Ethernet-Based Public Communication Services: Challenge and Opportunity

Lothar Zier, Detecon International

Wolfgang Fischer and Frank Brockners, Cisco Systems

ABSTRACT

Enterprises and residential customers increasingly take advantage of more sophisticated applications and evolve their business models. At the same time service providers face declining revenues from existing connectivity services. SPs respond to these trends and start to offer highly customized, high-bandwidth network services that complement and interwork with their existing leased line, ATM, or frame relay offerings. Ethernet is quickly becoming the customer UNI of choice. Ethernet, being a packet-based technology, complements IP-based services such as content, voice, data, video, and additional value-added services over a high-speed access connection. Ethernet also provides a flexible high-speed connection to the SP network and eliminates access bottlenecks. Using Ethernet as the common access interface, providers can employ flexible transport services that enable additional value-added services at layer 3 and above. At the same time the rollout of Ethernet services also challenges the installed base, given that bandwidth per customer is usually much higher and additional attributes for tight SLAs apply. This calls for an updated approach to network deployment and rollout: networks need to be planned and built in a service-centric fashion. Starting with a brief look at the driving forces for Ethernet-based public communication services, this article identifies five generic services for Ethernet MANs/WANs and discusses related deployment issues of the different service options, such as the degree of customization, geographic reach and bandwidth profile of the service, as well as evolution aspects of the installed base.

INTRODUCTION: EVOLVING TO A SERVICE-CENTRIC WORLD

Until recently, companies had a limited choice of core service offerings such as leased lines from their service providers (SPs). Today, to meet the performance requirements of new Internet applications, enterprises are demanding customized services that will enable them to align the network closely with their business objectives.

To deliver the tailored services customers demand, SPs are undergoing fundamental changes in their operational and organizational models, and their assumptions about how to design and deploy networks. Historically SPs have focused on building network infrastructure to deliver a specific service only. This led to technology-specific organizational silos. Typically, there is little or no interaction between silos, and the underlying vertical business model often did not allow SPs to offer flexible technology-agnostic services.

THE ETHERNET ADVANTAGE

Ethernet has become the dominant interconnection technology in enterprise networks due to its simplicity, cost effectiveness, scalability, and familiarity to enterprise customers. These factors have led many enterprises to consider Ethernet as the user-network interface (UNI) to access MAN and WAN services. Consequently, many SPs are also turning to Ethernet as a transport technology for metropolitan and wide area networks (Fig. 1).

While the majority of Ethernet deployments use Category 5 copper cables for speeds from 10 Mb/s to 1 Gb/s, Ethernet is increasingly transported over other media, like single-mode and multimode fibers, wavelengths in dense wavelength-division multiplexing (DWDM) and CWDM optical transmission systems, synchronous digital hierarchy/synchronous optical network (SDH/SONET) transmission systems, wireless systems, and digital subscriber line (DSL). Additionally, Ethernet can also be transported using packet-oriented layer 2 transport mechanisms such as frame relay (FR), asynchronous transfer mode (ATM), multiprotocol label switching (MPLS), and Layer 2 Transport Protocol (L2TP).

Depending on the particular requirements for a specific service, most of the above technologies are employed within the network, from native Ethernet transport and switching using Ethernet

