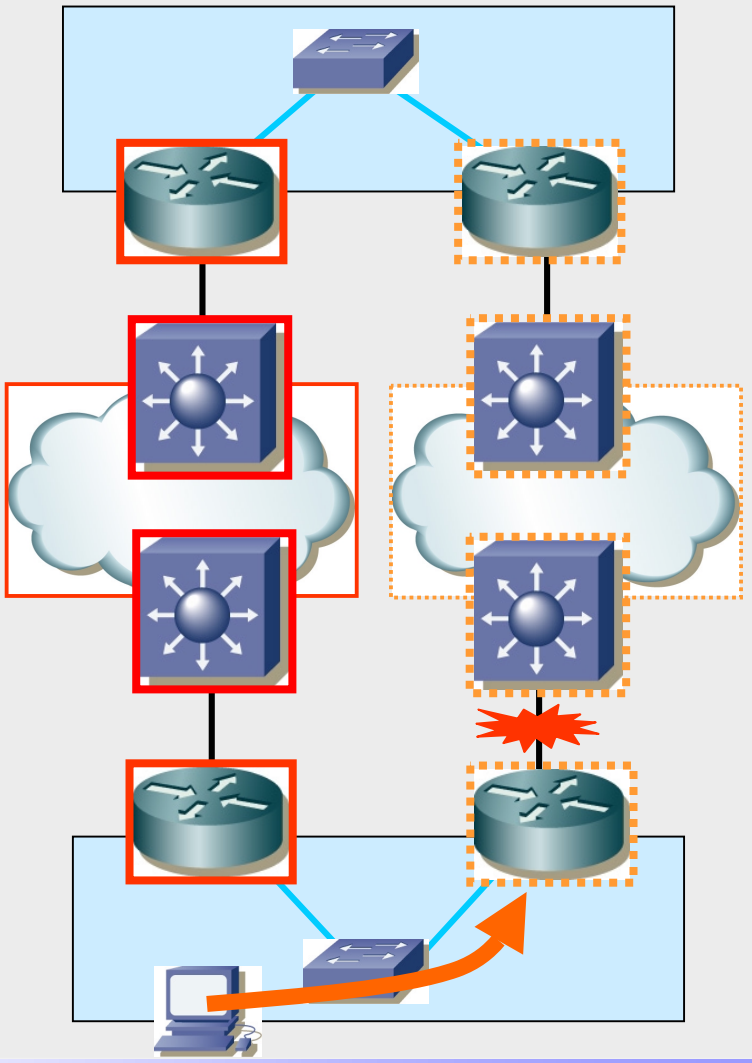

L3 VPN: Unmanaged CE Router, Dual SP

- Use same protocol to both PE's (OSPF, EIGRP)
  - Is there a common PE-CE protocol?
  - Need both doing same thing re internal or external routes
  - Don't want multi-site bi-directional redistribution, e.g. between OSPF and EIGRP
- Use eBGP to both PE's
  - More familiar model
  - Could end up supporting BGP at every site (?!)
- Use HSRP, VRRP, or GLBP (some FHRP) for PC default gateway failover

