Hybrid Transport Solutions for TDM/Data Networking Services

Enrique Hernandez-Valencia, Lucent Technologies

ABSTRACT

There is a growing demand for native data transport services for enterprises and corporations across public transport networks. Recently, equipment vendors have begun to incorporate a variety of LAN and storage area network interfaces, notably Ethernet, Fibre Channel/FICON, and ESCON, on traditional metro and long-haul transport equipment. Embracing Ethernet and SAN technology enables the introduction of flexible high-capacity transport services optimized for data networking. Transport operators may thus offer both enterprise-centric connectivity services, such as transparent LAN connectivity and virtual LAN services, as well as traditional bandwidth services, such as private lines, while preserving the operations and management infrastructure of the existing public networks. In this article we discuss the benefits of a hybrid Ethernet/TDM transport solution.

INTRODUCTION

With the onset of the information age, a wide variety of specialized connectivity, storage, content, and data distribution/processing services have sprung across the telecommunications/data communications landscape. These services vary from traditional telephony and private line connectivity services over the public switched telephone network (PSTN) infrastructure, to virtually switched circuits for WAN data networking (from most major telecommunications carriers), also to IP-oriented virtual private networks (VPNs) and residential/business Internet access services offered by Internet service providers (ISPs), and also to a new breed of Web hosting/storage and information processing services offer by application/storage service providers (ASPs/SSPs).

In order to address the need for connectivity, capacity, and content arising from the information age, service providers are quickly converging to a new view of the telecommunications and data communications business. Connectivity, capacity, and information services are each dealt with as one of the various services to be supported by a common public transport infrastructure. The relentless growth in bandwidth, connectivity, and content demand fueled by the Internet revolution has made bursty information transfer the main application driver of modern communications systems. As the packet switching technology that constitutes the basis of modern data communications networks matures, opportunities for integration with the common transport, multiplexing, and switching functions within the public transport infrastructure become more clearly defined. Fibre Channel, FICON, and ESCON are the most prevalent technologies deployed for storage networking. Ethernet is the de facto enterprise network standard for data communications. Extending enterprise and data center connectivity over the metro access/core network infrastructure is an obvious first step in deploying connectivity services over traditional transport networks. This new emerging paradigm is illustrated in Fig. 1. The transport network is multiservice. It enables a variety of basic connectivity services for both circuit and data applications over native data interfaces. Sophisticated data and content services that require finegrained peruser control are delegated to an intelligent services layer. For this services layer the transport network provides traditional traffic aggregation and distribution services on a wholesale basis.

SERVICE CONVERGENCE AS DRIVERS FOR HYBRID TRANSPORT SYSTEMS

The increased multiservice nature of a converged transport network infrastructure requires efficient handling of narrowband, wideband, and broadband traffic sources whether from voice, Web pages, electronic data exchanges, packetized digital audio, or video. Furthermore, service providers demand an enduring public transport network infrastructure that provides a flexible and affordable service evolution path despite unpredictable traffic patterns, service models, and technology evolution. Revenue generating services (e.g., private lines) must be preserved at the same level of expectation for survivability, manageability, and maintenance.

 0163-6804/02/$17.00

One trend toward a converged public transport network evolution is to convert every transport node in the public transport infrastructure into a packet switching device. Such an approach is attractive to greenfield operators, to smaller metro-oriented carriers with no other embedded communications infrastructure (e.g., metro C-LECs), or to new entrants with a highly specialized portfolio of services (e.g., E-LECs). This approach has turned out to be more complicated than initially envisioned given the configuration, operations, and management complexity associated with most packet-switched services, the relatively high cost of deployment of such an integrated services transport paradigm over the public switching/cross-connect facilities, and the narrowly focused customer base. It also carries a higher degree of uncertainty given the desirability to support profitable telecom services such as voice and private lines, and the relative immaturity of such transport services over packet-switch-centric technologies (other than asynchronous transport mode, ATM [1]).

Another trend is to include packet transport, multiplexing, and switching capabilities on synchronous optical network (SONET)/synchronous digital hierarchy (SDH) add/drop multiplexers (ADMs) and broadband crossconnect systems (BXCs). The goal here is to incorporate basic packet transport capabilities that help enable packet-oriented connectivity services over the existing transport infrastructure rather than on a fully collapsed transport, services, and applications layer. (For such an approach there are already well-defined transport mechanisms, e.g., ATM). Such an integrated TDM/data transport approach is attractive to established carriers as they can deploy new packet-switching technology to implement data transport services on an as-needed basis. It also facilitates the controlled introduction of operations, administration, management, and provisioning (OAM&P) procedures for those new services, and reuses the existing transport capabilities of deployed TDM networks.

Emergency of Ethernet Transport Services

Over the last few years we have seen Ethernet emerge as the dominant technology for LANs and enterprise networking. Ethernet has also begun to make inroads as a networking solution for storage facilities in corporate/hosting collocated data centers, and as an interconnect solution among ISP points of presence (POPs) in metropolitan networks. Among the drivers for Ethernet's popularity are its relative simplicity, maturity, and volume of sales (with corresponding lower manufacturing costs), particularly for short/medium reach data connections (up to a few kilometers). For example, compared to traditional packet over SONET/SDH (POS) interfaces, Ethernet interfaces can be as much as 50–80 percent lower in price (albeit with more limited OAM&P capabilities than typically required by public-grade telecommunications services). For best effort and non-mission-critical traffic, which dominates most corporate/enterprise data traffic today, TDM and ATM networks provide a level of performance, reliability, and service features beyond that required by these applications. Ethernet-based solutions, particularly on the access portion to the MAN, can fulfill this particular end user's needs.