 0163-6804/04/$20.00 © 2004 IEEE

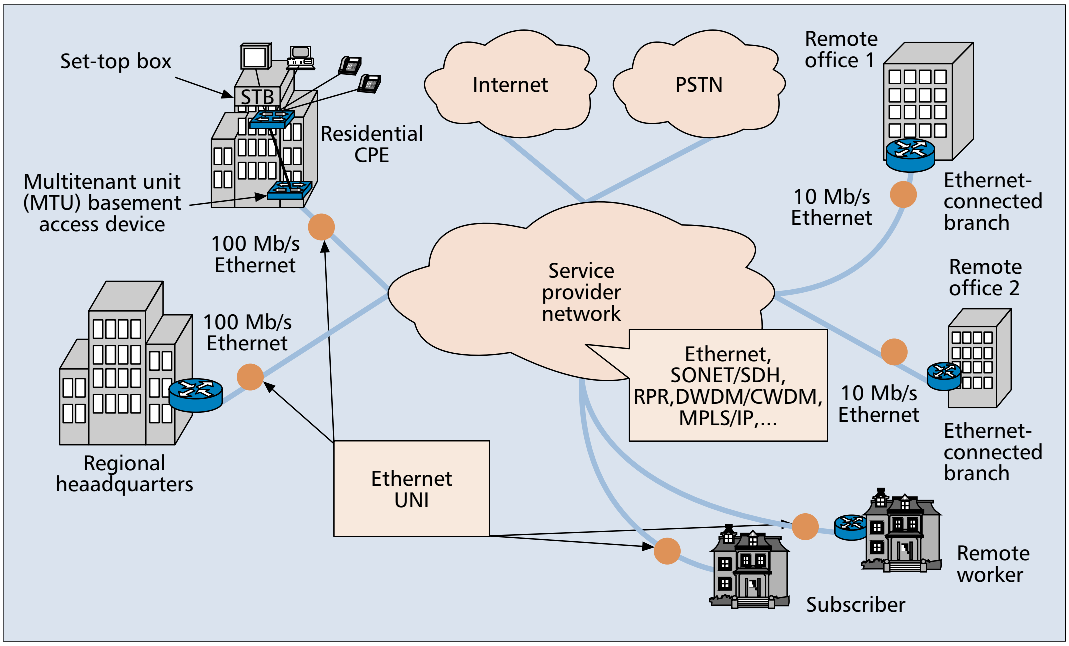

■ **Figure 1.** *Overview.*

While Ethernet offers significant advantages, the evolution of current service offerings toward customized Ethernet-based services also includes technological, structural, and deployment challenges.

over SDH/SONET to Ethernet over MPLS for wide area connectivity. In fact, most DSL deployments today simply transport Ethernet frames on top of ATM between customer premises and aggregation routers in the network. In many DSL networks, there is a clear trend for Ethernet transport to replace ATM for cost and manageability reasons.

Although initially true, Ethernet services are no longer limited to metropolitan areas, but are now offered nationally and internationally.

Depending on the target customers the SPs are addressing, two distinct service delivery paradigms have emerged.

Ethernet to the business (EttB) addresses the needs of enterprises for high-speed interconnection of multiple sites and Internet access. High-speed network access typically ranges from 10 Mb/s for remote sites to 1 Gb/s for corporate headquarters. There are stringent requirements in terms of availability, quality of service (QoS), and security. Mission-critical information, including real-time information like voice and video, is transported over these services integrating packet networks.

Ethernet to the subscriber (EttS) addresses the needs of home offices and residential subscribers. One of the dominating requirements for this kind of deployment is low cost and very high scalability. Compared to EttB there are typically two orders of magnitude more customers to serve in a given area, which has significant ramifications for the network architecture. The services that need to be supported by these networks are often summarized by the term *triple play*, referring to voice, video, and data delivery.

Increasingly, these service delivery models are blurring as many SPs start to address both market segments, facilitating applications like remote access to corporate networks.

While early metro Ethernet networks were exclusively based on native (Gigabit) Ethernet using an Ethernet switch infrastructure, more and more solutions include other transport mechanisms such as Ethernet over SDH/SONET, or DWDM as well as interworking with FR or ATM. "Ethernet-UNI" access to the service is the common denominator rather than the underlying architecture or technology.

THE ETHERNET CHALLENGE

While Ethernet offers significant advantages, the evolution of current service offerings toward customized Ethernet-based services also includes technological, structural, and deployment challenges. Although many organizations such as the Metro Ethernet Forum (MEF), IEEE, Internet Engineering Task Force (IETF), and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) are addressing the technological and architectural challenges, the deployment of Ethernet-based services often has additional implications.

Even though Ethernet services are becoming widely available, these services are still in their infancy. On a global scale, there is little experience in which service offerings are the most profitable and which would gain the most acceptance with customers. Deploying Ethernet-based services requires that security, QoS, availability, and scalability be considered. When UNI scalability is considered, the migration of customers to Ethernet usually involves a significant increase of the bit-rate supported at the UNI, i.e. customers migrating from low-speed FR to 10 Mb/s, or more, Ethernet-based services. While most of today's SONET/SDH transport networks were designed with T1/E1 scaling increments, this is no longer true in the case of Ethernet services being offered across these networks. For example, the equivalent of 168 STM-1 connections can carry either more than 10,000 E1 connections (for potentially 10,000 customers) or 24 Gigabit Ethernet connections. Note that bit rate is only one attribute of the UNI, and a customized service level agreement (SLA) is likely to include additional parameters.

