The Building Blocks of a Data-Aware Transport Network: Deploying Viable Ethernet and Virtual Wire Services via Multiservice ADMs

Enrique Hernandez-Valencia, Lucent Technologies
Gady Rosenfeld, Corrigent Systems

ABSTRACT

SONET/SDH technologies constitute the core transport infrastructure of major telecom service providers worldwide. As the percentage of packet-oriented traffic in the overall traffic demand continues to rise, prompted by the widespread adoption of the Internet Protocol suite, and recently by the fast adoption of Ethernet services, there is increasing pressure to improve the service provider's transport infrastructure in ways that make it data-aware and cost-effective for packet-oriented applications. Steps in this direction include the adoption of native physical interfaces, for Ethernet and storage area networks as service interfaces, or full integration of packet switching capabilities from Ethernet, Resilient Packet Ring, and MPLS technologies. This article discusses the emerging building blocks for next-generation data-aware transport networks and next-generation transport network elements.

INTRODUCTION

Today's metro transport networks were designed to carry fixed-rate telephony and private line traffic in a reliable and resilient manner. Over the last decade, service providers have experienced tremendous growth in demand for data-oriented services; today, in most carriers' networks the aggregate demand from the various data traffic types has started to become the dominating traffic profile on service providers' transport networks. Most of this data traffic is carried over synchronous optical network/synchronous digital hierarchy (SONET/SDH) or time-division multiplexed (TDM) circuits as layer 2 encapsulated traffic, such as IP traffic encapsulated in high-level data link control (HDLC)/packet over SONET [PoS], asynchronous transfer mode (ATM)/ATM adaptation layer 5 (AAL5), or frame relay (FR)/HDLC.

Ethernet is the dominant enterprise technology in the local area network (LAN), a fact that has also made it the preferred interface between the end user and the public network. Recently, new types of data services based on Ethernet technology, and Ethernet service interfaces, have been introduced by service providers. Three basic types of Ethernet transport services are readily identified. The first type, referred to as *Ethernet line services* by the Metro Ethernet Forum (MEF), addresses point-to-point connectivity applications including:

- Ethernet private line (EPL)
- Subrate (fractional and bursty) EPL (F-EPL)
- Ethernet virtual private line services

The second service type, referred to as *Ethernet LAN services* by the MEF, covers multipoint connectivity services based on packet-switched transport capabilities including:

- Virtual private LAN (PLAN) service
- Virtual private LAN services (VPLAN)

These services address transport features typically referred to as *transparent LAN services* (TLS) in the North American data communications market.

The third service type covers application-specific services that may include higher-layer (e.g., IP or multiprotocol label switching [MPLS]) or other transport layer (e.g., plesiochronous digital hierarchy [PDH] or SONET/SDH) components. The application-specific service include:

- Ethernet-enabled Internet access
- Access to managed IP service
- TDM circuit emulation services

These Ethernet-friendly service models create new challenges for data service providers. A

0163-6804/04/$20.00 © 2004 IEEE

multiservice add/drop multiplexer (ADMs) that can handle either constant bit rate or bursty traffic represents a new flexible approach to address the service provider challenge in delivering data-aware transport services.

Operational Challenges for Deployment of Ethernet Services

A carrier's transport network infrastructure is point-to-point and connection-oriented in nature. Traditionally, this transport network has served as a common resource to all services delivered by a service provider (voice, private lines, layer 2 virtual private networks, Internet access, etc.). All transport resources in this network are allocated on a per-TDM-connection basis. As such, traffic switching and/or aggregation is not applied to the packet traffic that may be carried in those point-to-point connections. Packet-level aggregation and switching are left up to the native services network. This simple client-server relationship between the transport network and the services networks, as illustrated in Fig. 1, has proved over the years to be easy to manage and hence operationally attractive for a large-scale carrier network. Any attempt to fundamentally change this operational paradigm is likely to incur great operational expenses that will overwhelm any potential savings in capital expenditure.

Nevertheless, the operational simplicity of operating a transport network that is completely unaware of the actual traffic carried on top of it has one main deficiency: it does not allow differentiation between lucrative high-end services and best-effort low-end services. In other words, all data is treated equally on today's transport network.

The inability to differentially allocate transport resources to low-end data services forces service providers to offer these services at lower margins, and in some cases even discourages a carrier from offering them at all. Indeed, although the volume of data traffic has been steadily growing in the past years, the profitability of data services offerings has not grown accordingly.