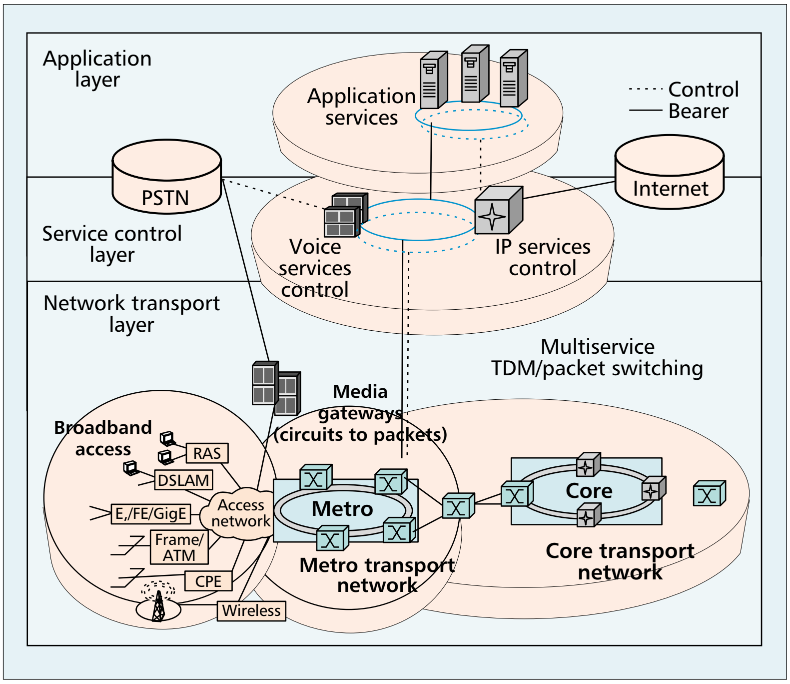

■ **Figure 1.** *The emerging converged network model for public transport networks.*

Networks built upon SONET/SDH and wavedivision multiplexing (WDM) are optimized for delivery of reliable cost-effective transport for voice, private line, and other mission-critical services that continue to dominate the access network revenue stream today. However, pure SONET/SDH and WDM networks are not yet optimized for addressing all the data transport needs. For instance, it has not been until recently that TDM solutions have been enhanced with flexible traffic aggregation and multiplexing mechanisms, or the granular bandwidth allocation schemes required for data communications. These shortcomings related to data transport have begun to be addressed with next-generation hybrid network architectures.

A Layered Hybrid Network Architecture Model

A recent development in service convergence is to integrate Ethernet technology into public transport networks. Support for standard Ethernet interfaces directly on network elements such as SONET/SDH ADMs/BXCs and dense WDM (DWDM) optical line systems (OLSs) enables native service interfaces for data transport. This approach exploits low-cost data interconnectivity on the enterprise equipment while maintaining the reliable and manageable transport infrastructure required in large public networks. We refer to this technology as *hybrid transport*.

Figure 2 depicts the functional layering model of the data transport capabilities in a hybrid Eth-

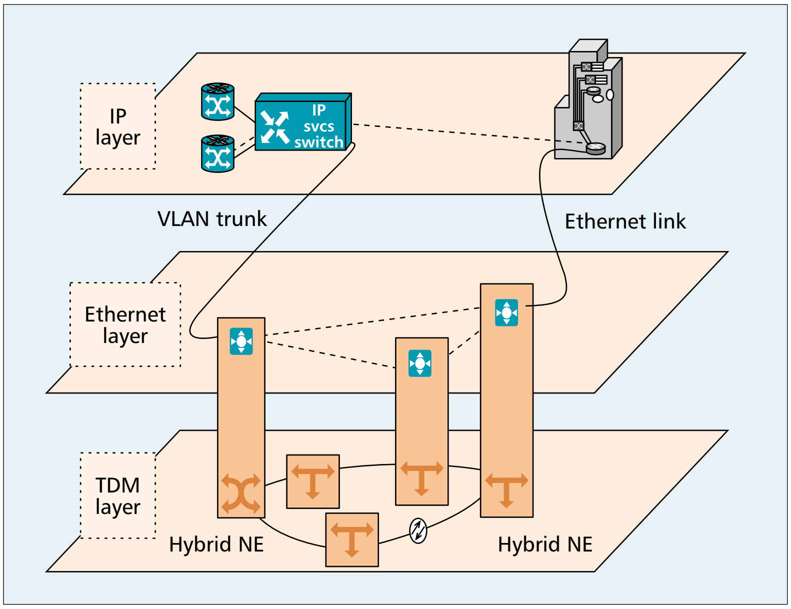

■ **Figure 2.** *The layered TDM/Ethernet transport model.*

ernet/TDM architecture. At the bottom of the hierarchy is a standard transport network, consisting of SONET/SDH ADMs and BXCs, that supports traditional time-division multiplexed (TDM) services such as voice and private lines. Incorporated into these network elements are Ethernet/IEEE 802.3-based connectivity [2], IEEE 802.1D/w [3] bridging functions, as well as IEEE 802.1Q/p-based virtual LAN (VLAN) network services and QoS capabilities [4]. These features facilitate the logical overlay of data-aware unicast, multicast, and broadcast transport services currently not available over a SONET/SDH network infrastructure. In this manner the transport network can easily be configured, and enhanced, to offer not only native data interfaces but also native transport service already familiar to the data communications community.

A number of new data transport services can now be offered to enhance the operator's revenue stream. Sample services include virtual leased lines and Ethernet-based virtual private networks (VPNs). The implementation model for these services is discussed later.

Enabling Ethernet Transport over SONET/SDH

Three key technologies help enable storage networking and Ethernet transport over SONET/SDH networks: virtual concatenation of SONET/SDH paths, virtual bandwidth allocation via the Link Capacity Adjustment Scheme (LCAS), and the Generic Framing Procedure (GFP) to adapt the MAC frames (e.g., the IEEE 802.3/Ethernet frames) to the octet synchronous SONET/SDH payload. A functional view of such a hybrid TDM/Ethernet network element model is illustrated in Fig. 3. How these three mechanisms interwork to support the implementation of data transport over SONET/SDH is discussed next with a focus on Ethernet solutions. The same concepts can be extrapolated to storage networking.