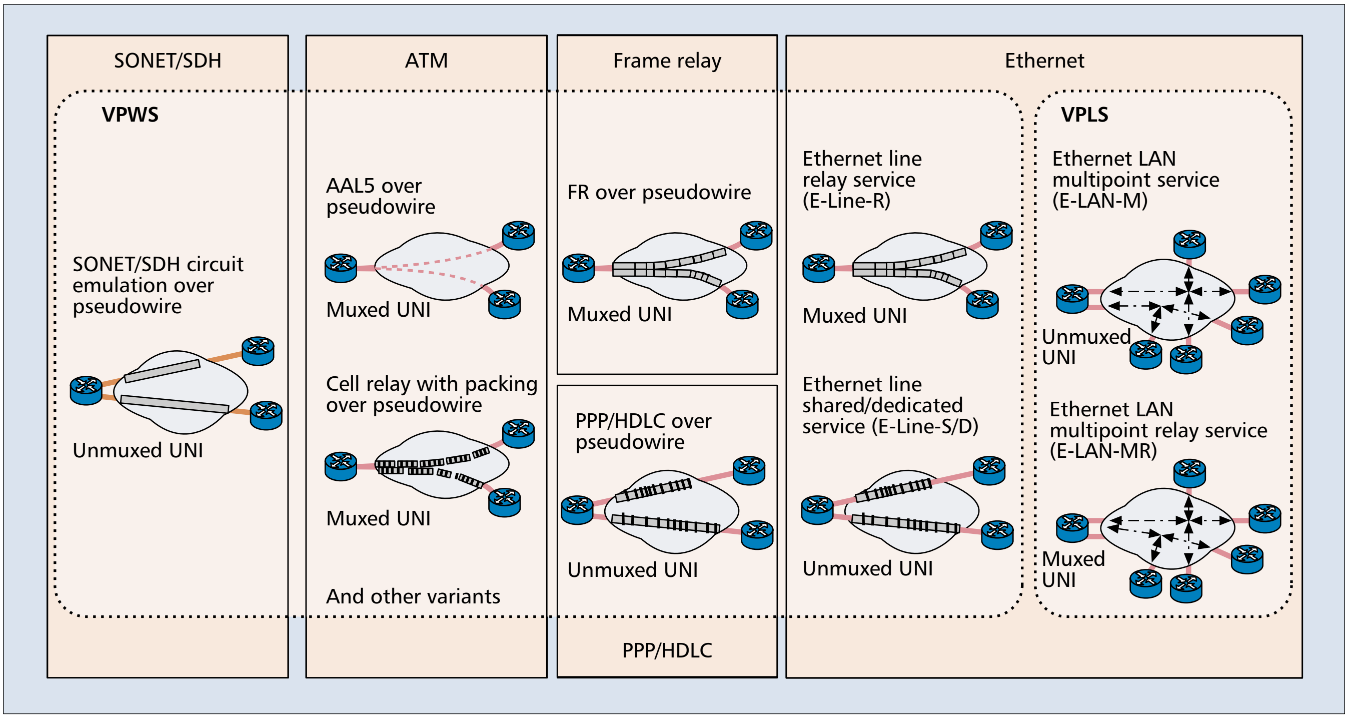

■ **Figure 2.** *Layer 2 virtual private network services: IETF terms in relation to generic Ethernet services.*

APPROACHING ETHERNET-BASED WAN SERVICES

Ethernet is becoming the natural evolution of UNIs for both layer 2 and layer 3 services. Ethernet-based WAN services are expected to complement existing services. Churn from legacy services such as leased lines, FR, and ATM is expected, but this will only happen over time. Following the service-driven approach, these legacy services should always be taken into account when discussing implementation and deployment of new Ethernet-based WAN services.

Although an Ethernet UNI is very often associated with layer 2 services, an Ethernet UNI may as well provide access to both layer 2 and layer 3 services, and allows the flexibility for a layer 2 service to be used to access a layer 3 network. One should note that layer 3 virtual private network (VPN) services, especially IP-VPN using Border Gateway Protocol (BGP)/MPLS (RFC 2547) or IPSec, are complementary to the layer 2 Ethernet services described below. The ability of SPs to offer a wide service variety that incorporates new and legacy services provides competitive advantages over pure legacy or pure Ethernet competitors.

ETHERNET WAN SERVICE DEFINITIONS

Different organizations have defined different terms for Ethernet-based WAN services. The MEF focuses on the capabilities and characteristics of the UNI [1] and has defined two service groups: a point-to-point service, or E-Line, and a multipoint-to-multipoint service, or E-LAN. The IEEE 802.1ad project [2] specifies a provider

Characteristic Term	Description
Ethernet virtual connection (EVC)	An EVC connects two or more UNIs, enabling transfer of Ethernet frames between them, at the same time preventing data transfer between sites that are not part of the EVC.
Service multiplexing	Service multiplexing describes the capability to allow different services on a single physical Ethernet interface. This is achieved by using the VLAN ID as a connection identifier. IEEE is in the process of defining a multiple services UNI that allows untagged, tagged, and double-tagged frames on the same physical interface.
Bundling	Bundling allows multiple VLANs to use the same EVC.
All-to-one bundling	All-to-one bundling provides transport of the entire content of a physical interface, irrespective of the presence of VLAN IDs.
VLAN transparency	VLAN transparency means that the VLAN ID does not determine an EVC, and VLAN IDs are not altered when packets are traversing the network.
PDU transparency	Layer 2 protocol data units (PDUs) are transported transparently across an EVC. Transparent transport of bridge PDU (BPDUs) that are used as part of the Spanning Tree Protocol (802.1D) is an example.

