Services and features not available in all areas. Equipment installed by Cox,
and the installation of networks may vary. LAN wiring and electronics are not included.
Other restrictions apply.

www.cox.com

DMEWP1109

Carrier Ethernet: Transforming Business Telecommunications

Introduction

There is no question that telecommunications network services for the modern business customer have significantly and rapidly evolved over the past decade. With business customers demanding greater bandwidth, more flexible connections and more economical options, legacy Private Line, Frame Relay and ATM services are yielding to an array of advanced Ethernet and IP-centric networking services. These new services greatly increase a business's network capacity, deliver flexible bandwidth options, support a wider range of business applications, improve productivity and provide significant cost savings.

But while this diversity of options is a benefit to business customers, it also presents some challenges: how to choose a service or combination of services that best satisfies a particular customer's wide-area networking (WAN) needs. Business customers have a vast array of options including traditional Private Line, Layer 2 Virtual Private Networks (VPNs) such as Frame Relay and Ethernet, and Layer 3 VPNs commonly referred to by names such as IP VPNs and IP MPLS.

In assessing these WAN alternatives, businesses need to account for several factors including geographic reach, service availability, traffic patterns and types, existing and future applications, bandwidth needs, and the desire to maintain or relinquish management control of the WAN. Often, no one technology is best suited to meet all of an organization's WAN service reach needs, particularly for larger enterprises. Therefore, a hybrid solution utilizing more than one service technology is often necessary.

WAN Alternatives

Layer 1: Private Line Services

To start with, there are tried-and-true Private Line services, which include Digital Signal (DS) and Optical Carrier (OC) circuits which are based on Time Division Multiplexing (TDM) signaling. Private Line services offer dedicated connections at preset data rates, the most common speed (1.5Mbps) referred to as T-1. As such, they are used to set up point-to-point connections for voice, data or video applications.

Although they are well-understood and dependable services, Private Line connections have their drawbacks. These include inflexible bandwidth options, limited scalability, higher cost, and the fact that the point-to-point connections require bandwidth be dedicated to a particular customer whether or not it is being used.

The bandwidth limits lie in the fact that TDM-based circuits are based on telephone-centric technology and are only available in fixed bandwidth increments. For example, a DS-1/T-1 circuit offers 1.544Mbps symmetric – and no more. Technology allows bonding up to eight T-1 circuits for a total of about 12Mbps. After that, customers are required to make a sizeable stair-step upgrade to more costly DS-3/T-3 lines – or fractional DS-3 options – with a maximum 45Mbps throughput. These incremental bandwidth upgrades typically require hardware changes to customer equipment and on-site visits by technicians to reconfigure the circuits and the equipment, often at a cost of several hundred dollars to the customer.

Then there is the cost of the circuits themselves. Private Line circuit pricing is often based on distance, so the greater the distance between the two endpoints, the higher the price. The cost is also generally higher than newer packet-based services given the bandwidth is always dedicated to a particular customer. Moreover, when deploying Private Line circuits to facilitate hub-and-spoke or meshed network topologies, the number of circuits required and costs increase substantially.

These limitations gave rise to Layer 2 VPNs such as Frame Relay and ATM services, both of which bring switching to the party and therefore offer greater bandwidth flexibility and topology options.

Layer 2: Frame Relay and ATM Services

Frame Relay is a Layer 2 technology, meaning it operates within the data link communications layer of the Open Systems Interconnection (OSI) stack, in contrast to the physical connection of Private Line circuits. Introduced in the early 1990s, Frame Relay was also one of the initial packet-based WAN technologies, in that it parses customer traffic into variable-length frames which are then transmitted across a service provider's shared core network.

To keep each customer's traffic separate, Frame Relay introduces the concept of a virtual circuit. Essentially, each frame of data is tagged with an identifier corresponding to the virtual circuit for a specific customer. This allows multiple virtual circuits from many different customers to traverse the same path in the core of the service provider's network. The most common type of virtual circuit in Frame Relay is called a permanent virtual circuit (PVC) where the virtual connection paths between locations are pre-established and more permanent in contrast to the dynamic nature of switched virtual circuits or packet-based routing.