Virtual Concatenation

SONET/SDH systems were initially optimized for the transport of telephony services. Given the constant bit rate (CBR) nature of voice and private line traffic, a coarse fixed-rate multiplexing hierarchy was most efficient for the transport of these CBR signals. Data traffic, however, is inherently bursty. The bulk of this traffic demands elastic bandwidth allocation. This demand can easily be accommodated with a best effort delivery service. Statistical multiplexing, via packet switching technologies, provides far better utilization of the transport medium for this type of applications. Neither the signal rate nor the nominal data rate of popular physical interfaces for data networks makes efficient use of the existing SONET/SDH channel sizes. A flexible mechanism to interact with the SONET/SDH multiplexing hierarchy was required.

Virtual concatenation [5, 6] is an inverse multiplexing technique that combines an arbitrary number of SONET/SDH transport channels to create a single-octet synchronous byte stream. It is an alternative to standard contiguous concatenation, which only supports aggregation and multiplexing in $4^N \times$ STS-3cs (SONET) or $4^N \times$ VC-4 (SDH) containers. With virtual concatenation, network operators can bundle an arbitrary number (X) of either low-order (e.g., VC-12s or VC-3s in SDH or VT1.5s in SONET) or high-

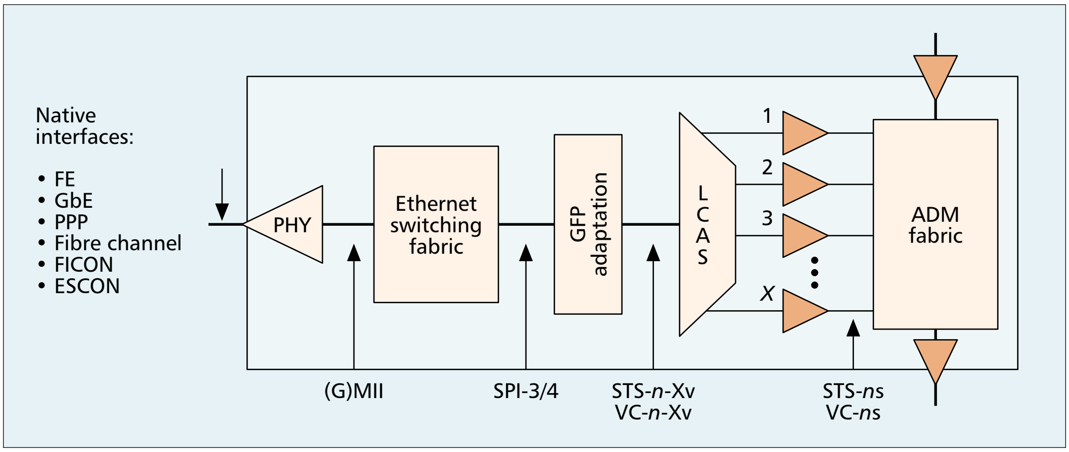

■ **Figure 3.** *The hybrid ADM/Ethernet network element model.*

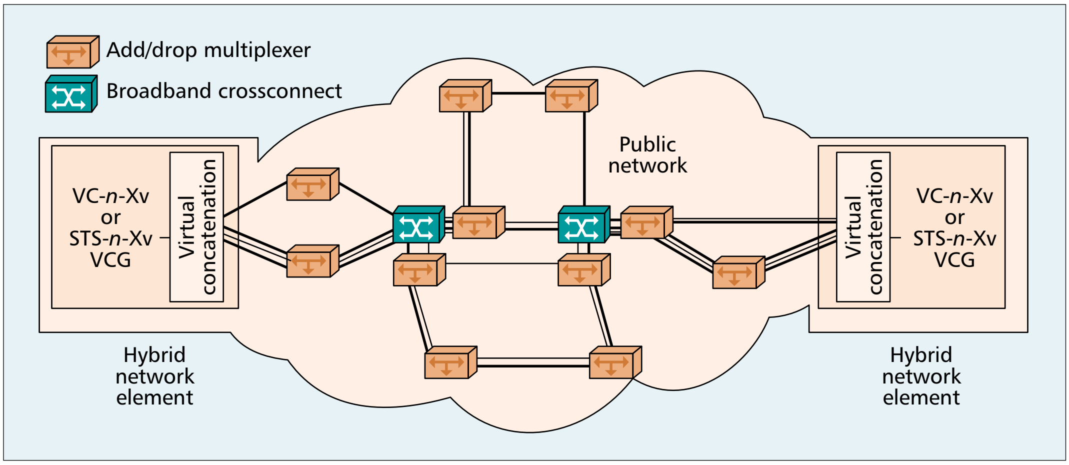

Figure 4. *Transport of constituent VCG components over SONET/SDH.*

LCAS is an extension to virtual concatenation that allows dynamic changes in the number of SONET/SDH channels in a connection under management control of the initiating/terminating network elements such that hitless performance of the VCG is guaranteed.

order (e.g., VC-4s in SDH or STS-1s/STS-3cs in SONET) channels to create a single virtual concatenation group (VCG) signal (e.g., VC-12-Xv/VC-3-Xv/VC-4-Xv in SDH or VT1.5-Xv/STS-1-Xv/STS-3c-Xv in SONET). An important aspect of virtual concatenation is that the individual transport paths that constitute the VCG can be transported independently over the SONET/SDH network. As illustrated in Fig. 4, only VCG initiating/terminating equipment at the edge of the transport network (typically implemented in a line card) needs to support this function. Virtual concatenation works seamlessly with legacy SONET/SDH equipment. The rest of the transport network simply transports the component TDM channels independent of each other. In addition, virtual concatenation provides mechanisms to manage not only the constituent paths of the VCG, but also compensation for the differential delays among those paths across the SONET/SDH network. Thus, virtual concatenation addresses bandwidth allocation constraints associated with the coarse multiplexing hierarchy of traditional SONET/SDH systems.

