14b. MPLS VPNs

MPLS (MultiProtocol Label Switching)—A technology tailored to the needs of service providers and the internet backbone. It's marketed to as a service which delivers OSI L_3 packets, e.g. IP, between customer locations. Invisible to the client, MPLS tags packets with a small label of its own which it uses for quickly and easily switching the data through the service provider's network. At the other end, the label is removed, and the packet delivered to the customer.

The service provider *will* need to know about the customer's subnets in order to make forwarding decisions and will use a routing protocol, e.g. BGP, to learn them.

In a Cisco diagram, you can expect MPLS to appear as routers in a cloud, reminding you that it is sold as a layer 3 service.

MPLS VPN (Virtual Private Network)—Even though a service provider has many customers, each appears to have a private L_3 connection between sites. Often, the term "VPN" refers to encryption to provide privacy across a public medium, like the internet. In this case, there's no encryption, just separation from the service provider's other customers. Encryption is in Chapter 15.

TOPOLOGY

CE (Customer Edge)—Typically a router; connects to the link leading to the service provider

PE (Provider Edge)—The device at the service provider that connects to links to customers

Because the edge devices are routers, a wide variety of L_1 + L_2 technologies (Metro Ethernet, leased line, etc.) can be used for the "last mile" link between the customer and the provider. The edge routers will discard any L_2 encapsulation anyway—routers always do. It's even possible to make different choices at opposite ends of the VPN, based on availability and cost in each city.

QoS (Quality of Service)—MPLS has the ability to prioritize "kinds" of customer packets. For example, if you're using IP telephones, you can mark their packets as more time-critical than your OSPF updates and the service provider's network will respect that.

ROUTES ACROSS MPLS

Because MPLS VPNs are sold as an L_3 service, the provider network needs to know about the customer's subnets. A wide variety of routing protocols are usually supported for the access link.

- CE and PE routers on opposite ends of the access link become routing protocol neighbors
- The provider network will *redistribute* the routes into its own routing protocol, usually MPBGP [see below], and share them across the PE routers
- Remote CE routers learn routes from their own local PE routers. CE routers don't directly become neighbors with each other.

Redistribution—Translating a route from one routing protocol into another.

MPBGP (MultiProtocol BGP)—has the ability to keep subnets learned from one customer separate from subnets learned from another. This allows the service provider to run an L_3 VPN instead of an internet.

OSPF WITH MPLS

Combining the fact than an MPLS VPN is an L3 service instead of just a link with OSPF's normal needs, you end up with a few requirements to meet when you design your OSPF areas.

- Every area that isn't area 0, the "backbone," must connect to it
- Area 0 must be contiguous (in one piece)
- The service provider's MPLS VPN routers will need to be in some OSPF area

If different sites are in different areas—useful for summarizing routes to keep routing tables small—this really only leaves one choice. The provider gets area 0.

OSPF Super Backbone—The OSPF area that includes the PE routers. It becomes area 0 and the customer can choose to add more of their own (contiguous) routers to it. This was added to OSPF to cope with L3 WANs like MPLS, where the CE routers can't simply be neighbors across the provider's link, ignoring the provider's equipment in between.

EIGRP WITH MPLS

While OSPF groups its routers into areas, EIGRP groups its routers into Autonomous Systems (AS). Once again the issue arises that when using an L3 WAN, you're dropping someone else's network into the middle of yours in a way that it can't be ignored by your design.

Normally, your network would be one AS, but with the provider's MPBGP network chopping it in half, you can suddenly have different sites in different autonomous systems.

If you choose to keep the same AS number at each of your sites, route redistribution in and out of MPBGP will automatically heal the MPLS division and hide the complexity of the provider's network, keeping your EIGRP metrics looking much as they would have over an L2 link.