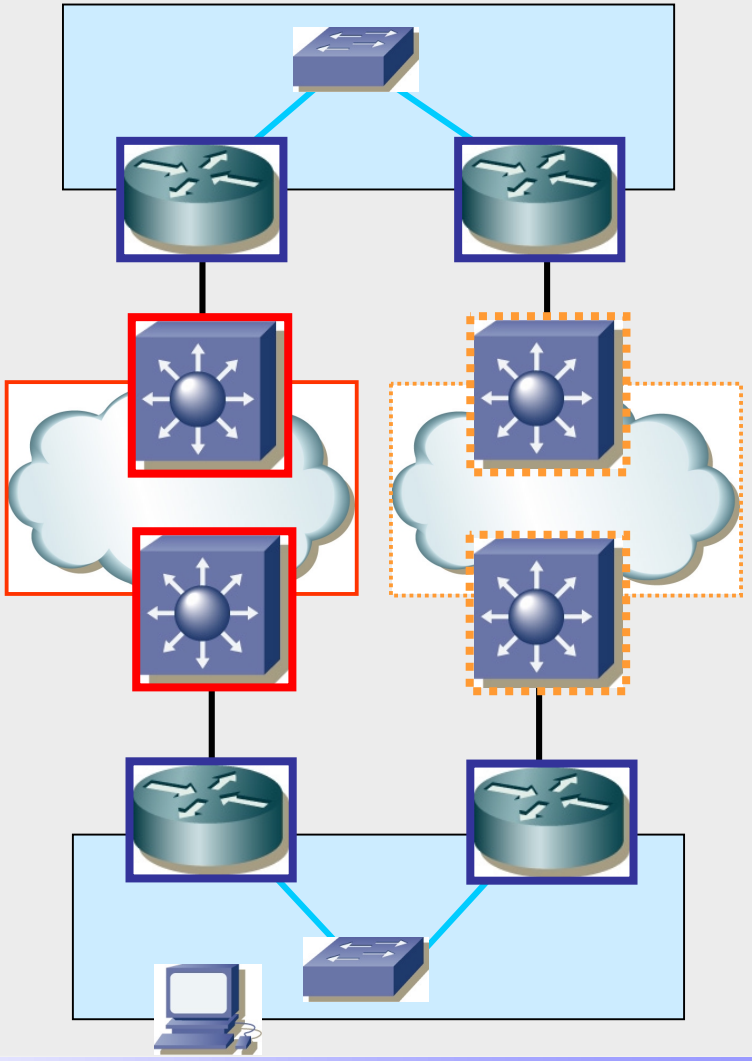

L3 VPN: Other Questions to Ask Your SP

- MPLS L3 VPNs use MP-BGP between PE's
  - How robustly architected?
    - Dual route reflectors?
    - Massive peering?
  - Timers and convergence
    - Essentially have IBGP WAN, doesn't behave like an internal OSPF or EIGRP WAN
    - BGP timers and scan timers can be tuned by SP
- Oversubscription and SLA's
  - What are you guaranteed for bandwidth in SP core?
- Security measures, esp. for managed CE's
  - SP NOC is one potential weak point – due diligence
  - Need due diligence info about SP re-configuration controls, compartmentalization, routing, audit trail, and other security measures

L3 VPN: QoS

- L3 MPLS VPN is a routing full mesh
- Considering doing own QoS?
  - SP cloud drops were possible in typical SP FR or ATM WAN deployment
    - Some SP FR or ATM trunks were designed to run at 50-70% average utilization
  - FR / ATM PE egress drops were mainly due to head-end and trunks being higher speed than remote site – could traffic shape
  - MPLS VPN PE to CE link can carry traffic from several sites – can't traffic shape at CE egress (not pt-pt link)!
  - How do you prevent congestion causing VoIP drops in SP cloud or on PE egress?
  - Conclusion: MPLS VPN QoS needs SP QoS

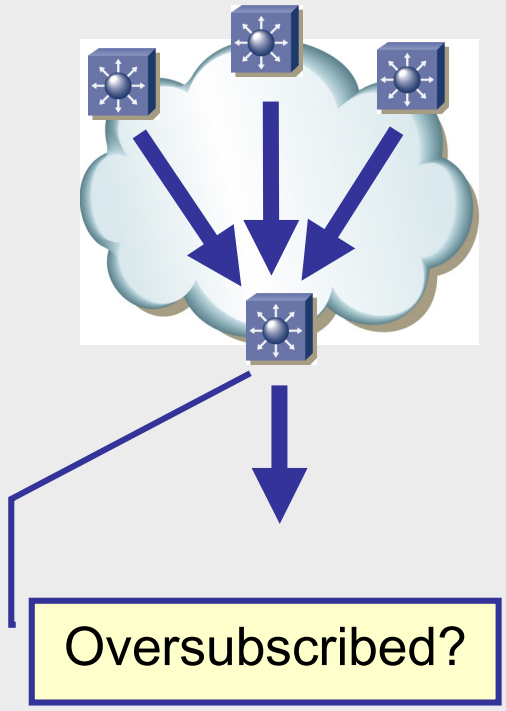

L3 VPN: QoS – 2

- Some sites want to be able to adjust their edge (CE) QoS, possibly at fairly frequent intervals
  - They also want management visibility into the devices
- This is at odds with the usual SP-managed router change control and deployment process
  - Weeks to months to get changes deployed by SP
  - Charges / overall cost
  - Lack of ability to fine-tune configuration commands on the fly: if not exactly right the deployment fails
  - Lack of SNMP or other access to the router/CE device
- One customer is deploying Peribit boxes for this reason (plus the WAN compression benefits)
- Potential market for SP's to offer a co-managed router service?

L3 VPN and Security

- Typical L3 MPLS VPN acts like a full mesh network between sites
  - Traffic routed directly
  - NOT through hub as typical in FR or ATM
- Security impact: central IDS at HQ can't see all traffic
  - Can use hub & spoke MPLS VPN design
  - If your SP provides that
  - Or use stub routing to force traffic flows to HQ
  - Or route summaries for non-voice traffic, assuming VoIP/IPT addressed w/in different prefix (Best Practice!)