With virtual concatenation bandwidth can be allocated as needed to accommodate the precise bandwidth requirements of the end systems. For example, an enterprise might need a 100 Mb/s Ethernet pipe to interconnect sites in a given metro area. This task can be accomplished by allocating either 2 STS-1 channels (SONET) or VC-3 channels (SDH), and then combining these two channels into a single VCG as a VC-3-2v byte stream of roughly 2*48 = 96 Mb/s. This is s substantial improvement, about 33 percent, over allocating a conventional VC-4 path at roughly 155 Mb/s for the same purpose, and hence without the associated waste of the unused channel capacity. The local connection at the enterprise site can be done with conventional Fast Ethernet interfaces.

Better yet, virtual concatenation affords network operators with a new mechanism to provide value-added connectivity services, such as fractional or subrate Ethernet transport services. Here, enterprises may attach to the transport network with an inexpensive short-reach Gigabit Ethernet interface. However, customers may only request enough transport capacity to meet the anticipated interoffice traffic volume, say 150 Mb/s. The network operator may configure such service by only allocating a single VC-4 or 3 STS-1s to the associated VCG between the two sites. When demand goes up, the enterprise may request that the capacity of the VCG be upgraded, say in VC-4 or STS-1 increments, until the 1 Gb/s limit is reached for the available GigE interface (and assuming the additional transport resources are available).

Link Capacity Adjustment Scheme

Modifying VCG size by adding or removing constituent channels may render the data path useless if proper coordination among endpoints is not provided. LCAS [7] is an extension to virtual concatenation that allows dynamic changes in the number of SONET/SDH channels in a connection under management control of the initiating/terminating network elements (NEs) such that hitless performance of the VCG is guaranteed. LCAS also allows dynamic removal (addition) of failed (recovered) constituent paths. Channels can be added or removed by management actions while in service. The VCG capacity modifications will occur without scheduling any facility downtime to reconfigure the data service and without losing any customer traffic.

VCG/LCAS provides the equivalent of an intelligent link aggregation facility for SONET/SDH much in the same way the IEEE 802.3ad specification provides link aggregations facilities for Ethernet segments. It also allows the implementation of connectivity services with graded levels of performance, (e.g., higher transport throughput when there is no failure in any of the constituent channels). When there is a failure in one of the constituent channels, the available bandwidth will be lower without incurring complete failure of the transport service. This is achieved by ensuring that only the failed channels of the VCG are withdrawn from service while the remaining channels will continue carrying live customer traffic.

The VCG/LCAS approach is advantageous for QoS transport services such as DiffServ for IEEE 802.1Q/p VLANs or IP/MPLS networks [8]. In such networks, packets are appropriately

Virtual concatenation by itself is not sufficient to create a transport link that fits the exact bit rate of the native data signal into the SONET/SDH payload areas. A mechanism is still needed to map the native bitstream into the SONET/SDH channel, providing for signal rate adaptation and minimal OA&P functions.

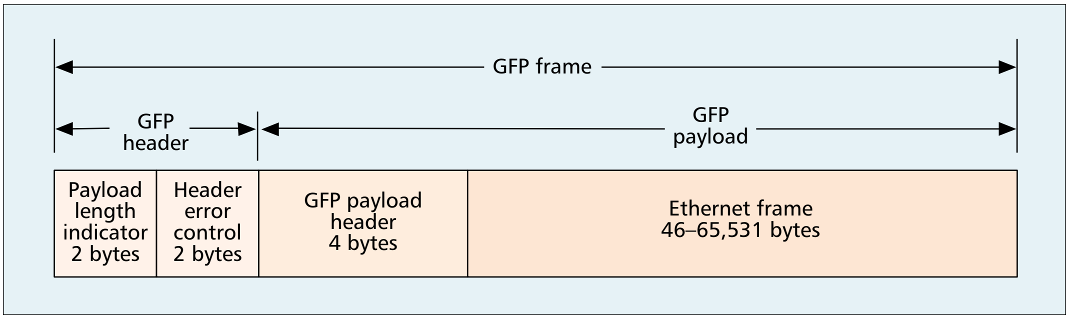

Figure 5. *The frame format for Ethernet over GFP using a null extension header.*

classified and marked to reflect different levels of handling priority with respect to packet loss and resiliency. An IP router or Ethernet VLAN switch will notice that less bandwidth has become available on an LCAS-enabled link experiencing a channel failure event. If temporary congestion arises from such an event, only the traffic with lower handling priority would be affected. In addition, for IP-based networks, the logical network topology need not be affected by such a change because IP-level connectivity is still maintained. Therefore, IP routing protocols do not need to reconverge; hence, no service interruption occurs.

In addition, traditional SONET/SDH protection schemes for highly reliable transport services may still be enabled. These are implemented as top-quality protection services over the SDH/SONET or WDM layer. They further allow service providers to offer graded levels of protection services and treat different traffic sources according to their revenue cost structure.

Transporting Packets in Circuits: The Generic Framing Procedure

Virtual concatenation by itself is not sufficient to create a transport link that fits the exact bit rate of the native data signal into the SONET/SDH payload areas. A mechanism is still needed to map the native bitstream into the SONET/SDH channel, providing for signal rate adaptation and minimal OA&P functions. The Generic Frame Procedure (GFP) fulfills this role. GFP is a lightweight adaptation protocol that provides a flexible mechanism to map different bitstream types to an octet-synchronous channel. The adaptation mechanism is frame-based and allows the segmentation of the physical channel into fixed or variable size containers, GFP frames. Two modes of signal adaptation are provided with GFP.

The transparent-mapped adaptation mode (currently defined for 8B/10B encoded signals only) is particularly suitable for full-rate point-to-point applications. (By full-rate is meant that the entire capacity of the local physical interface is supported). Adaptation is accomplished by mapping link-layer codewords into GFP frames. This mode is intended for applications that seek to emulate a native physical interface with very strict packet delay, loss, and throughput requirements (e.g., Fibre Channel, FICON, and ESCON).