■ **Table 1.** *Characteristic terms for Ethernet services.*

UNI characteristic	Ethernet virtual connection (EVC)* type	
	Point-to-point	Multipoint-to-multipoint
Service multiplexing*	Ethernet line relay service (E-Line-R) frame relay replacement with Customer equipment (CE) likely to be a router	Ethernet LAN multipoint relay service (E-LAN-MR) CE likely to be a router
All-to-one bundling*	Ethernet line shared service (E-Line-S) Ethernet line dedicated service (E-Line-D) Private line replacement with CE router or bridge	Ethernet LAN multipoint service (E-LAN-M) LAN extension with CE router or bridge
	Ethernet line shared service (E-Line)*	Ethernet LAN service (E-LAN)*

* Metro Ethernet Forum (MEF) terminology

■ **Figure 3.** *Ethernet services: MEF definition and service characteristics in relation to generic Ethernet services.*

bridge that includes a UNI supporting multiple different service types at a time. The IETF takes an architectural approach to services that also describe two service groups: virtual private wire services (VPWS) and virtual private LAN services (VPLS). The IETF definitions are directly related to an underlying MPLS or IP packet-switched network and the use of pseudo-wires (PWs) to transport frames end to end ([3, 4]). It should be noted that VPWS is not limited to Ethernet, but may also support FR, ATM, or time-division multiplexing (TDM) emulation services. Figure 2 shows a schematic representation of IETF and MEF service definitions. Figure 3 places MEF and IETF services in the context of the generic Ethernet WAN services used in this article. SG15/Q12 of the ITU-T takes a network-centric approach to Ethernet services. Service definitions for Ethernet services are still in discussion but will be based on the "Architecture of Ethernet Layer Networks" [5] currently in the approval process.

Following the definitions in [6], this section takes an SP-centric approach and further refines the two basic service definitions of the MEF. As the MEF definitions represent only a very coarse classification of Ethernet services, there is often a need to enhance these classifications to represent the services being deployed and utilized widely today. As standardization progresses these terms will certainly change. These definitions may be thought of as generic transport services: *Ethernet line shared* (E-Line-S); *Ethernet line dedicated* (E-Line-D); *Ethernet line relay* (E-Line-R); *Ethernet LAN multipoint* (E-LAN-M); and *Ethernet LAN multipoint and relay* (E-LAN-MR). They may be bundled with higher-layer services (hosted telephony, security, etc.) and corresponding SLAs to form the actual service offering of the SP. It is important to note that SPs will define their own unique service definitions, but the underlying service will conform to the attributes detailed below. The characteristic attributes used in the service definitions are shown in Table 1. Table 2 provides a detailed definition of the five generic services.

Ethernet Service Delivery: Deployment Considerations for Point-to-Point and Multipoint Service

Point-to-point (P2P) services follow a model similar to FR or ATM whereby each peer requires a P2P connection. The complexity of configuring the network service grows as the number of peers increases. However, the relative simplicity of P2P networks allows for comparatively simple troubleshooting and policy definition. Point-to-point connections also engender predictable traffic patterns that ease network traffic engineering as well as QoS and security policy definition.

Multipoint-to-multipoint (MP2MP) services introduce a new WAN broadcast model. The multipoint nature engenders a somewhat simpler plug-and-play configuration; as the MP2MP service emulates an Ethernet switch, each device will discover other peers using broadcast/multicast protocols. The relative simplicity and inherent peer-to-peer connectivity makes an E-LAN-M-based layer 2 VPN service attractive to some enterprises as an alternative to a layer 3 VPN service. Although a layer 2 VPN service allows the enterprise to retain control over the routing topology of the network, the E-LAN-M, being multipoint in nature, requires careful consideration with respect to routing protocol deployment, traffic engineering, QoS, and multicast deployment, and can also be challenging when it comes to troubleshooting. Additionally, most layer 2 MP2MP implementations will treat multicast traffic the same as broadcast traffic, which results in very inefficient use of network resources. A more detailed discussion can be found in [7].

All the service definitions described above have immediate application in EttB scenarios for providing layer 2 VPN services. In addition, E-Line-R, E-Line-D, and E-LAN-M have applicability as access mechanisms to layer 3 VPNs that are typically based on MPLS connectivity between corporate sites, and also to EttS scenarios to provide access to layer 3 Internet services.

The expanded reach of the Ethernet service UNI can be achieved by interworking with existing services such as FR or ATM and the introduction of new technologies that enable native Ethernet transport in areas where it was not previously possible. A key area of focus is the first mile.

Ethernet access to layer 3 VPNs is a particularly attractive proposition as the layer 3 mechanisms relieve the SPs of the complexities of interworking at layer 2. It is irrelevant whether sites of the same VPN use FR, ATM, or Ethernet, as the layer 2 access mechanism is local to the PE router. This observation also applies to leased line and DSL access scenarios.

