



CHAPTER 3

WAN Core—MPLSoL2 Service


As stated earlier WAN Core typically consists of connecting large corporate islands such as MANs or large campuses. These are typically dedicated high bandwidth point-to-point connections. If these large islands are already MPLS enabled, then it is recommended to convert the edge devices into Provider (P) routers. This creates an integrated but flat MPLS network.

Platforms

The edge platforms should be capable of performing a P role, beyond which the selection is based on the interface/throughput requirements. Thus any Cisco platform above the 7200VXR series can be used as the edge router.


In [Figure 3-1](#), we have a combination of GSR, 7600, and 7200 connecting to two large MPLS networks. Each network has its own sets of P, PE, and RR routers. But by connecting the P routers together as shown, we can create a single integrated network where the VPN traffic can access both networks seamlessly. No major changes are required in the existing networks. Functionally, these edge routers act like the other P routers in their respective locations and label switch traffic between locations.

Figure 3-1 MPLSoL2



While our test network only had two locations connected together, as shown in [Figure 3-2](#), any number of networks could be connected by connecting the edge P routers. In certain cases, the edge P router could become transit for other sites. In such case care should be taken to ensure that the platform and the links are properly scoped for such scenarios. Additionally, enough redundancy should be built in to account for any single point of failure. Having multiple links between two sites allows the traffic to be load balanced. The links can be chosen based on the IGP metric to the PE next-hop address. In case of a single link failure, the traffic should reconverge and all traffic should transit the remaining up links.

Figure 3-2 MPLSoL2—Multi-Site Topologies




The only change that may be required is extending the IGP across the two networks. Since the assumption is that the two networks are under the same administrative control, the IGPs should be the same but may require some redesign. The ultimate goal is to provide reachability between the PE next-hop addresses. Typical recommendations are to allocate a subnet address with netmask of /32 from the same /24 subnet and not summarize it anywhere in the network. This ensures that a label is assigned for each /32.

Another change that may be required is in the Route Reflector placement or configuration. Figure 3-3 shows two of the options:

- Option A has a RR per MAN but every PE is peered to both the RRs. This provides RR redundancy where a loss of RR does not affect the VPNs. The RRs do not need to be peered to each other.
- Option B has two RR per MAN. The PEs in each MAN peer with their local RRs. To be able to learn the VPN routes from the other MAN, the RRs are peered with each other in a full mesh to provide additional redundancy. This option creates an additional level of hierarchy that can have an impact on convergence (more RR hops through which the VPN route needs to be propagated).

Option A provides the simplest implementation while option B provides better scale and partitioning capabilities that may be required when servicing large number of PEs and VPN routes.

Figure 3-3 Route Reflector Peering

**Note**

While in large networks (SP-type) it is recommended to have dedicated RRs which are placed out of the forwarding path, in smaller enterprise networks it is possible to use an existing PE as a RR. Care should be taken to scope the additional processing overhead on the PE.

The basic P configurations have been discussed in phase 1 of the design guide and are applicable here as well.

For multicast, the recommendation is to use MVPN as discussed in the phase 1 guide. If the VRFs are using anycast RP, then they should be configured in each of the sites to provide local accessibility.

QoS recommendations from phase 1 hold true here as well. The only additional capability that may be required is the ability to shape the outgoing traffic. For example, in Figure 3-1 P1 and P3 are connected via GE ports, but the underlying service may only be a sub-rate GE. Thus outgoing traffic on both sides needs to be shaped to match the sub-rate.

There could be scenarios that may require the use of Inter-AS to connect the two MPLS networks. This could be when the two networks are part of two different administrative domains or are large enough that a flat network is not recommended. Since we have not seen such requirements yet from enterprise networks, this is not discussed here but future versions may be updated to include it if required.

