Comparing Private Line, Frame Relay, ATM, Ethernet and IP VPNs

Executive Summary

Implementing the best-performing, most cost-efficient mix of WAN services requires that IT departments compare the characteristics of a broad array of physical network types and network service alternatives. Different networks work best depending on the geography of the enterprise and the application(s) at hand. A key decision is at what Open Systems Interconnect (OSI) layer – or operational function – an enterprise wishes its WAN service provider to participate. There are compelling reasons to use Layer 1, 2 and 3 services, and a given enterprise might use just one or a mix of service types. This paper takes a comprehensive look at each of the layers, the service types that accompany them and a number of sample use scenarios.

at&t

A number of access and backbone services are available to accommodate enterprises' many and varied connectivity requirements. There are decisions to be made about physical and service access connection types, as well as backbone service choices. Some networks will be selected to connect users to corporate data resources, while other services will serve different applications, such as machine-to-machine communications. Part 1 of this paper identified situations that might motivate an enterprise to upgrade or change WAN services as well as help the enterprise recognize when an upgrade may not be needed.

WAN Services by 'Layer'

- **Layer 1: Private line services**
- **Layer 2: Packet-switched frame relay, ATM and Ethernet services**
- **Layer 3: Routed IP services**

As discussed in "Understanding Your WAN Choices," Open Systems Interconnect (OSI) "layers" specify technical functions, not branded WAN services. Though OSI technical layers and network services are not synonymous, the terms tend to be used interchangeably. This often complicates the task of sorting out what are enabling technologies and what are actual services that are built on them. This paper will explain those differences and then put them into context as they apply to creating an effective overall WAN service strategy.

Layers, Functions and Service Options

A key WAN service decision for the enterprise to make is at what OSI layer it wishes its WAN service provider to participate in its access and backbone network operations. Generally speaking, the lower the layer number, the less the carrier will be participating and the fewer enhanced services the provider will be supplying.

Layer 1

Layer 1 services imply that the carrier will be providing a network pipe – a point-to-point transmission medium only. To connect any pair of sites, a Layer 1 access service at each site linking to a Layer 1 long-haul private-line circuit in the middle must be physically set up for an end-to-end connection with bandwidth dedicated exclusively to that connection (see Figure 1). This works well for small networks or pockets of high-speed communications needed directly between two sites. The reason is that the network remains relatively simple, Layer 1 networks are easy to troubleshoot, and the entire communications circuit can be used exclusively by the enterprise.

Layer 1 services are often considered highly reliable and secure for the basic reason that the equipment supporting them is mature, comparatively simple and capacity is dedicated to one entity. Because carriers are not providing equipment that participates at the other network service layers, there are fewer components between point A and point B to "break" or open up security vulnerabilities. These Layer 1 services characteristics are why many of the world's trading networks use private lines between brokerage and stock exchanges.

For highly meshed enterprise networks, private lines grow more expensive as the number of sites to be connected and the distance between those sites increases. The total number of enterprise network site-to-site connections can be calculated using the formula $n(n-1)/2$, where n represents the total number of sites that need full-mesh interconnection. A five-node network, then, would need 10 circuits; a 30-node network would require 435 circuits.

Another consideration is private line pricing (specifically, T1/T3 pricing) is largely distance-sensitive: the farther away two sites are from one another, the greater the connectivity service fee is per month. The distance-oriented pricing model of private lines has driven many hierarchical network designs, whereby enterprises connect multiple sites to a relatively nearby aggregation point and multiple aggregation points ultimately feed into the data center.

Figure 1

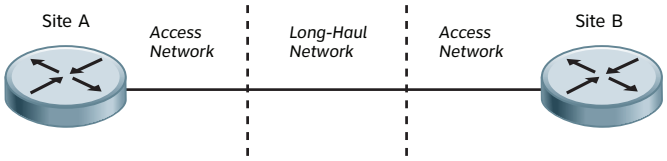

To interconnect 5 sties with private lines, 10 end-to-end physical circuits must be purchased.