The frame-mapped adaptation mode is a more flexible adaptation mode that is suitable for either full-/subrate point-to-point and multipoint applications. Adaptation is accomplished by mapping upper-level protocol data units (PDUs), such as Point-to-Point Protocol (PPP) frames or IEEE 802.3 MAC frames, rather than link-layer code words, into the GFP frames. The frame structure for mapping an Ethernet/IEEE 802.3 frame on a GFP frame (assuming a null extension header) is illustrated in Fig. 5. For applications where both the transport and bridging capabilities of Ethernet are integrated into the transport NEs, the frame-mapped mode is the preferred mode of adaptation since the physical layer aspects of both the SONET/SDH and Ethernet interfaces (layer 1) are segregated from the media access control (layer 2) aspects. Since the same mode of adaptation is applied to either point-to-point or multipoint configurations, service providers can deal with these two styles of application with the same provisioning and management procedures. Thus, for instance, if a customer wishes to migrate from a point-to-point transport service to a multipoint transport service, both of these services can be delivered from the same service interface without further reconfiguration of the preexisting endpoints.

Although many proprietary mechanisms abound for adaptation of native data traffic into SONET/SDH channels, GFP is the only international standard supported by both the American National Standards Institute (ANSI) and International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) [9]. GFP is a highly efficient encapsulation protocol with a fixed, but small, overhead per packet. Unlike most framing protocols, GFP scales very well to higher transport rates, which is one of the reasons GFP is being widely adopted for various high-speed applications such as optical channels in ITU-T optical transport network (OTN) architecture [10].

Hybrid Ethernet Transport Services

Based on the transport enhancements to SONET/SDH aggregation and multiplexing just described, new data-centric connectivity services can easily be instantiated over the public transport network infrastructure. Basic hybrid Ethernet transport services can be classified, in terms

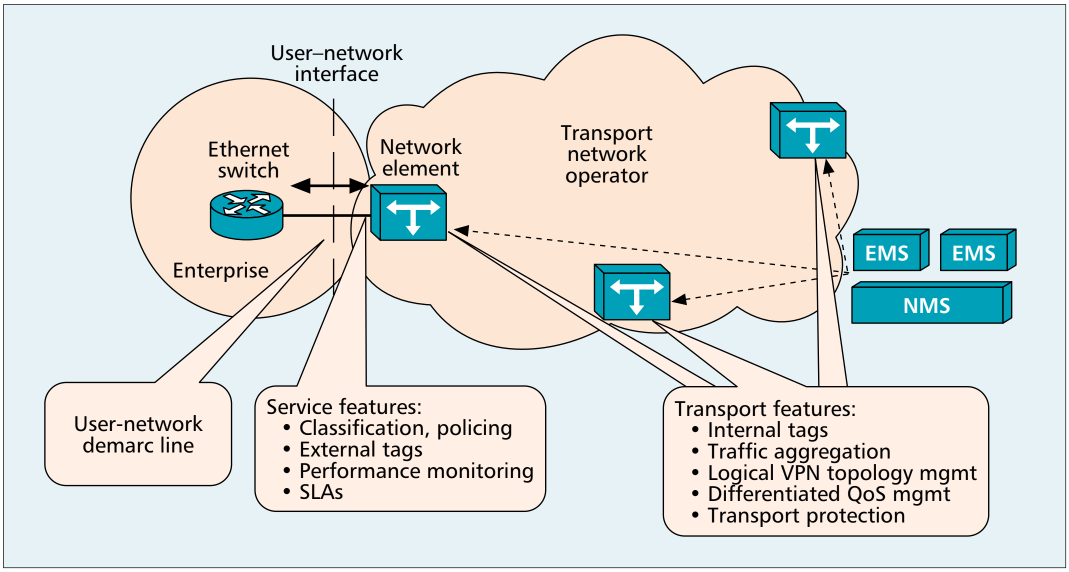

Figure 6. *A reference model for services over the public transport network.*

of transport feature complexity, into three generic groups, namely:

- Ethernet private line services
- Transparent LAN interconnect services
- Access to managed IP services

Below we highlight service and transport features that can be associated with those services. For the purposes of discussing these services it is useful to think of the transport network as a black box. The network elements must distinguish between end-user and network transport services. End-user service attributes cover service characteristics negotiated between the customer and the service provider (typically using a user-network signaling or management interface). Transport network services cover service attributes that are based on well-known transport and networking mechanisms, and enable the delivery of the contracted services according to negotiated service level agreements (SLAs). All traffic management capabilities would reside on this layer, and any sophisticated QoS mechanisms would be implemented on the packet-switching components of the NE. Initially such end-user and network transport services and features may be configured via element/network management systems (EMS/NMS). Such a view is illustrated in Fig. 6. In the future, it may be possible to request such services and features via a well-defined user–network interface. Efforts in that direction are already underway in the Metro Ethernet Forum.

Ethernet Private Lines

Ethernet private line (EPL) refers to the simplest of the Ethernet connectivity services. It offers point-to-point connectivity between two remote sites by emulating the transport service delivered by an Ethernet segment, but over the public transport network. This is a particularly useful service that can be used to extend the distance limitations of standard 10/100 Mb/s and 1 Gb/s Ethernet interfaces by reusing the connectivity services delivered by SONET/SDH paths over metro and long-haul networks. Such an approach would be substantially more efficient and cost effective than traditional solutions based on deploying optical transponders at both ends of dedicated lateral fibers to interconnect the end-user sites to the operator's infrastructure. The cost of the equipment and fiber access infrastructure is shared across multiple end users.

In the past transport of native Ethernet signals on SONET/SDH proved inefficient due to the capacity mismatch between the nominal signal (or

Ethernet private line (EPL) refers to the simplest of the Ethernet connectivity services. It offers point-to-point connectivity between two remote sites by emulating the transport service delivered by an Ethernet segment, but over the public transport network.