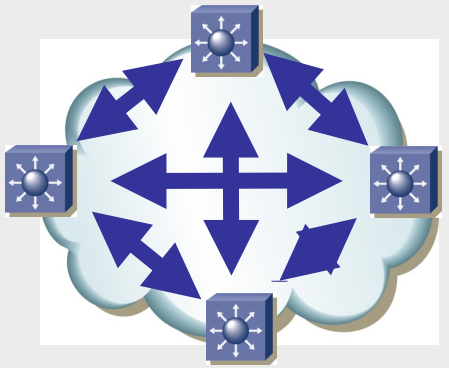

L2 VPN: Different Architectures

- See my online Metro Ethernet article for a taxonomy of L2 VPN LAN-like services…
  - Is it point-point or multipoint? Is it transparent to BPDU's, CDP, etc.?
  - What does it do with multicast?
  - Is it delivered using Ethernet switches or edge devices and optical transport of Ethernet, or using L2 over MPLS technology?
  - Does it use QinQ? Does the SP core use STP, routing, or what?
- Verizon TLS – multi-point Ethernet between sites
  - Now QinQ Cisco 6500-based w/in LATA (offering now ~11 years old)
  - National pt-pt Ethernet MPLS VPN, using QinQ net for local access
- Point-to-point VPLS over MPLS uses a VLAN per pseudo-wire
  - This is somewhat like FR DLCI's, ATM VPI/VCI !!
  - Multicast stays within a VLAN

L2 VPN

- L2 VPN is probably an unmanaged service
- No SP routing, pure L2 (to customer)
- If some customers connect via switches, their STP loops could consume bandwidth on carrier inter-switch trunks
- Question to ask: how does the SP protect your traffic?
  - Do they require customers to connect via routers?
  - Connecting via routers reduces MAC learning burden for PE
  - Do they police inbound traffic?

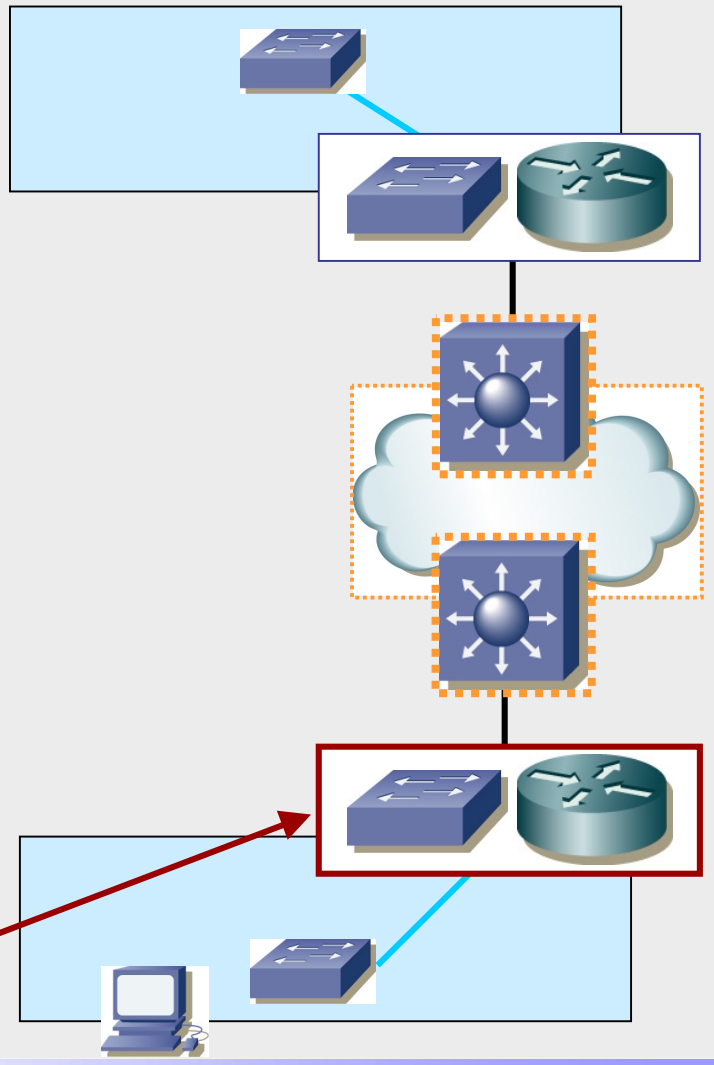

L2 VPN: Dual SP's

- If CE is all switches or transparent multi-point, STP loop protection is up to the customer
- STP blocks redundant path
  - Classic benefit of routing
- So CE should be a router
  - Picture shows 2 CE routers, could use just 1
  - Less redundant, of course
- Redundancy: usual two router techniques