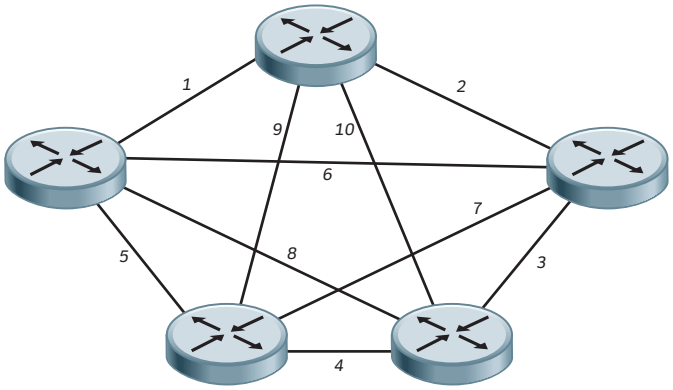

Figure 1 shows a Point-to-Point Private Line Network

Because of the characteristics described, Layer 1 private line services are often used in the following sample situations:

1) In networks with a relatively small number of nodes that all require direct connections to all other nodes.

2) For point-to-point high-speed connections dedicated to bandwidth-intensive operations, such as database synchronization or disaster recovery between two sites.

3) For high-speed, time-critical transactions that require the highest level of security, such as brokerage-to-trading floor communications.

4) To support broadcast television signals, which require high reliability and have nearly zero tolerance for rerouting. Here, circuits are often implemented in a redundant dual-circuit/failover design.

5) In organizations with a do-it-yourself philosophy that are not interested in statistically sharing bandwidth with other customers. For example, a worldwide manufacturing company might run private lines among its engineering sites to support high-bandwidth computer-aided design, or CAD, applications.

6) As an access service connecting sites to shared, higher-layer network services, which are described in the below subsections.

Layer 2

Layer 2 services are responsible for framing packets, providing an access mechanism to the physical circuit and error detection. This layer performs its portion of network duties using software-based protocols running on the WAN access devices – installed and configured either by the enterprise or the enterprise's router vendor. As another option, Layer 2 software can be installed and configured by the enterprise's carrier as part of a "managed" service that includes installing and managing CPE. A sampling of Layer 2 protocols includes frame relay, ATM, Ethernet and the Point-to-Point Protocol (PPP). In addition to running on the enterprise's WAN access equipment, Layer 2 protocols are also implemented in the carrier's WAN backbone infrastructure equipment.

When implemented as a WAN backbone service, some Layer 2 services add switching capabilities. Switching tends to appeal to companies that wish to take advantage of the economies of scale and the pervasive coverage of a carrier's nationwide or global backbone, which is shared by multiple customers. Here, an enterprise minimizes the number of access and backbone connections it purchases in the Layer 1 example. It can purchase a single physical circuit that connects a business location to an access point at the edge of the carrier backbone cloud, which contains a national or global mesh of high-speed circuits. From the carrier access point, the traffic from one site can be switched over to any other enterprise site (see Figure 2).

Figure 2

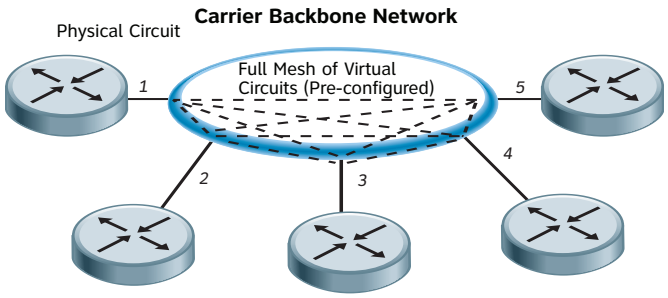

Five (5) physical circuits are purchased for full mesh connectivity instead of 10. Destination sites are predefined in customer router software.

Figure 2 illustrates a Layer 2 switched service

However, like private lines, virtual circuits also have to be pre-configured for every network connection, so configuration complexity remains about the same. The efficiencies are in the cost savings associated with purchasing fewer physical access and long-haul circuits.

Below is a discussion of some common Layer 2 services.