The virtual circuit concept helps to lower network costs in two ways:

- It allows for shared physical circuits in the core of the service provider's network while keeping each customer's traffic distinct and protected.
- It allows one physical circuit at a customer location to contain multiple virtual circuits terminating at multiple remote customer locations.

Frame Relay has been tremendously successful because it is substantially more cost-effective than Private Line services for connecting networks with several sites. However, it also has a few key limitations.

First, Frame Relay generally only supports bandwidth options up to 45Mbps, and it relies on TDM-based circuits with the same rigid bandwidth constraints for access to a provider's Frame Relay network cloud. Additionally, given the variable-length nature of the frames, Frame Relay has primarily been used for data networking and is not ideal for carrying more delay-sensitive customer traffic such as voice or video. Finally, it differs from the standard Ethernet technology deployed in customer LANs, thereby requiring more costly customer interface equipment and unique customer expertise to manage the WAN.

Another Layer 2 technology, Asynchronous Transfer Mode (ATM), attempted to address the speed and other limitations of Frame Relay. ATM increased the supported bandwidth to 622Mbps and used small, fixed-length cells rather than variable-length packets. While not always optimal for bursty data traffic, fixed-length cells allow ATM to support data, voice and video across a single converged network. Like Frame Relay, ATM also uses the concept of virtual circuits.

Although ATM solves some issues, it has its own limitations. The small, 53-byte fixed-length cells generate considerable overhead which can significantly reduce the throughput of a customer's actual network traffic, particularly for bursty data traffic.

Also, because multiple standards have been developed over the years to address the use of ATM in different endpoint devices, it is complex to provision and maintain. As with Frame Relay, ATM also requires more costly customer premise equipment and unique expertise to maintain the WAN.

For these and other reasons, ATM did not widely proliferate in the marketplace and has predominantly been used for backbone transport connectivity for large enterprises and service providers.

Layer 2: Carrier Ethernet Services

Frame Relay and ATM services are waning as next-generation Ethernet and IP-VPN services have entered the market as very worthy replacements. An outgrowth of the familiar local-area network (LAN) technology, Ethernet has evolved into a wide-area network (WAN) technology able to link endpoints across town or across the world at data speeds anywhere from 1Mbps up to 10Gbps. Thanks to advances in technology and the standards work of the Metro Ethernet Forum (MEF) industry consortium of service providers and equipment vendors, Ethernet has evolved well beyond its best-efforts roots to offer true carrier-grade capabilities. The term Carrier Ethernet has emerged to highlight the inclusion of these carrier-grade capabilities including reliability, scalability, and quality of service.

The fact Ethernet is already used in the customer's LAN greatly simplifies the LAN/WAN interconnection between the customer and service provider. At each customer endpoint location, the service provider supplies a standard Ethernet interface delivered over fiber, coaxial cable or copper phone lines (the latter options restricted somewhat by the physical medium's own bandwidth limitations) directly into the customer's existing office router or Ethernet switch. In most cases, the customer does not need to buy added equipment, thereby lowering the initial connection costs.

Like its Layer 2 peers, Carrier Ethernet services are based on virtual circuits rather than physical circuits that are switched across the service provider's network. Dubbed the Ethernet Virtual Connection (EVC), it defines the endpoints to be connected to deliver a particular service such as voice or data. As a result, the service provider creates a closed VPN, allowing the customer's data to flow securely between all of the customer's defined endpoints.

For every EVC, the service provider also provisions and assigns user network interfaces (UNIs), which define the individual customer ports or endpoints assigned for each customer location. The service provider uses the UNIs in conjunction with the EVCs to set the bandwidth profiles for that customer, including the overall speed and performance characteristics of individual service flows.

The service provider can work with the customer to design highly flexible networks using three basic standard topologies as defined by the MEF:

- **Point-to-Point: Ethernet Private Line (E-Line EPL),** which creates a point-to-point connection similar to a traditional Private Line service.