In particular, for new Ethernet-based services these limitations may pose material barriers to widespread deployment by service providers over their traditional transport networks. A scheme that on one hand does not change the existing operational paradigm of the transport network, but on the other hand allows carriers to differentiate between various service types and allocate transport resources as appropriate to carry them, is clearly missing and necessary to enable both competitive price points and continued carrier profitability.

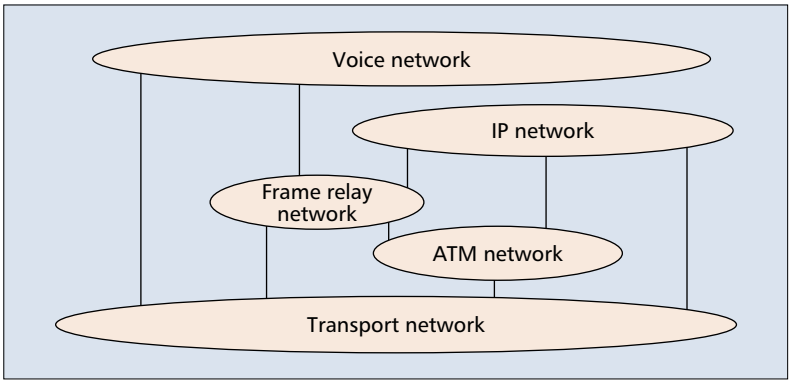

Figure 1. *Client-server relationship between various networking technologies.*

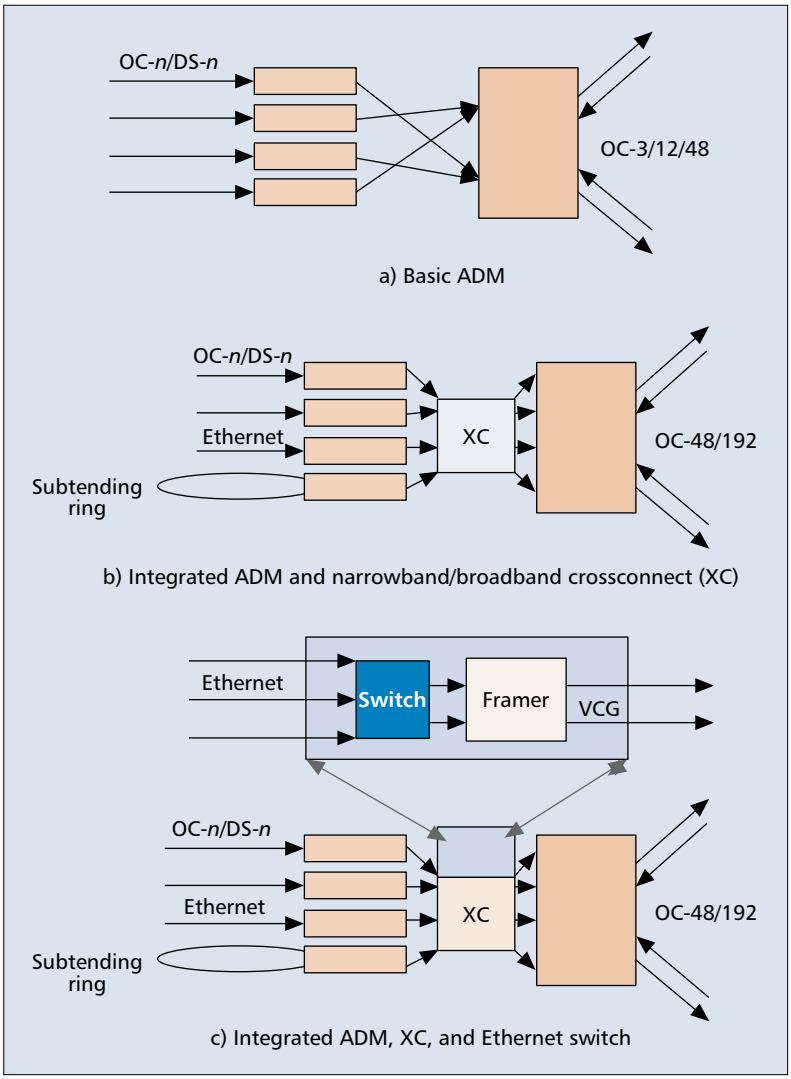

Figure 2. *The evolution of SONET/SDH ADMs.*

The Evolution of SONET/SDH ADMs

During the past five years, SONET/SDH ADMs have evolved to address some of the challenges carriers face in deploying data services, and lately in deploying Ethernet services as well. Figure 2 illustrates the development of SONET/SDH ADMs.