Frame Relay Access and ATM Backbone Services

These Layer 2 switched services rely on permanent virtual circuits (PVCs), configured between all participating end points through the carrier cloud, to segregate a given enterprise's traffic flows from those of other customers. In addition to reducing the number of physical access connections required at each corporate site, frame relay and ATM offer a "bursting" advantage over private line services, which can further reduce network costs. The way this works is that the pricing models for these services are such that the customer selects and pays by the month for a port speed and one or more PVC speeds per access connection. The PVC speed(s) might be significantly lower than the port speed (often T1, or 1.544 Mbps, for example) and less expensive than a full T1 access line. But most of the time, there is enough network capacity in the frame relay/ATM network backbone to allow traffic on those lower-speed PVCs to "burst" up to their port speed and take full advantage of T1 network speeds across the entire access line and backbone.

From a configuration standpoint, a PVC must be pre-defined and configured between each pair of sites that might wish to connect across the backbone. Layer 3 routing, which will be discussed in the "Layer 3" subsection, greatly simplifies the setup and configuration of highly meshed environments.

Ethernet Services

These Layer 2 services can function as switched services to mimic a LAN. Provisioned over fiber, they are highly scalable, in that Ethernet technology runs at speeds from 2 Mbps up to 10 Gbps and customers can often simply upgrade the interface on their premises-based Ethernet switches to "turn up the speed" of their bandwidth. Skilled resources for Ethernet local networks are widely available, so there is little or no training required for staff to apply those same skills to Ethernet in the MAN or WAN.

Note that fiber access is available only to about 13% of North American businesses. So Ethernet access service availability may not reach all the sites where a given organization wishes to deploy it.

Layer 2 MPLS Services

These services primarily exist to allow customers to gain the benefit of MPLS's meshed connectivity in the carrier backbone and retain non-IP interfaces at customer sites that don't yet require an immediate CPE upgrade. For example, Ethernet, frame relay and ATM can all be "tunneled" through MPLS, allowing the enterprise to have a single consistent backbone service (MPLS) with a mix of endpoints and access services at their various locations.

Layer 3

Layer 3 services embrace routing, which is a more dynamic form of traffic forwarding than switching. Routing relies on IP addresses and intelligent routing protocols to determine the "best" (fastest or least expensive) paths that traffic should take through the network at any given moment. For subscribers seeking more options from their networks than forwarding traffic from one point to another, Layer 3 services – because of the intelligence afforded them by information in the IP packet header – also offer routing-related advanced capabilities.

For example, Layer 3 services can offer inherent rerouting around failures, encryption, voice over IP (VoIP), high-availability/failover protocols and features, and bandwidth management for per-application and per-user control of each traffic flow.

From an enterprise configuration standpoint, adding sites to the network becomes easier than at Layers 1 and 2, because no circuits – neither physical nor virtual – must be predefined between any two end points. Rather, Layer 3 sites function independently of one another. A new site can simply be "plugged in" to the network service and that site will automatically be able to communicate to every other site.

Note that making the "best path" decision requires routers to auto-discover one another across the network, share route information and let one another know about path and router availability. As such, routing is inherently slower than Layer 2 switching. A combination of Layer 3 and Layer 2 that blends the strengths of each is available in MPLS VPN services. Because MPLS VPNs use IP route information and intelligence, they are considered to be Layer 3 services.

MPLS VPN Services

MPLS is generally deployed by organizations requiring a great deal of dynamic meshed connectivity among many distributed sites. Sometimes the need for meshed connectivity is driven by the deployment of real-time, multimedia applications that demonstrate much better performance without having to traverse a central hub site (making an extra "hop") to get from point A to point B. MPLS VPN services use a mix of IP routing and fast switching: IP routing at the edges of the backbone network determine where packet streams need to go and relieves network administrators from having to predefine site-to-site circuits (see Figure 3). However, across the backbone, fast-switching Layer 2 technology ensures very high forwarding speeds across most of the backbone equipment hops. In this way, MPLS VPNs set up virtual circuit-like partitions for inherent security and are sometimes considered a hybrid Layer 2/Layer 3 service.