- **Hub-and-Spoke: Ethernet Virtual Private Line (E-Line EVPL),** which creates connections in a hub-and-spoke arrangement similar to typical Frame Relay network deployments. This service establishes point-to-point EVCs between two locations, but enables these individual EVCs to be service multiplexed onto a single UNI port, such as at a hub site location.

- **Multipoint-to-Multipoint: Ethernet LAN (E-LAN),** which creates multipoint-to-multipoint connections. This service provides connectivity to and from any of the UNI endpoints within an EVC associated with a particular E-LAN configuration. Therefore, it enables the same Ethernet topology and traffic flow that are used in customers' Ethernet LANs to be supported in their MAN/WANs.

ETHERNET LINE (E-LINE) SERVICE

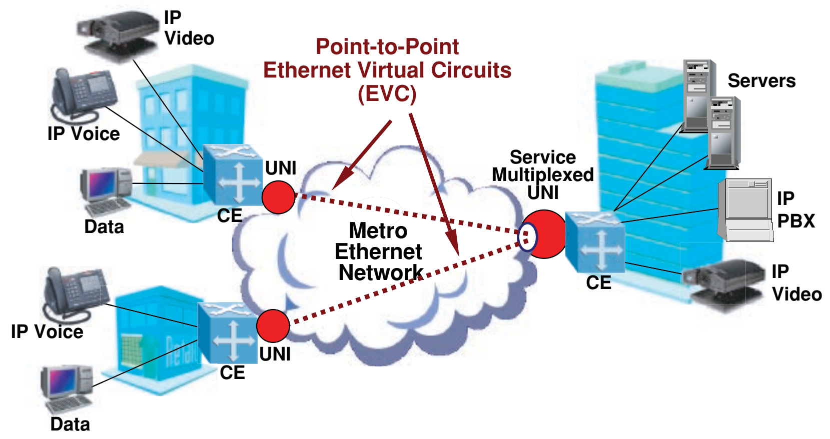

ETHERNET LAN (E-LAN) SERVICE

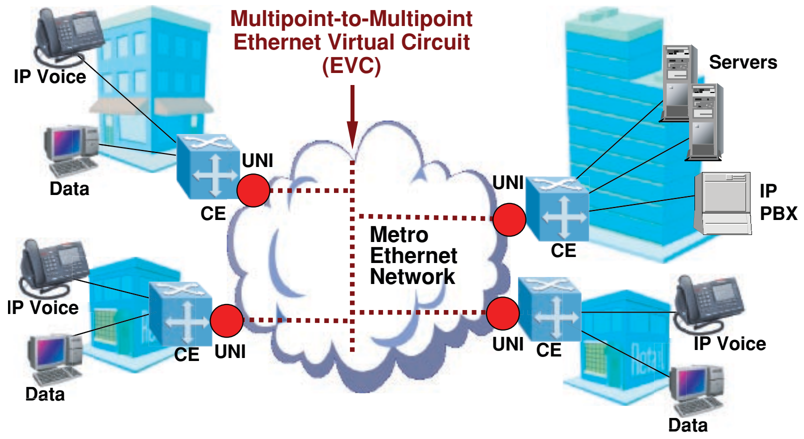

Beyond supporting multiple topologies, there are key additional benefits associated with the use of Layer 2 Ethernet VPN services as outlined below.

Highly scalable and flexible bandwidth: Ethernet offers substantially greater bandwidth scalability and flexibility, with the ability to quickly adjust throughputs. The service speeds can range from as little as 1Mbps to 10Gbps, allowing the service provider far more flexibility in assigning speeds to meet an individual customer's needs. Additionally, the bandwidth can be readily changed as per the customer's requirements, typically remotely and without the need for any new customer premise equipment.

Ethernet simplification and economics: The use of Ethernet in the WAN greatly simplifies the LAN/WAN boundary, as modern organizations have Ethernet LANs and maintain Ethernet expertise within their IT staff. This allows customers to reduce the costs associated with training IT staff to service other WAN technologies and allows them to work more effectively with providers in provisioning and troubleshooting the services. Moreover, the overall simplification and wide-scale proliferation of Ethernet technologies results in significantly lower costs to providers and customers for Ethernet circuits, equipment, interfaces, and ongoing maintenance.