EttS scenarios usually provide IP services, such as Internet access, voice over IP, video on demand, broadcast video, and video telephony. This mix of bandwidth-hungry services is most commonly delivered by providing 10/100 Mb/s symmetrical Ethernet connectivity to individual offices or apartments from Ethernet switches located in multitenant units. These Ethernet switches are then connected via Gigabit Ethernet over fiber to the building (FttB) architectures to central office facilities. Where fiber deployment is not economically viable (e.g., single family homes), EttS scenarios can utilize DSL as an alternative first mile transmission technology.

Customized Service Delivery: Increased Degree of Customization

Enterprises require customized services in order to be able to align their networks with their applications and business objectives; these are implemented by defining SLAs between the customer and the SP. SLAs provide additional attributes to, and are an integral part of, the service specification. Examples of these attributes are availability of service, security attributes, and delay and performance characteristics of the service. As discussed in the introduction, customization is a means of differentiation that allows the SPs to develop new profitable services. In addition, SLAs allow for better network engineering and optimization of resources. The SLA attributes influence the architecture and deployment of the network. If, for example, the attribute *availability* is considered, different protocols, technologies, and equipment types offer different levels of resiliency and restoration times; network-level security is directly linked to availability because an unsecure network can suffer from denial-of-service attacks that may violate the availability SLA. The following sections briefly discuss some additional dependencies.

Traffic Separation

In many cases, customers specify the way they expect their traffic to be separated in the provider network. They restrict the level of transparency provided by the UNI definition. As such, they influence the way the SP implements the EVC. The customers' requirements for traffic separation will be met by one of two fundamental network paradigms: shared resource networks (SRNs) or dedicated resource networks (DRNs). SRNs generally, but not necessarily, separate customer traffic at layer 2 or above, while DRNs rely on layer 1 to do so. DRNs are implemented using layer 1 technology. SRNs share bandwidth among customers and can provide gains from oversubscription for the SP.

Security

Security is a primary consideration in any public network. SPs need to be able to ensure that different customers on a common infrastructure cannot affect each other, and denial-of-service attacks or other malicious actions cannot interfere with SLA compliance.

Although Ethernet brings tremendous flexibility to the service portfolio, SPs must consider the security implications of implementing Ethernet services. Traditional point-to-point WAN connections are easier to secure than multipoint networks based on Ethernet bridging mechanisms. Publicly available hacker software can enable users to exploit Ethernet mechanisms without any expert knowledge if these networks are not properly configured. Therefore, SPs should ensure that the robust security features available in state-of-the art Ethernet devices available today are employed.

Consistent End-to-End Quality of Service

QoS allows for tighter bandwidth control and stringent SLA delivery. QoS mechanisms can control traffic attributes such as delay, jitter, or packet loss to efficiently utilize network resources and ensure consistent application performance during congestion situations. Effective QoS enables customers to match the network to their business needs and helps SPs to differentiate themselves from competitors and scale capacity more effectively.

In point-to-point connections, traffic attributes can be enforced at each network ingress point, and bandwidth profiles can be defined, implemented, controlled, and measured. Layer 3 devices such as routers can offer a rich set of QoS features. For multipoint-to-multipoint layer 2 services, a meaningful QoS definition is more complex to define, and SPs need to enforce traffic contracts at both the network ingress and egress points. QoS also increases the transport efficiency of the network. Using intelligent packet processing together with QoS, SPs can oversubscribe their networks to make better use of their existing interfaces and bandwidth.

Availability

Availability is another important network requirement that can be achieved in several ways. Equipment redundancy can improve overall network availability. Using appropriate network protocols, such as IP routing protocols, MPLS, Fast ReRoute, Rapid Spanning Tree, EtherChannel, or SONET/SDH protection mechanisms, can improve the availability of a network service.

Geographic Service Reach

Service acceptance and take rates are very dependent on the geographic reach of the service. The expanded reach of the Ethernet service UNI can be achieved by interworking with existing services such as FR or ATM and the introduction of new technologies that enable native Ethernet transport in areas where it was not previously possible. A key area of focus is the first mile.