Traffic type	SONET		SDH	
	Contiguous	Virtual	Contiguous	Virtual
10 Mb/s Ethernet	STS-1 (20%)	VT-1.5-7v (89%)	VC-3 (20%)	VC-12-5v (92%)
100 Mb/s Fast Ethernet	STS-3c (67%)	STS-1-2v (100%)	VC-4 (67%)	VC-3-2v (100%) or VC-12-46v (100%)
200 Mb/s (ESCON)	STS-6c (66%)	STS-1-4v (100%)	VC-4-4c (33%)	VC-3-4v (100%) or VC-4-2v (66%)
1 Gb/s (FC/FICON)	STS-21c (85%)	STS-1-18v (95%)	VC-4-16c (35%)	VC-4-6v (95%)
1 Gb/s Ethernet	STS-24c (83%)	STS-1-21v (95%)	VC-4-16c (42%)	VC-4-7v (95%)

Table 1. *Transport efficiency of contiguous vs. virtual concatenation transport of Ethernet private line services over SONET/SDH.*

The simple EPL service model over a dedicated TDM transport channel supporting the peak interface rate can be enhanced in various ways to deliver a variety of value-added Ethernet connectivity services. As a starting point, operators may choose to support Ethernet transport services at subrates of standard 10/100 Mb/s and 1 Gb/s access interface rates.

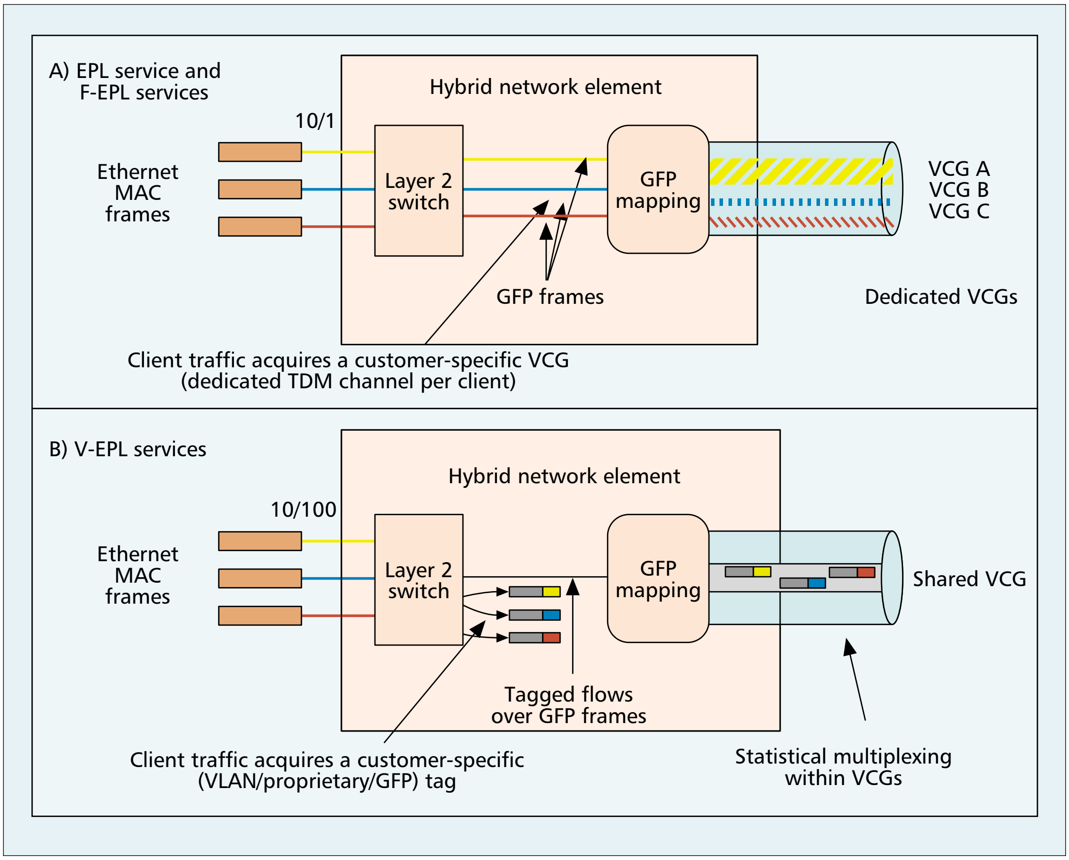

■ **Figure 7.** *Ethernet private line services (EPL, F-EPL, and V-EPL) over TDM channels.*

data) rates and the capacity of the SONET/SDH payload. This packing inefficiency is shown in Table 1. For instance, transporting a 100 Mb/s signal over a contiguously concatenated STM-3c/VC-3 would only result in 67 percent transport efficiency, that is, 50 Mb/s would be wasted for each 100 Mb/s connection. Virtual concatenation and GFP help address these past limitations in SONET/SDH payload granularity by providing transport containers closer to the nominal data rate and a low overhead mapping mechanism of the MAC frames into the SONET/SDH payloads. For the case of 100 Mb/s EPL, efficiency in the nominal data rate is close to 100 percent. (Note that these numbers assume the elimination of interpacket gaps, IPGs, which provides another 3–5 percent gain in efficiency for Ethernet payloads. They do not include bandwidth allocation for service management functions.)

Figure 7a illustrates a typical instantiation of EPLservices over a hybrid SONET/SDH network element. Each user flow is allocated to a segregated TDM channel. The TDM channel itself is composed of multiple STS-Nc/VC-N paths virtually concatenated into a VCG (VCGs A, B, and C in Fig. 7a). Each VCG would be sized to meet specific end-user requirements as specified in an SLA (typically the full capacity of the access line rate).

Fractional and Virtual Ethernet Private Lines — The simple EPL service model over a dedicated TDM transport channel supporting the peak interface rate can be enhanced in various ways to deliver a variety of value-added Ethernet connectivity services. As a starting point, operators may choose to support Ethernet transport services at subrates of standard 10/100 Mb/s and 1 Gb/s access interface rates. Such services are useful for enterprises that generate long-term traffic demands at a fraction of the links' peak rate. A fractional Ethernet private line (F-EPL) service can easily be implemented by configuring the transport path at a fraction of the access link rate. Typically, the path is configured at either VT1.5/VC-12 (1.5 Mb/s), STS-1/VC-3 (50 Mb/s) or STS-3c/VC-4 (150 Mb/s) granularity using standard virtual concatenation procedures.