Customer-maintained routing control and protocol transparency: Layer 2 WAN services such as Ethernet provide a complete separation of the customer's Layer 3 routing and addressing from the service provider's underlying WAN transport. Thus, customers are able to retain network and routing control, which is often preferred by customers who have the in-house expertise and do not want to give up control of this critical function. Ethernet and other Layer 2 services are also network protocol agnostic with the ability to easily support any protocol including IP, SNA, IPX and others. This is attractive to customers with legacy applications that have not been converted to run over IP.

In addition to the improvements in reliability and security, Ethernet has evolved beyond the confines of fiber-optic connections, particularly in the past few years. Ethernet services are increasingly being delivered over additional mediums such as hybrid fiber-coax cable networks, copper facilities and most recently point-to-point microwave wireless links. This is greatly extending the availability of Ethernet services beyond larger commercial buildings to satellite or branch locations and other mid-sized or smaller business customers.

Service providers are also extending Carrier Ethernet's reach beyond their own networks by establishing Network-to-Network Interfaces (NNIs) with other service providers. Individual providers are already forging these NNIs amongst themselves, and the MEF is pursuing NNI standard specifications to foster wider-scale interconnection agreements between providers.

On the horizon are more improvements and standards to add embedded mechanisms for service monitoring – an attribute business customers may demand to ensure that the services they are receiving are meeting minimum levels of performance and uptime.

MPLS-based Ethernet VPNs

As IP technologies and services continue to proliferate, service providers are transitioning to Multi-Protocol Label Switching (MPLS)-based technology in their core backbone and metro networks. The goal is to leverage a common, robust and efficient IP infrastructure to provide a host of voice, Internet, video and transport services such as Ethernet using MPLS as the common underlying technology.

MPLS is a packet switching technology that combines the benefits of both connectionless and connection-oriented network technologies. This allows it to efficiently and cost-effectively support very high-speed data traffic while also providing stringent reliability and quality of service mechanisms necessary to support more performance-sensitive traffic. It also adds sub-50ms protection mechanisms similar to traditional SONET-based network architectures.

It is important to note that MPLS is used by service providers for two distinct purposes. As indicated above, it is used as an IP/MPLS service and transport infrastructure technology which provides robust traffic engineering capabilities to carry a variety of business or consumer services over a converged network platform. Additionally, this same MPLS infrastructure is also commonly used as a means to deliver Layer 2 Ethernet or Layer 3 IP-VPN WAN services. Much of this is invisible to customers, and they generally do not participate nor do they necessarily need to be aware of how the service is implemented by the service provider.

As it happens, there is a relatively new technology augmentation to MPLS that enables it to be used to implement Layer 2 Carrier Ethernet services for multipoint-to-multipoint (E-LAN) topologies: Virtual Private LAN Service (VPLS). In a nutshell, VPLS is used to create Carrier Ethernet VPNs provisioned over a provider's MPLS network that allows multiple customer locations within a common VPLS instance to interconnect as if they were on the same LAN, regardless of their physical location.

Here's how it works. VPLS-capable routers placed at the edge of the provider's network link the local access connections to customers. These routers are capable of learning and managing the Ethernet-addressable information transmitted via individual customer VPLS networks. These VPLS-capable Provider Edge (PE) routers are then connected in a full mesh of MPLS label switched paths. Several individual customer VPLS instances or services can then share these label switched paths, with the system keeping each customer's traffic separate based on the MPLS labels attached to their data packets. When a VPLS customer's data hits the PE router, the router attaches labels identifying, among other things, the VPN transport and service labels for proper delivery of the customer's traffic. It then sends the packet down the label switched path. At the other end, the egress PE router strips off the MPLS label and sends them on to the customer's destination endpoint as Ethernet frames.

