



## CHAPTER

# 1

## Problems/Solution Description

Cisco enterprise customers have in the past relied heavily upon traditional WAN/MAN services for their connectivity requirements. Layer 2 circuits based on TDM, Frame Relay, ATM, and SONET have formed the mainstay of most low-speed WAN services. More recently, high-speed MAN solutions have been delivered directly over Layer 1 optical circuits, SONET, or through the implementation of point-to-point or point-to-multipoint Ethernet services delivered over one of these two technologies.

Today, many enterprise customers are turning to Multiprotocol Label Switching (MPLS)-based VPN solutions because they offer numerous secure alternatives to the traditional WAN/MAN connectivity offerings. The significant advantages of MPLS-based VPNs over traditional WAN/MAN services include the following:

- Provisioning flexibility
- Wide geographical availability
- Little or no distance sensitivity in pricing
- The ability to mix and match access speeds and technologies
- Perhaps most importantly, the ability to securely segment multiple organizations, services, and applications while operating a single MPLS-based network


Although service providers have been offering managed MPLS-based VPN solutions for years, the largest enterprise customers are now beginning to investigate and deploy MPLS in their own networks to implement self-managed MPLS-based VPN services. The concept of self-managed enterprise networks is not new; many enterprise customers purchase Layer 2 TDM, Frame Relay, or ATM circuits and deploy their own routed network for these circuits. The largest of enterprise customers even manage their own core networks by implementing Frame Relay or ATM-based switching infrastructures and “selling” connectivity services to other organizations within their companies.

Both of these solutions have had disadvantages; deploying an IP-based infrastructure over leased lines offers little flexibility and segmentation capabilities that are cumbersome at best. Deploying a switched Frame Relay or ATM infrastructure to allow for resiliency and segmentation is a solution within reach of only the largest and most technically savvy enterprises.

As noted, the self-managed MPLS-based network is typically reserved for larger enterprises willing to make a significant investment in network equipment and training, with an IT staff that is comfortable with a high degree of technical complexity. A self-managed MPLS VPN can be an attractive option if a business meets these requirements and wants to fully control its own WAN or MAN and to increase segmentation across multiple sites to guarantee delivery of specific applications. The level of security between separated networks is comparable to private connectivity without needing service provider intervention, allowing for consistent network segmentation of departments, business functions, and user groups.

Corporations with a propensity for mergers and acquisitions benefit from the inherent any-to-any functions of MPLS that, when the initial configuration is completed, allow even new sites with existing networks to be merged with the greater enterprise network with minimal overhead. Secure partner networks can also be established to share data and applications as needed, on a limited basis. The self-managed MPLS is also earning greater adoption as an important and viable method for meeting and maintaining compliance with regulatory privacy standards such as HIPAA and the Sarbanes-Oxley Act. A typical description of this model is “an enterprise acting as a service provider.” Figure 1-1 shows a typical self-managed MPLS MAN deployment.

**Figure 1-1** Typical Self-Managed MPLS MAN Deployment



The following chapters of this guide:

- Explore the technologies necessary to implement a self-managed MPLS-based VPN.
- Describe the evolution of a traditional IP-based enterprise network into an MPLS-based segmented MAN network.
- Discuss the implementation of advanced features such as high availability, QoS, security, and common network services such as NAT, DNS, DHCP, and messaging.
- Explore the management of MPLS VPNs.
- Describe key MPLS-based VPN services such as multicast VPNs and Ethernet over MPLS pseudowires.

143057

- Describe the test bed, test scenarios, and configuration guidelines/caveats associated with recent testing of the MPLS-based VPN MAN topology based on the reference topology described in [Figure 1-1](#).

## Deploying VPNs

While the technology enables you to create the logical separation across networks, it is important to understand the reasons for creating these logical networks. Enterprise customers increasingly require segmentation for a number of different reasons:

- Closed User Groups (CUG)—The CUGs could be created based on a number of different business criterias, with guest Internet access for onsite personnel being the simplest example. Providing NAC/isolation services also creates a need to separate the non-conforming clients. While this can be done using VLANs within a Layer 2 campus network, it requires Layer 3 VPN functionality to extend it across Layer 3 boundaries. CUGs could be created with partners, either individually or as a sub-group, where the segmentation criteria are resources that are to be shared/accessed. This simplifies the information sharing with partners while still providing security and traffic separation.
- Virtualization—Segmentation to the desktop is driving virtualization in the application server space. This means that even existing employees can be segmented into different CUGs where they are provided access to internal services based on their group membership.
- Enterprise as a Service Provider—With some of the Enterprise networks expanding as their organization expands, IT departments at some of the large Enterprises have become internal Service Providers. They leverage a shared network infrastructure to provide network services to individual Business Units within the Enterprise. This not only requires creating VPNs, but also requires the ability of each of the BUs to access shared corporate applications. Such a model can be expanded to include scenarios in which a company acquires another company (possibly with an overlapping IP addressing scheme) and needs to eventually consolidate the networks, the applications, and the back office operations.
- Protecting critical applications—Another segmentation criteria could be based off the applications themselves rather than the users. An organizations that feels that its critical applications need to be separated from everyday network users can create VPNs for each or a group of applications. This not only allows it to protect them from any malicious traffic, but also more easily control user access to the applications. An example of this is creating separate VPNs for voice and data.

Beyond the segmentation criteria, the overarching considerations should be based on the need to share. The VPNs create a closed user group that can easily share information but there will always be the scenario that requires sharing across the VPNs. For example, a company-wide multicast stream would need to be accessible by all the employees irrespective of their group association. Thus the VPNs should be created based on practical considerations that conform to the business needs of the organization.

**■ Deploying VPNs**