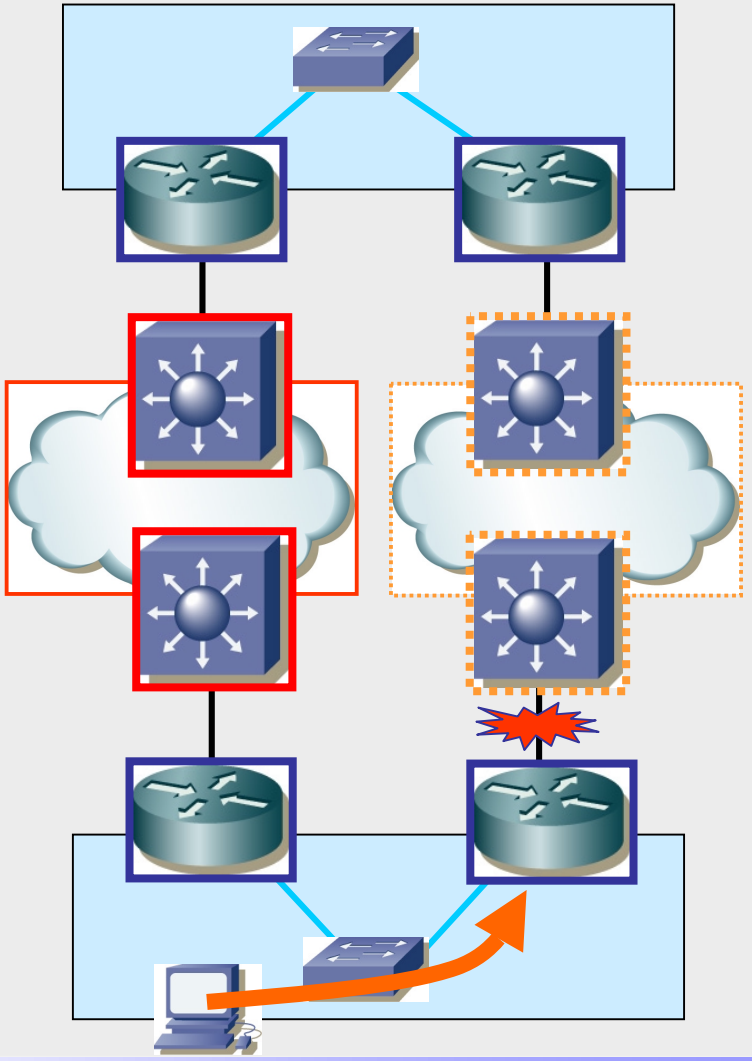

L2 VPN: Possible Issues

- **Design**
  - **We minimize L2 STP in campus designs for High Availability**
  - **Multipoint L2 VPNs can't; multi-enterprise STP … stability?**
- **Routing & Multipoint L2**
  - **Might have too many router adjacencies on the "WAN LAN"**
- **QoS**
  - **See L3 QoS above,** multi-point means need SP QoS
  - Vendor QoS features for Gigabit ports may be limited
- **Multicast**
  - **QinQ or transparent multipoint L2 VPN floods multicast to all CE's**
    - **Even if you're using VLANs to enforce hub-spoke behavior for unicast**
  - **Can use GRE tunnels instead if want point-to-point behavior**
    - **But: possible fragmentation CPU problem, IPmc DoS**

Perspective

- Trend seems to be going towards increased mix of L2 and L3 technologies (hybrids)
  - I/we don't trust Spanning Tree even for Enterprises
  - Mis-handled dynamic routing can also become a problem!
- L3 MPLS VPN may be accessed via a L2 technology
- Ultimately, what matters is stability of the SP network and how fast problems get fixed – and the cost
  - Same as FR, ATM – betting on ability of SP to make it work well
- To evaluate objectively, we'd need to know what technologies the SP's are using, robustness of their designs, plus outage frequencies and durations for each
  - That's NOT going to happen
  - May see reports of conspicuous / long-lived outages
  - May compare our own data if we have two SP's

Summary

- There's a great new world of less costly bandwidth
  - Especially NYC and Long Island, Washington DC, Baltimore, and major urban areas
  - Example: Verizon is now offering Ethernet access to its L3 MPLS VPN services
- Do your homework:
  - Some technical issues to think through
  - Some questions to ask the SP (repeatedly)
- If you need QoS, you have to ask for QoS, and be prepared to pay for QoS
  - Can't DIY (Do It Yourself) as you could with pt-pt WAN links
  - Your traffic is subject to statistical muxing within the SP cloud
- No time to talk about managing MPLS VPN services, tools, SLA's
- Thanks for coming!

Any Questions?

- **For a copy of the presentation, email me at pjw@netcraftsmen.net**
- **About Chesapeake Netcraftsmen:**
  - **Half of our technical experts possess a CCIE**
  - **7.6 Cisco certs per person on average**
  - **Cisco Specializations:**
    - **IP Telephony**
    - **Network Management**
    - **Wireless**
    - **Security**
    - **Routing and Switching**
  - **Expertise in other areas as well**

Cisco Systems | Premier Certified Partner