VPLS-enabled Carrier Ethernet Services provide the same standard Ethernet service and interface to the customer with complete separation of the customer's Layer 3 (IP) routing network from the service provider's network. Thus, customers maintain the same control over networking and routing as is the case for traditional Carrier Ethernet Services. New locations can be easily added with no change to the network or routing configurations for the existing endpoints.

Layer 3: MPLS-based IP VPN Services

In many cases, Layer 2 Carrier Ethernet services can accommodate large corporate networks with diverse locations. But for companies requiring large nationwide or global networks mixing a variety of connection types, there is yet another alternative in Layer 3 MPLS IP VPNs. MPLS-based IP VPNs, which came to market before MPLS/VPLS-based Layer 2 Ethernet services, have been deployed worldwide by large national and global service providers and achieved significant market share among large enterprises.

The fundamental difference with Layer 2 Ethernet at the customer premise is the way in which it processes traffic. In contrast to Layer 2 services, MPLS-based IP VPNs are established via Layer 3 router technology, controlling data flow via routing rather than switching. Data packets entering the network are given labels by the provider edge routers identifying the service and transport labels for proper delivery of customer packets. Along the way, core routers examine and swap these labels, routing the traffic to more efficiently deliver it across the larger backbone. At the destination, the edge router strips off the label and delivers it to the recipient.

Layer 3 MPLS IP VPN's primary advantage is its ability to connect numerous far-flung locations because of its extensive reach across major network backbones and because it is largely connection agnostic. As a Layer 3 IP service, MPLS-based IP VPNS can typically support a wide variety of technologies for the local access connections, be it Private Line, Frame Relay, Ethernet, DSL, or a cable modem. For large corporations that must connect geographically dispersed remote offices using a mish-mash of connections, this is attractive. However, it is important to note, the characteristics associated with any particular access technology still apply, such as the rigid bandwidth limitations associated with Private Line or Frame Relay-based access.

The technology can also extend between service provider networks if the providers involved forge NNIs for interconnectivity. The IP MPLS Forum, an industry consortium of service providers and equipment makers, is working on an NNI inter-carrier interconnection specification that seeks to bring more uniformity to this process.

There are some considerable drawbacks to Layer 3 MPLS IP VPN services as well, primarily greater costs, the inability to support non-IP protocols natively, and loss of network routing control by the customer. Because of the heavy lifting at the router level, much greater coordination of IP routing with the service provider is required. While customers can relieve their IT staff of the weighty task of managing the VPN, they also have to give up some level of routing control to the service provider – which often raises reliability or security concerns.

Meanwhile, the managed routing functionality that MPLS-based IP VPNs inherently provide comes at a price, with premium rates that generally exceed Layer 2 Carrier Ethernet options. This price premium is sometimes substantial, particularly at higher bandwidths or where a provider must backhaul access circuits considerable distances to connect into their IP VPN network service nodes, which are often only located in the major backbone hubs of national providers.

Internet-based VPNs

There is yet another set of alternatives a business customer can consider – Internet-based VPNs, including IPSec and SSL VPN services.

For these services, rather than relying on a service provider to set up a private network for data traffic, customers can tunnel a VPN across the open, lawless public Internet. Able to harness cable modem, DSL or other Internet access connections, these products encrypt the data using either standard IP Security (IP Sec) or Secure Socket Layer (SSL) protocols, thereby making them unreadable to any but the customer's approved access points. Client software downloaded to customer PCs and other devices acts as the "key" to unlock the encrypted data.

For many businesses, such secure Internet-based VPNs fall short because they still rely on the shared bandwidth of the public Internet. The VPN traffic is therefore prone to slowdowns and outages that affect the larger Internet cloud. Many business customers therefore will use Internet VPNs to connect employees who are traveling outside the office or working from home, but will opt for more reliable private network services such as Carrier Ethernet for interoffice connections.

Conclusion

Based on this analysis of WAN services, it is clear that the business owner has a lot to consider when choosing among the various WAN alternatives. The devil is in the details, requiring the customer to match their particular corporate networking needs to the available options.