Many access networks today use one of the

E-line shared (E-Line-S)	E-Line-S is the Ethernet equivalent of the private line (or private wire) service. Routers, bridges, or hosts may be used as customer edge devices that establish point-to-point connections between two sites. An E-Line-S is characterized by all-to-one bundling, VLAN transparency, and PDU transparency. In the case where the service user wants to build a flat layer 2 network, E-Line-S can connect two customer equipment (CE) bridges and allow the VLAN architecture of the service user to be extended between sites.
E-line dedicated (E-Line-D)	E-Line-D is a special case of an E-Line-S. E-Line-D services exclusively use dedicated layer 1 transport (SONET/SDH or wavelengths) to implement the service. The E-Line-D service is introduced because several SPs are interested in differentiating services that are delivered over a packet multiplexing infrastructure (like E-Line-S) from those that are based on layer 1 transport (E-Line-D).
E-line relay (E-Line-R)	E-Line-R is the Ethernet equivalent of the FR service. Here routers or hosts are used as CE devices to establish point-to-point connections between two sites. Like FR, service multiplexing is key in E-Line-R as multiple logical point-to-point connections are provisioned on a single physical interface that allows multiple remotes sites to be reached from a single UNI. Instead of the data link connection identifier (DLCI) used in FR, E-Line-R uses VLAN IDs to identify the logical point-to-point Ethernet connections within a UNI and, as in FR, these values may be locally significant. Layer 2 control protocols may be peered, tunneled, or discarded by the network-side UNI, making E-Line-R a non-transparent or "opaque" service. In particular, PDUs may be dropped, due to IEEE 802.1q Spanning Tree operation. Typical uses of E-Line-R are hub and spoke enterprise connectivity or Internet SP (ISP) to customer connectivity.
E-LAN multipoint (E-LAN-M)	E-LAN-M is the WAN analog of the multipoint Ethernet LAN capability. E-LAN-M delivers each unicast frame to its designated destination site using standard Ethernet bridge self-learning and forwarding behavior. Broadcast/multicast frames or frames with destination unknown MAC addresses are replicated to all sites (except the originating site). In fact, it can be argued that the operation of an E-LAN-M emulates that of a layer 2 Ethernet switch. A CE can be a router, bridge, or host. The service is characterized by all-to-one bundling, VLAN transparency, and PDU transparency.
E-LAN multipoint and relay (E-LAN-MR)	Lastly, E-LAN-MR is a combination of E-LAN-M and E-Line-R. This special case is given its own name because of the need for a multipoint-to-multipoint service that can be service multiplexed. E-LAN-MR is an interesting option for customers who require multipoint layer 2 services but in addition to this would like to access point-to-point services such as Internet access over the same UNI.

Table 2. *Generic Ethernet services.*

different flavors of DSL (ADSL, ADSL2+, VDSL, or SHDSL) to achieve broadband connectivity. To maximize service footprint SPs combine these existing access network services with new Ethernet-based WAN services. Enterprises also desire end-to-end layer 2 transport services between disparate UNIs. Service interworking at layer 2, however, is not a trivial exercise. It requires far more than just the translation of different frame formats.

Each layer 2 protocol has a different frame format: Ethernet has layer 2 source and destination addresses; FR and ATM have a destination address only; FR and ATM have routed and bridged encapsulations; HDLC and PPP have no addresses. And to complicate things further, each layer 2 protocol has different address resolution processes: Ethernet, for example, uses IP ARP, which means that the target layer 3 address is known but not the layer 2 address. For ATM and FR multipoint interfaces it is the opposite, so Inverse ARP is used. The service interworking function must therefore mediate between the different address resolution mechanisms by spoofing the appropriate ARP/Inverse ARP responses and translate between the different frame formats.

As mentioned above, the necessity for interworking between different layer 2 technologies disappears when layer 3 VPN services are used.

While early deployments of Ethernet-based services primarily used native Ethernet transmission to the customers, these networks are increasingly being extended using other transport mechanisms in the access (e.g., DSL).

The 802.3ah Ethernet in the First Mile (EFM) task force will help Ethernet natively "bridge" the first mile for EttS as well as EttB services. The 802.3ah standard, which is expected to be ratified mid-2004, will include subscriber access network topologies and physical layers for P2P copper, P2P over a single optical fiber, and P2MP fiber. Operations, administration, and maintenance (OAM) for EFM will also be defined.

Service Capacity: High-Capacity, Low-capacity Services

Service capacity (i.e., the bandwidth requirement of a particular service) continues to be an important service attribute that has a major impact on the cost of the service and the way it is implemented. This is becoming particularly relevant for Ethernet-based SP networks since the service bit rates range from a few hundred kilobits per second up to 1 Gb/s and above. Networks that deliver high-capacity services follow different design approaches and have different scalability metrics than networks that deliver low-capacity services. Although the total number of high-capacity services per network might be orders of magnitude lower than the total amount of low-capacity services, the high-capacity services may determine the architecture and design of the network.