MPLS VPNs also provide class-of-service (CoS) prioritization capabilities to ensure that real-time traffic performs as expected in the presence of data traffic. CoS becomes particularly important in Layer 3 IP and MPLS VPN networks, because these types of network services are used most often for supporting multiple types of traffic converged on a single network that reaches a fairly distributed set of sites with meshed connectivity. Each traffic type has different requirements and thresholds for maintaining performance. Because these different types of services are sharing a common backbone, it is important for the backbone to know how to prioritize and treat the packet types as they come across so that the services perform consistently well.

The way this happens is that a network administrator at the customer location configures its WAN edge router with industry-standard IP priority markings. These are known by a number of names, including Differentiated Services Code Point (DSCP), DiffServ and Type of Service (ToS). At the edge of the carrier's backbone network, the IP markings are communicated to the carrier's MPLS switch and stored in a label that is passed across the network to preserve and enforce the customer's priority markings end-to-end. If a Layer 2 access protocol, such as frame relay, ATM or Ethernet, is used, the Layer 2 protocol header is stripped off at the ingress of the MPLS switch, and the IP header, with its priority markings, is exposed to the WAN for enforcement.

Internet VPN Services

These services involve buying Internet access through a public Internet service and encrypting all data and authentication traffic with IP Security (IPSec) to partition a company's traffic from the rest of the Internet (see Figure 4). They are less expensive than MPLS services, but, because they traverse the public Internet, the same levels of performance are not guaranteed. Performance can vary depending upon whether the service purchased keeps traffic on a single carrier's network or whether it traverses the Internet segments managed by multiple carriers. The single-carrier service will be much more reliable and better-performing, simply because its operations fall within the purview of a single operator, who can manage it and troubleshoot it.

Figure 3

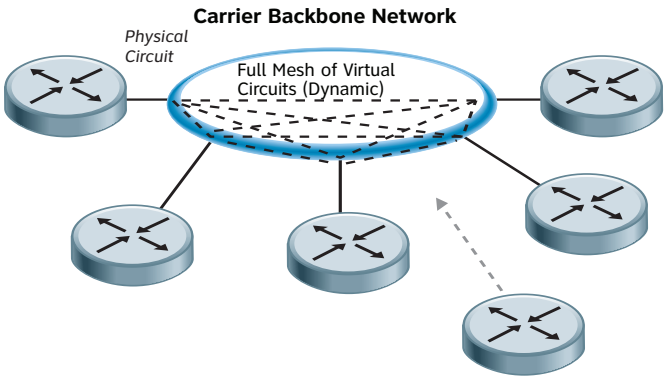

Five (5) physical circuits are purchased for full mesh connectivity instead of 10. No pre-configuration of site destinations is needed; new sites can "plug in" to the network and automatically become part of the mesh.

Figure 3 depicts a MPLS VPN Service

Figure 4

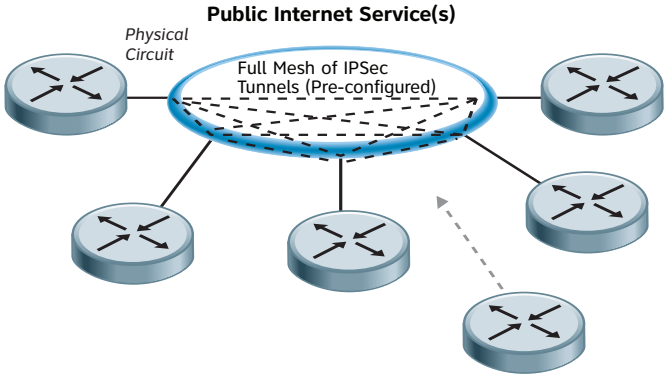

Five (5) physical circuits are purchased for full mesh connectivity instead of 10. IPSec encrypted tunnels must be predefined among sites for privacy. Performance is generally better across a single provider's network than multiple ISP networks.