Furthermore, operators may choose to configure arbitrarily sized transport paths and share the TDM capacity across multiple such F-EPL users. This approach allows service operators to achieve further transport efficiency via statistical multiplexing and graded levels of performance. Such virtual EPLs (V-EPLs) can easily be implemented using either GFP-based tags, such as those from the GFP linear extension header, IEEE 802.1Q/p VLAN tags (assuming these tags are not already in use by the end systems), or stacked VLAN tags by tagging the different flows with IEEE 802.3Q-compatible VLAN tags as illustrated in Fig. 7b. These features allow differentiated traffic treatment and QoS that can be indicated via the IEEE 802.3 user priority field. Note, however, that these services do require additional transport overhead for the virtual link identifiers that need to be accounted for as part of the traffic engineering requirements.

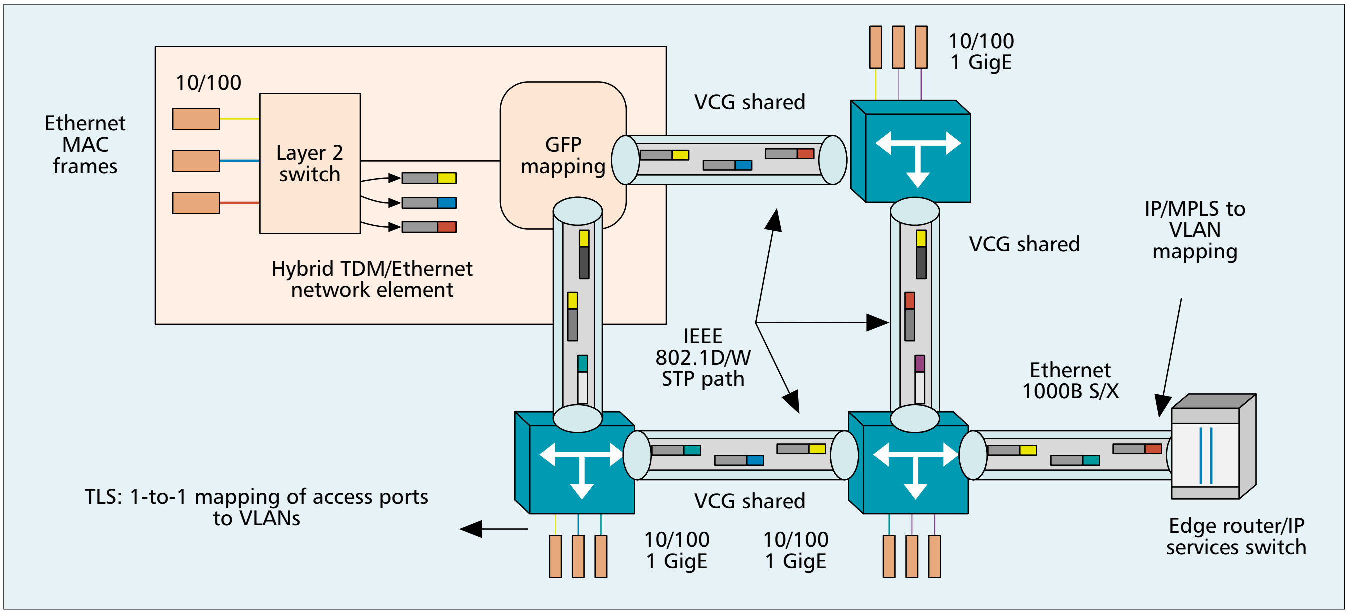

Figure 8. *Access to managed IP services via transparent LANs (over shared TDM channels).*

F-EPL and V-EPL services are intrinsically bursty packet-switched services that exploit the gaps in the end-user flow for statistical multiplexing gain. This style of service presumes at least a minimal set of resource management capabilities integrated into the transport network element (above the SONET/SDH transport layer) to help a service operator meet QoS commitments. QoS-based packet scheduling (e.g., Strict Priorities, Class Based Queuing, or Weighted Fair Queuing) and active queue management (e.g., Random Early Discard, RED, Band Weighted RED), typically not subject to standardization, are required. Other traffic management mechanisms may include flow control across the UNI via mechanisms such as IEEE 802.3x (Pause Frame) or shaping/policing of the user flows according to a committed information rate (CIR)/peak information rate (PIR).

Transparent LAN Interconnect Services

With the growing need to share resources across multiple enterprise sites over larger and larger distances, many enterprises find themselves faced with the need to interconnect their LANs using either private or public transport facilities. Often LANs are connected via a dedicated private line circuit (e.g., T1/E1s, DS3/E3s, or *n*64 kb/s) or packet switching technologies like X.25, frame relay, and ATM. These solutions are traditionally optimized for unicast services, and based on transport solutions different from the LAN technology prevailing in the enterprise environment that provide both unicast and multicast transport services. To address this limitation, it is simpler to integrate not only the Ethernet physical interfaces directly onto the network transport equipment, but also the media access control (learning and bridging) capabilities.

Transparent LAN services thus refer to a generic set of transport features that enable multipoint connectivity services to extend Ethernet learning and bridging functions over the public transport network. These services facilitate sharing of enterprise/corporate resources over metro access/core networks by emulating the transport services offered by the Ethernet/IEEE 802.3 MAC layer.