High-capacity services. Common services are

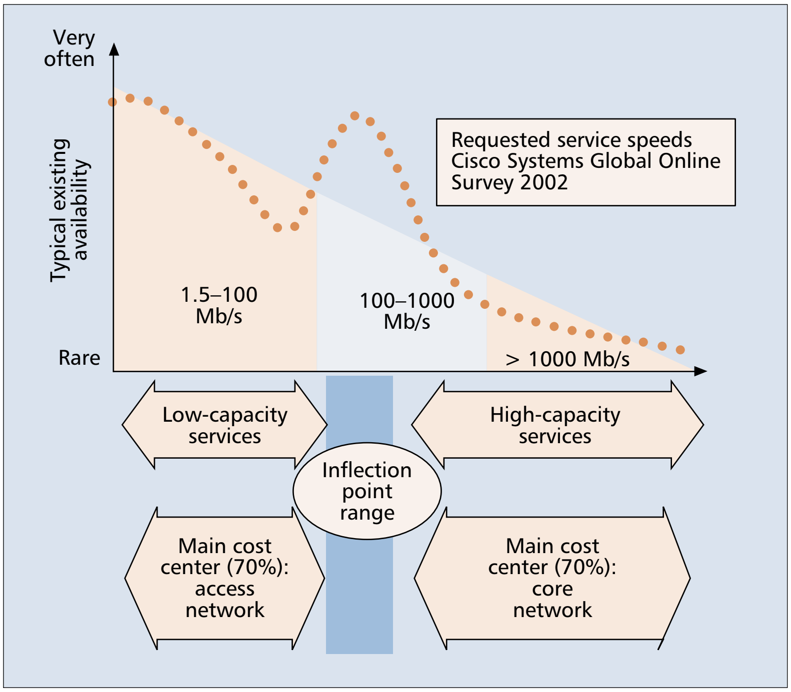

■ **Figure 4.** *Requested and available bandwidth in perspective service capacity and network cost centers from a qualitative perspective.*

fractional or full line rate Gigabit Ethernet private line, typically delivered using DRNs. These services are commonly requested by Fortune 1000 businesses, as well as government or educational institutions across regional, national, and even international boundaries.

Low-capacity services provide connectivity for many different types of business customers as well as residential customers and are typically delivered using SRNs. The number of customers using low-capacity services is generally expected to be orders of magnitude higher than those who use high-capacity networks. The resulting network architecture is required to support a high volume of geographically distributed customers.

While evaluating a network architecture design, SPs must consider how the cost of the network architecture changes when the number of devices increases, the number of customer locations increases, the geographical distribution of customers changes, or the bandwidth and service requirements for customers change/increase.

High-capacity services have a significant impact on the cost of aggregation and connectivity across the core, whereas low-capacity services have a much greater impact on the cost and scale of the access piece of the network. Modeling for numerous SP networks has demonstrated that when low-capacity services dominate the service mix, the access portion of the network equates to typically 70 percent of the total system cost. Alternatively, when high-capacity services dominate, the situation is reversed, with typically 70 percent of total system cost being associated with the core and distribution.

What has been seen using service cost modeling is that 100+ Mb/s per connection becomes a major *inflection point* in looking at network cost and scalability. Figure 4 correlates the inflection point to the availability of existing services, as well as the capacities customers expect from new offerings (as per a study performed in 2002). The inflection point also depends on the physical bit rates present in the core network to support the aggregated service bandwidth. As such, high-capacity services are a major consideration for the currently deployed networks if their number grows significantly.

Qualitatively, this effect can be explained based on the level of statistical multiplexing that can be achieved. It can be seen that the bandwidth of the high-capacity services is a large proportion of the backbone bit rates. In this case the law of large numbers, which is the basis for statistical averaging, no longer holds. This means that the transmission systems carrying high-capacity services cannot be oversubscribed to any significant degree without violating QoS SLAs. Therefore, high-capacity services are very often carried using dedicated resources. This is in contrast to low-capacity services that are a small proportion of the core bandwidth. The law of large numbers allows for very effective statistical multiplexing and a high degree of oversubscription without violating traffic contracts.

In designing next-generation networks, it is imperative for the SPs to identify the services they are planning to sell, and design the network with the scalability to meet those service demands. When determining the scalability requirements, it becomes clear that the SP must design architectures that can support high-capacity services while still being able to offer low-capacity services cost effectively.

Service Integration: Evolving the Installed SDH/SONET Base

Evolution Scenarios

Many SPs have an installed base of SDH/SONET transmission systems to provide leased line services based on TDM connectivity. With the emergence of Ethernet-based WAN services, there is an obvious need to leverage this installed base as much as possible. Based on service considerations outlined above, an analysis of this installed base needs to be undertaken before determining the evolution of these platforms. Such an analysis has to take into account many factors, including:
- The service requirements in terms of bit rates, projected over time
- Topologies and traffic flows
- Fiber availability
- Existing traffic

Capabilities of the existing network elements in terms of:
- Upgrade to higher bit rates
- Granularity of bit rates
- Enhancement by Ethernet interfaces

Depending on the outcome of this analysis, different scenarios can be envisaged.

Leverage the Installed Base Only

When an installed base of SDH/SONET equipment with sufficient spare bandwidth is available, it can be exploited by using external equipment to present TDM interfaces transporting Ethernet frames to the transmission platforms.

This external equipment can be in the form of customer premises equipment or central-office-based systems that provide both native TDM and Ethernet interfaces. The Ethernet interfaces will be mapped to SONET/SDH virtual containers that can be concatenated (contiguously or virtually) to provide the required bit rates. The network service delivered is a DRS that typically provides E-Line-D services. The architecture does not allow for statistical multiplexing among the individual traffic streams as each stream is carried within a fixed portion of bandwidth and as such is best suited for high-capacity services.