The first generation of SONET/SDH ADMs consisted of a single-stage multiplexer/demultiplexer that aggregated various lower-rate inputs into a high-speed OC-*N* signal. At an add/drop site, only those signals that need to be accessed are dropped or inserted, and the remaining traffic continues through the network element with-

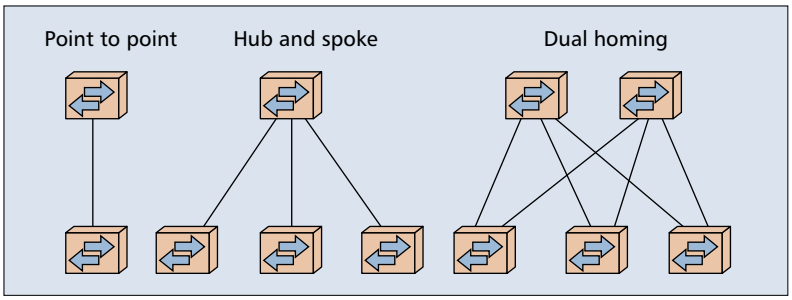

Figure 3. *Ethernet interconnect topologies.*

out requiring special pass-through units or processing.

The second generation of SONET/SDH network elements introduced crossconnect functionality, previously located in a distant centralized hub. It enabled carriers to perform STS-1/STS-3 (or VC-3/VC-4) switching at the edge of the network, and obviated the need for expensive backhauling of local traffic for grooming at a distant hub location. In addition, Ethernet interfaces were added to enable simple and transparent mapping of Ethernet traffic to a SONET/SDH circuit.

The most recent generation of SONET/SDH ADMs introduced native Ethernet switching and processing functionalities, such as VLAN tagging, user priority classification, local multiplexing, and bridging, on an Ethernet tributary. In addition, virtual concatenation was introduced to improve the granularity at which SONET/SDH circuits are provisioned to carry asynchronous payloads.

These enhancements to the original SONET/SDH hierarchy and traffic multiplexing options are extremely useful and advantageous when a transport network is still predominantly used to carry fixed-rate synchronous voice traffic, but has to be able to also carry some asynchronous data traffic. However, as the relative portion of data traffic steadily grows and surpasses that of synchronous voice traffic, improving the circuit-level granularity of a synchronous fixed-size traffic channel is no longer enough.

New Building Blocks to Enable Data-Aware Transport

Next-Generation SONET Technologies and Their Role

Transport networks have been using SONET/SDH technology for several years, and now the installed base of SONET/SDH is high enough to be able to say that SONET/SDH is almost everywhere in service providers' networks. Up to now SONET/SDH networks have been used to transport only TDM-oriented traffic such as voice and private lines. New technologies have been developed to enhance SONET/SDH with efficient transport of packet-based traffics like Ethernet. The three major technologies are the generic framing procedure (GFP), virtual concatenation (VCAT), and the link capacity adjustment scheme (LCAS).

Generic Framing Procedure — Specified in International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Recommendation G.7041/T1.105 [1], GFP is the most popular encapsulation protocol today to adapt a variety of character and packet-oriented payloads into SONET/SDH containers. The main characteristics of this protocol are:

- Low and deterministic overhead
- Complete transparency
- Robust delineation mechanism
- Extensibility to support client-specific management functions

Wide industry support has also facilitated extensive availability of components from a multitude of components manufacturers, simplifying interoperability concerns.

Virtual Concatenation — Virtual concatenation can be viewed as an inverse multiplexing scheme in which X independent STS-N/VC-N signals form a virtually concatenated group (VCG), denoted VC-N-Xv, yielding a transport capacity X times the capacity of a single STS-N/VC-N member. A crucial property of VCAT is that it does not place any new requirements on the existing SONET/SDH network since the STS-N/VC-N containers that make up the STS-N/VC-N-Xv group travel independently from the source over the SONET/SDH network to their common destination. Hence it is sufficient to support VCAT only in both termination points. This allows a very smooth upgrade path, since one only needs to deploy a set of new plug-in units in the existing SONET/SDH systems at both endpoints to establish a virtually concatenated connection. The rest of the network can remain entirely unaware of this fact. The individual STSs/VCs that form a VCG can travel by completely different routes between the VCG termination points. The VCAT mechanism is now fully standardized within ITU-T Recommendation G.707/T1.105 [2].