Figure 4 illustrates an IPSec VPN Service

	Service Type(s)	Typical Applications	Attributes	
			Benefits	Possible Limitations
Layer 1	Private Lines (physical point-to-point circuits)*	Database synchronization	Reliable; few parts to "break"	Distance-sensitive pricing can can add up in large distributed networks
		Site-to-site disaster recovery	Available at very high speeds	All physical connections must be predefined; configuration-intensive in large networks
		Transactional B2B networking	Organization uses all capacity; no statistical capacity sharing with other entities	
		Distributed engineering network with bandwidth-intensive applications		
		Access connection to shared backbone service		
Layer 2	Frame Relay/ATM	Connecting multiple distributed sites to consolidated data center resources	Shared network cost efficiencies: reduction in private lines required	All virtual connections must be predefined; configuration-intensive in large networks
		Branch to regional site connectivity	Ubiquitous network footprint	
		Some inter-site mesh connectivity		
	Ethernet	Create metro or wide-area "LAN" with internal VLAN user group configurations intact	Fast, easy bandwidth provisioning and speed changes	Low availability (relatively small fiber footprint)
			Available at very high speeds	Generally not recommended for more than 50 meshed sites or 200 hub-and-spoke sites due to interior gateway protocol limits
	L2 MPLS	Retain existing WAN interfaces at some sites while adding new MPLS services	Avoids forklift CPE upgrade at all sites	Minimum scalability
Layer 3	MPLS VPN	Fast failover among multiple data centers	"Plug-in" connections simplify large-network configuration	Customer shares IP route information with carrier
		Full-mesh connectivity among large number of highly distributed sites	Enforces priorities of multimedia traffic end-to-end	
			Intelligent value-added IP services available	
	IPSec VPN	Use of the public Internet as a corporate network	Relatively inexpensive	Less predictable performance
		Teleworker and extranet access		Camouflaging IP information can suppress value-added capabilities
				Configuration-intensive in large networks
				Customer shares IP route information with carrier
	SSL VPN	Mobile worker access to corporate Web-based applications	Web browser is only client needed	
		Dynamic access, including disaster recovery	Controls access at application layer so is sometimes seen as more secure than IPSec from an access standpoint	Dynamic access can encourage use of unmanaged devices, which can increase risk of network infections

*Enabling service delivery technologies: Sonet/SDH, DWDM, T-Carrier

Source: AT&T

Figure 5 matches networking requirements to WAN Services.

Secure Sockets Layer (SSL) VPN Services

SSL is an encryption protocol designed for a wide range of web applications. SSL secures two applications communicating with each other, and addresses the requirements of mobile users by controlling user access at the application level. No special client software is required by the user; browser-based devices provide access to network resources. Users traverse the public Internet through tunnels encrypted at the application layer. Access is granted or denied at a SSL VPN termination device hosted and managed by the network operator (see Figure 5). Although it is not typically used for site-to-site communications, SSL is being mentioned in this paper as an option for remote access or traveling users.

Summary

There are compelling reasons to use Layer 1, 2 and 3 services and any given enterprise may use one or a mix of service types. Generally, private line networks, which operate at Layer 1, are used for point-to-point application traffic where the user organization (or part of an organization) is not interested in statistically sharing bandwidth with other customers. This is often the case for very high-speed applications such as machine-to-machine communications. Private line services are often considered because of their reliability, performance and security characteristics, as well.

Layer 2 services allow enterprises that are fairly distributed to reduce WAN access circuit costs by sharing the carrier's backbone service with other customers. Enterprise can also get extra bandwidth across the backbone when enough capacity is available in the backbone network to allow them to burst up to their full port speeds end-to-end.

Layer 3 services are oriented toward highly distributed organizations and are often selected by those running voice, data and video on a single network infrastructure. Layer 3 intelligent routing allows any site to find any other site simply by plugging into the backbone network, vastly simplifying configuration. Layer 3 intelligence also allows for value-added services such as encryption and other security features, as well as automatic failover between devices for redundancy and traffic management.

There is no one-size service that fits all organizations or even necessarily a single organization. As they plan their WAN strategies going forward, enterprises will do well to keep in mind the descriptions and sample applications discussed in this paper and to consider their own profiles carefully. Each service type has its place, and a given organization will build its optimum WAN strategy based on its application mix, traffic patterns, connection speed requirements and inherent "build or buy" networking philosophy.

For more information contact an AT&T Representative or visit www.att.com/business.

04/08/08 AB-1216-01

at&t
Your world. Delivered.