Still, there are some general observations that can be made, starting with the fact that legacy Private Line services can still sometimes be a sound choice for simple point-to point connections. However, even point-to-point circuits are increasingly being delivered with Ethernet handoffs and Private Lines are generally no longer adequate to connect multiple office locations due to their high costs, limited bandwidth scalability, and their inefficiency in supporting converged networks carrying a mix of data, voice and video traffic. They also are not well positioned to support the growing array of IP applications and managed services many businesses seek.

At the other end of the spectrum, Layer 3 MPLS-based IP VPNs offer advantages to large businesses trying to connect numerous far-flung locations that require the use of multiple local access technologies, and for businesses willing to outsource greater management and maintenance of their WAN. But these services are also generally more costly, often substantially so for larger circuits. Additionally, they introduce far more complexity and require that the business customer give up essential routing control to the service provider. For these reasons, Layer 3 MPLS-based IP VPNs are often not well suited for a wide range of metro or regionally oriented businesses, as well as for other businesses where larger circuits are needed between sites located within the same metro or regional area.

For such needs, Layer 2 Carrier Ethernet services strike a strong balance between reliable, secure connectivity and flexibility in network options, IP applications support and scalable bandwidth levels – all at lower costs. As a technology augmentation to MPLS-based networks, VPLS offers added scalability and flexibility to Carrier Ethernet services requiring meshed connectivity between many locations.

In the end, the customers want a reliable, scalable and reasonably priced multisite network service that is available for their locations and is flexible enough to adapt to their changing service needs. So the reality is that the underlying technology used to deliver the service is less important to many customers, and no one technology is the perfect fit for all business customers.

With a greater field of options – thanks to the more recent addition of next-generation Carrier Ethernet services and underlying VPLS technology – customers stand a better chance of finding the private wide-area network service or mix of services that best fits their needs for today and the future.

Comparison of WAN Alternatives

	Private Lines	Frame Relay	Carrier Ethernet	MPLS-based IP VPNs	Internet VPNs
Customer Handoff	TDM (Layer 1)	Frame Relay (Layer 2)	Ethernet (Layer 2)	IP (Layer 3)	IP (Layer 3)
Key Benefits	Tried-and-true reliability and extensive availability over fiber, copper and wireless facilities Bandwidth is dedicated to a particular customer Available at very high speeds	Tried-and-true reliability and extensive coverage More cost-effective than private lines for multisite networks	Ethernet simplicity eases implementation and troubleshooting Very scalable, granular bandwidth options that can be easily changed Lower cost of bandwidth and CPE equipment Customer maintains complete network routing and IP addressing control Transparently supports all network and application protocols	Strong national/global coverage due to large deployments and ability to leverage multiple local access technologies Provides a fairly turnkey network service for customers looking to outsource IP network control Additional IP-based network services often available from the provider as extensions to the service	Ubiquitous availability of the Internet Relatively inexpensive to implement
Possible Limitations	Higher costs, often distance-based pricing Inflexible, rigid bandwidth options Point-to-point nature inefficient for multisite networks	Primarily suited for data, lacks hard CoS/QoS to support convergence of voice and video applications Relies on TDM for access with rigid bandwidth options Being phased out by most providers	Primarily available over fiber to date, but increasingly offered over copper and Hybrid Fiber Coax (HFC) facilities	Customer must coordinate IP network and routing with the provider Only supports IP protocol natively Often more costly, particularly for larger circuits	Less predictable reliability and performance Complex to implement and maintain for larger networks
Typical Applications	High-bandwidth connections between major sites such as data centers Local access connections to shared WAN service Interoffice circuits where TDM handoffs required	Large hub-and-spoke data networks Branch office to regional or headquarters site connectivity	Metro and regional oriented wide-area networks High-bandwidth connectivity between large data centers or other enterprise locations Frame Relay and Private Line replacement/augmentation Site-to-site disaster recovery	Larger regional, national and global enterprise networks Full-mesh connectivity between large number of geographically distributed locations Customers seeking to outsource IP routing control and management for some or all of their locations	Extend intranet and extranet connectivity to remote access users and smaller branch locations Establish temporary or backup WAN connections