Link Capacity Adjustment Scheme — VCAT has been expanded with a protocol called LCAS in ITU-T G.7042/T1.105 [3] that enhances the VCAT scheme with hitless in-service addition and removal of STSs/VCs to/from the VCG. Additionally, the LCAS protocol provides load sharing protection by dynamically removing failed members from the VCG if and when they experience faults. Of course, this temporarily reduces the bandwidth of the end-to-end service, but the applications that use the channel are usually capable of adapting to such varying bandwidths. This important LCAS functionality allows a provider to significantly improve the resiliency offered to end users by provisioning diversely routed SONET/SDH paths that belong to the same VCG.

The combination of VCAT and LCAS is a very powerful addition to the SONET/SDH standard as it solves the bandwidth granularity problem. Their definition allows gradual introduction in existing SONET/SDH networks and thus leverages the investments already made in these networks. In fact, they increase the value of existing SONET/SDH networks by allowing bandwidth-efficient introduction of new services.

ETHERNET OVER SONET/SDH

Ethernet over SONET/SDH (EoS) is used to refer to a set of solutions that incorporate transport and switching of end users' Ethernet frames over native SONET/SDH networks. EoS uses IEEE 802.1/802.3 [4] compliant Ethernet processing and switching capabilities, with new SONET/SDH-specific enhancements in terms of capacity adjustment (via GFP, VCAT, and LCAS) and reach of the links, as well as some enhancements to Ethernet bridging to provide better protection and restoration times (e.g., Rapid Spanning Tree Protocol, RSTP) or increased security and data privacy (e.g., VLANs or layer 2 filters). These switches are interconnected with STS-*N*/VC-*N* SONET/SDH paths to carry Ethernet traffic, or use dense wavelength-division multiplexing (DWDM) systems with EoS transponders. Capacities of these links may be allocated on standard IEEE 802.3 physical layer rates (e.g., 100 Mb/s, 1 Gb/s, or even 10G b/s) or allocated on any combination of STS-*N*/VC-*N* grouped into STS-*N*-*Xv* or VC-*N*-*Xv* VCGs.

Like their counterpart native Ethernet switches, EoS switches may be interconnected using a variety of different interconnect topologies depending on the requirements of the services in terms of protections and traffic patterns. When high availability is required a dual homing architecture is the best option, while for looser protection requirements point-to-point or hub-and-spoke architectures may fit well. In these topologies the links can also be protected, although both active and protection links go between the same two switches. The three topologies are outlined in Fig. 3.

RESILIENT PACKET RING

The IEEE 802 LAN/MAN standards committee is finalizing Project 802.17, the Resilient Packet Ring (RPR) standard [5]. RPR is a ring-oriented media access control (MAC) protocol. RPR effectively transforms a chain of point-to-point SONET/SDH paths between nodes to a single virtual shared medium. The shared transport ring created by RPR can then be used over multiple SONET/SDH nodes to carry connection-oriented transport services, and enable optimal and fair use of bandwidth for bursty services through highly efficient statistical multiplexing, overbooking, and spatial reuse transport mechanisms.

The 802.17 MAC was particularly designed to serve as the traffic management layer for SONET/SDH networks. It supports the full range of SONET/SDH rates and GFP as the adaptation layer to the synchronous network. As such, it is the only standards-based scheme optimized from the ground up to perform packet-based traffic management on a SONET/SDH ring.

The RPR MAC introduces the concept of a transit path. At each node on an RPR ring, traffic that is not destined for the node simply passes through, avoiding the queuing and scheduling on a hop-by-hop basis. The MAC on each node performs three functions: *add* for insertion of local traffic from the node, *drop* for removal of traffic destined to a local port on the node, and *pass* for direct transfer of transit traffic from one network link to another. The transit path effectively becomes a part of the transmission medium and makes the RPR ring behave as one continuous shared medium.