ENHANCE THE INSTALLED BASE BY ADDING ETHERNET INTERFACES

Some recent platform developments allow sharing of bandwidth in the transmission network by implementing Ethernet switching functionality in the Ethernet modules that plug into the transmission systems. This functionality allows tributary traffic on the Ethernet interfaces to be add/drop multiplexed and mixed with pass-through traffic on the same virtual container. Usually, these switching modules are arranged in the form of shared capacity rings. Control of the ring structure can be performed by either established mechanisms for Ethernet networks (e.g., Rapid Spanning Tree) or the emerging IEEE 802.17 Resilient Packet Ring (RPR) mechanism.

BUILD AN OVERLAY NETWORK FOR NEW SERVICES

In cases where an existing transmission network does not have enough spare capacity for new Ethernet-based services, or cannot be upgraded, it can be economically viable to keep the existing traffic on the installed network and build a parallel network optimized for the required services. In case of fiber shortage the existing transport network can share the fibers with the new network using DWDM mechanisms.

HYBRID APPROACHES

It is entirely feasible to combine the different evolution scenarios based on the technical and economic considerations outlined. Analysis will be required to determine the best architecture for a given environment.

SUMMARY

By deploying a broad array of technology and transport options, SPs can build highly scalable and flexible networks that adapt quickly to changing market conditions.

A high degree of flexibility ensures that the network can offer customized services that meet the business objectives of the customer; it can been seen that the "services drives network" approach discussed here has a direct bearing on how transport networks are deployed.

Looking to the future, three major trends have been identified:

- An organizational shift whereby the SPs will shift their organizational model from vertical to horizontal
- A network paradigm shift whereby service considerations will drive the transport network, rather than transport driving service selection and deployment
- An equipment capability shift where integrated platforms that support hybrid network designs will tightly integrate IP/MPLS/Ethernet and optical transport to leverage scalability mechanisms at layers 1, 2, and 3 rather than be limited to scalability per layer.

Ethernet provides a unique foundation for the delivery of service-driven networks. Ethernet as a WAN service is a challenge and an opportunity at the same time.

It is entirely feasible to combine the different evolution scenarios based on the technical and economic considerations outlined.

REFERENCES

[1] MEF, "Metro Ethernet Services — A Technical Overview," White paper, http://www.metroethernetforum.org/metro-ethernet-services.pdf
[2] IEEE 802.1ad Provider Bridges project; http://www.ieee802.org/1/pages/802.1ad.html
[3] IETF Layer 2 VPN WG; http://www.ietf.org/html.charters/l2vpn-charter.html
[4] IETF Pseudo Wire Emulation Edge to Edge (PWE3) WG; http://www.ietf.org/html.charters/pwe3-charter.html
[5] ITU-T Draft Rec. G.8010/Y.1306, "Architecture of Ethernet Layer Networks."
[6] P. Gandhi and B. Klessig, "Metro Ethernet WAN Services and Architectures," *Int'l. Eng. Consortium's Annual Rev. of Commun.*, June 2003.
[7] F. Brockners, N. Finn, and S. Phillips, "Metro Ethernet — Deploying the Extended Campus Using Ethernet Technology," *Proc. IEEE LCN*, Oct. 2003.

ADDITIONAL READING

[1] IETF Layer 3 VPN) WG; http://www.ietf.org/html.charters/l3vpn-charter.html

BIOGRAPHIES

LOTHAR ZIER [M] (lothar.zier@ieee.org) is a networking consultant at T-Systems/Detecon International GmbH (a subsidiary of Deutsche Telekom), Germany. He received his diploma in computer science in 1991 from the University of Bonn. For 10 years he worked for the German National Research Center for Information Technology (GMD) in the field of national and international IP and ATM-based networks. Currently he is involved in the design and implementation of an MPLS-based carrier network.

FRANK BROCKNERS (fbrockne@cisco.com) is focusing on Ethernet and optical technologies in Cisco's corporate technical marketing organization. Earlier he worked in several roles, including product management for cross-platform technologies in metro DWDM, product marketing management for EMEA, as well as networking consultancy. He has authored several papers, published in journals and in the proceedings of international conferences. He holds a diploma degree in electrical engineering (University of Technology, Aachen, Germany) as well as a Ph.D. degree in information science (University of Cologne, Germany).

WOLFGANG FISCHER [M] (wfischer@cisco.com) is responsible for technical consulting in Cisco Systems EMEA, in the area of broadband access solutions and technologies, including optical and storage technologies. Before joining Cisco in 1999 he was responsible for system architectures of ATM core and access products at Siemens AG in Germany. He holds a Ph.D. in electrical engineering from the University of Stuttgart, Germany. His professional interests include the entire spectrum of broadband access technologies and quality of service aspects.