Multiple flavors of this service can be instantiated by exploiting transport features from either the TDM or IEEE 802.3 transport layers. In Fig. 8, GFP and virtual concatenation are employed to create traffic-engineered paths to interconnect Ethernet virtual switch instances. In its simplest instance, the service could be completely transparent with respect to any other end-user information other than the point of attachment of the various end-user sites to the transport network and the layer 2 devices reachable across that interface (for MAC address learning purposes). Any other information such as customer-generated BPDU frames (required for spanning tree configuration) or IEEE 802.3Q/p VLAN tags would be ignored by the transport network. In another instance of the service the network operator may share ownership of the IEEE 802.3Q/p VLAN tags with the end customers. CPEs could tag their local traffic with IEEE 802.3Q/p-based VLAN information to indicate to the transport network information about both a membership to a locally defined VPN or a desired QoS level, as illustrated in Fig. 8. The transport network could then use such information to forward traffic over the appropriate path across the public transport infrastructure, reclassify (even remark) user flows on ingress for traffic forwarding and multiplexing purposes and in accordance with contracted SLAs, or even encapsulate the user frames with stacked IEEE 802.3Q/p-like VLAN tags to convey end-user information unmodified across the transport network.

Access To Managed IP Services

The connectivity services defined so far operate strictly either at layer 1 (physical port) or layer 2 (VLAN tag). Often the main reason customers

The hybrid approach enables network operators to incorporate basic connectivity services such as Ethernet private/virtual lines, transparent LANs, and Internet traffic backhaul to cater to both point-to-point and point-to-multipoint connectivity needs.

contract transport services is to gain access to various kinds of information (content services) or the Internet. Dedicating a separate point-to-point connection for each customer service eliminates the need for the transport network to be aware of the higher-layer service associated with each traffic flow. However, this approach is also costly and painful to manage, for both users and service providers, since the various traffic types must be physically segregated in advance. It increases the port count requirement per node, as well as the number of cables to be dealt with to interconnect the CPE to both the transport network access equipment and the switches/routers at the ISP POP.

To limit the number of cables and ports on the access switch/router it is convenient to aggregate the data traffic from the various customers prior to handing off the traffic to the content/service provider. Integrating both layer 3–7 classification and virtual Ethernet switching/bridging capabilities in the transport network element enables SONET/SDH ADMs to perform both TDM and packet functions. The same network element can terminate different TDM channels, classify and tag each channel(s) into separate packet flows, and aggregate them into a single statistically multiplexed packet flow for the edge router. We refer to this service as a *VLAN trunking service*. Since multiple VLANs are exchanged at the handoff point, as illustrated in Fig. 8, an intelligent mechanism is required to map traffic flows to VLANs. Typically this function would be provided by an intelligent IP services switch that is aware of the IP services provided to the various end users. A VLAN trunk lowers the number of interfaces on the router and SONET/SDH multiplexer.

By enhancing simple transport features with more advanced packet classification functions it is possible to create a new set of value-added transport services that require very little additional routing and forwarding intelligence. This approach enables a flexible service layer architecture that taps into intelligent services devices deployed at the edges of the transport network. For instance, an enterprise may require both multipoint private LAN connectivity for disperse geological sites as well as connectivity to an ISP for Internet access. A network operator could instantiate two internal VLANs, one among all the enterprise ports for the private LAN interconnect service and one VLAN between the same ports and the ISP port for the Internet access portion of the service. Traffic from the enterprise ports can be mapped to either to these two VLANs strictly on L3-4 header information and without participation in the private enterprise of ISP routing protocols. It also does not require the transport operator to support more sophisticated label switching solutions such as ATM or MPLS.

These enhanced transport services afford a multitude of value-added transport services to be offered from a common service interface while allowing the traffic from a given enterprise site to reach different services offered in distinct geographical locations via simple connectivity and classification features from the transport network.

Conclusion

Hybrid TDM/Ethernet solutions enable the extension of native unicast, broadcast, and multicast data transport services, based on Ethernet switching, bridging, and networking capabilities, over a public SONET/SDH transport infrastructure. The hybrid approach enables network operators to incorporate basic connectivity services such as Ethernet private/virtual lines, transparent LANs, and Internet traffic backhaul to cater to both point-to-point and point-to-multipoint connectivity needs. Typical application scenarios that can benefit from such services include:

- Inter-POP connections
- Corporate LAN interconnection
- Ethernet VPNs
- Internet services access

Since data transport is compatible with the public SONET/SDH transport infrastructure, these value-added connectivity services can be provided beyond the immediate geographical area where the equipment is installed. Data services can be provided with reliability, interoperability, and manageability already found in today's public transport networks. The same approach can be applied to SAN technologies.

Acknowledgments

The author thanks the referees for their valuable comments and suggestions.

References

[1] ITU-T Rec. I.361, "BISDN ATM Layer Specification," 1993.
[2] IEEE 802.3, "Information Yechnology —Telecommunications and Information Exchange between Systems — Local and Metropolitan Area Networks —Specific Requirements –Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," 2002.
[3] IEEE 802.1D, "IEEE Standard for Information Technology —Telecommunications and Information Exchange between Systems — IEEE Standard for Local and Metropolitan Area Networks —Common Specifications — Media Access Control (MAC) Bridges," 1998.
[4] IEEE 802.1Q, "IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks.," 1998.
[5] ITU-T Rec. G.707/clause 11, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)," 2000.
[6] ANSI T1.105/clause7.3.2 (ed. 2001). "Synchronous Optical Network (SONET): Physical Interfaces Specifications."
[7] ITU-T Draft Rec. G.7042/Y.1305, "Link Capacity Adjustment Scheme (LCAS)," 2001.
[8] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF RFC 3031, Jan. 2001.
[9] ITU-T Draft Rec. G.7041/Y.1303, "Generic Framing Procedure (GFP)," 2001.
[10] ITU-T Rec. G.709, "Interfaces for the Optical Transport Network (OTN)," 2001.

Biography

Enrique Hernandez-Valencia [M] (enrique@lucent.com) is a distinguished member of technical staff at Lucent Technologies Bell Laboratories. He received his B.Sc. degree in electrical engineering from the Universidad Simon Bolivar, Caracas, Venezuela, and his M.Sc. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena. Since joining Bell Laboratories in 1987, he has worked in the research and development of high-speed data communications protocols and systems. He is a member of the ACM and Sigma Xi.