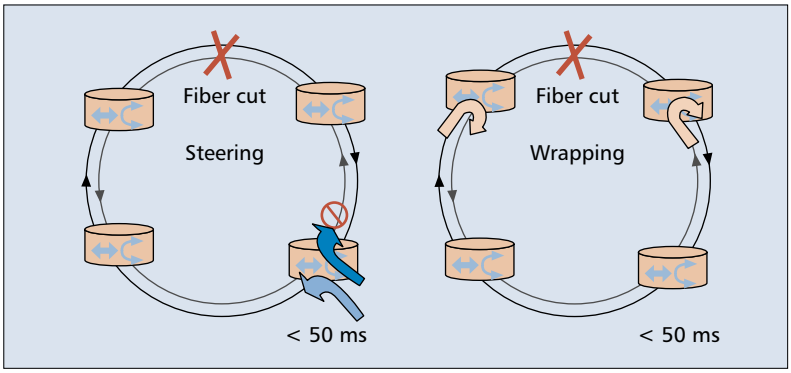

Figure 4. *RPR protection techniques.*

The effective shared medium created by the RPR MAC enables a carrier to perform efficient statistical multiplexing on a per-class-of-service (CoS) basis, on traffic coming from multiple sources on the ring. Instead of mapping unutilized connections to fixed-size pipes as in legacy SONET/SDH, RPR enables dedicated bandwidth for high-priority traffic, and statistical multiplexing and overbooking of low-priority traffic. Alternatively, a service that is defined with a committed rate and an excess best effort rate will use dedicated bandwidth according to its committed rate over the shared medium, and share bandwidth in excess of that, according to its excess rate, with other users over the ring.

Another important attribute the RPR MAC delivers is its ability to guarantee sub-50 ms protection and restoration time for each service carried over the ring. Two protection mechanisms are defined per the RPR standard: steer and wrap. Both guarantee sub-50 ms protection and restoration times, can have relative benefits depending on the particular service characteristics, and can be used on a per-service basis. Figure 4 illustrates the operation of both schemes.

Whereas SONET/SDH protects services on a per-channel/timeslot basis (and unprotected traffic is typically preemptable), the RPR MAC protects only the committed rate per each service (e.g., by reserving only an amount of bandwidth equal to the sum of committed information rated for all high CoS connections), allowing a service provider to guarantee no service interruption for high-priority services, the committed rate per the service level agreement (SLA) with the end user, and make efficient use of the extra bandwidth to provision low-priority best effort services. Moreover, even best effort services are not preemptable over the RPR shared medium, although they may experience bandwidth degradation during a failure event.

RPR's efficient statistical multiplexing capability, together with its support for multiple CoSs with strict performance guarantees, make it an ideal traffic management layer for Ethernet-based services. Ethernet services are expected to be characterized by a high degree of burstiness,

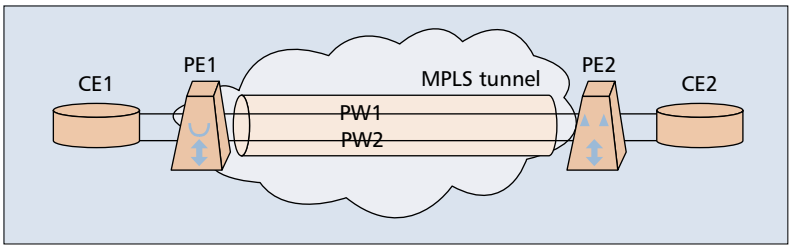

Figure 5. *A PW reference model.*

and an efficient multiplexing scheme that can guarantee end users' SLAs is vital to make them economically viable for service providers.

MPLS Pseudo-Wires as a Service Interface Extension

Pseudo-wires (PWs) are a new Internet Engineering Task Force (IETF)/PW emulation end-to-end (PWE3) defined mechanism that emulates the essential attributes of a point-to-point service over a packet-switched network (PSN) [6]. Figure 5 illustrates the reference network model for point-to-point PWs. Provider edge (PE) devices (PE1 and PE2) allocate one or more PWs on behalf of their client customer edge (CE) devices (CE1 and CE2) to enable the client CEs to communicate over the packet-oriented network. A transport tunnel, typically constructed from MPLS label switched paths (LSPs), is established to provide a direct forwarding path for the PW between the communicating PEs. As such, the PW traffic is invisible to the PSN nodes; consequently, the PSN is invisible to the communicating CEs. Native data units (bits, cells, or packets) presented to the PW end service are encapsulated in a PW protocol data unit (PW-PDU) for transparent transport across the underlying PSN technology. The PEs perform the necessary processing to create the PW-PDUs, as well as handle any other traffic adaptation and resource management functions required by the PW service, such as:

- Encapsulation of service-specific PDUs or circuit data arriving at the PE-bound port (logical or physical)
- Mapping and classification of the encapsulated data to a transport tunnel
- Establishment of the PW including exchange and/or distribution of the PW identifiers used by the tunnel endpoints
- Ingress traffic policing and/or egress traffic shaping
- Managing the signaling, timing, order, or other aspects of the service at the boundaries of the PW
- Service-specific status and alarm management

IETF-PWE3 Schemes for Virtual Wire Services

In the past year the IETF has been defining a unified approach to carry all types of data traffic over a packet-based transport network as point-to-point PWs. PWE3 in the IETF defined a common signaling and encapsulation scheme, commonly known as Martini encapsulation, to carry layer 2 (L2) packet services over MPLS tunnels to effectively provide "virtual wire services." This scheme allows a carrier to extend existing service interfaces from the customer premises to the service access point using a packet-centric transport network in a way that conserves the existing operational paradigm of a point-to-point-oriented transport network. Driven by high end-user demand, most PW implementations to date have been focused on providing the ability to deliver a scalable mechanism for mass deployment of Ethernet virtual wire services (Fig. 6).

Existing MPLS signaling protocols, such as the Label Distribution Protocol (LDP), are used to automatically provision a transport service as a PW end to end. In this approach, a carrier can provision a virtual wire service just by choosing its two endpoints, while the MPLS signaling protocols automatically negotiate the path and resource reservations.

The IETF is in the process of defining PW signaling and encapsulation schemes for the following protocol types: Ethernet, HDLC/Point-to-Point Protocol (PPP), FR, and ATM [7]. A common IETF-PWE3 formal workgroup draft defines the control and signaling for the establishment of a PW [8], while separate IETF-PWE3 workgroup drafts define the encapsulation scheme for each of the above layer 2 protocols.

Ethernet Pseudo-Wires

The most prevalent use of the Martini PW scheme today is for the transport of Ethernet virtual wire services over a PSN such as MPLS. When an Ethernet frame enters a PE network element, it is handed over to the native service processing (NSP) function that performs Ethernet-specific frame processing functions before the frame is forwarded to the PW termination point. Such functions may include stripping, overwriting, or adding VLAN tags, physical port multiplexing and demultiplexing, Ethernet bridging, L2 encapsulation, shaping, and policing. The NSP in turn hands the Ethernet frame to the PW endpoint, where the preamble and its frame check sequence (FCS) are stripped off (end-user FCS retention is optional), a PW control word is prepended to the resulting frame, a proper PW demultiplexer label (VC label) and a proper MPLS tunnel label are prepended, and the packet is transmitted over the MPLS tunnel as a PW. At the egress the same procedures are applied, and regeneration of the FCS is performed at the NSP. Figure 7 illustrates the encapsulation procedures for Ethernet according to the Martini scheme.

The use of the Martini-based scheme for transport and provisioning of Ethernet services promises to be a valuable tool to extend Ethernet services across a variety of packet-oriented networks. Numerous interoperability demonstrations have validated the readiness of the PW concept for the transport of Ethernet services for mass deployment in carriers' networks.

Ethernet Virtual Private LAN Services

The PWE3 virtual private LAN service model (VPLS) is intended to emulate a multipoint switched Ethernet service. Each customer site connects to the service provider network using an attachment circuit to a virtual forwarding/switching function (VFF) within a PE device. Each PE may contain many VFFs. All VFFs within the provider network that have sites belonging to the same VPLS instance are required to have full mesh connectivity among them. The VC label, used as local connection identifiers, identifies an association between two VFFs. Under the PWE3 MPLS architecture framework, the tuple (PE1, PE2, VC label, VC type) uniquely identifies each virtual connection, where VC type indicates the type of connection emulated, in this example VPLS.

A PE may contain multiple VFFs, and two PEs may have a number of VCs that connect VFFs within the PEs. Such VCs are carried within tunnels that are established between the two PEs. Note that there can be multiple tunnels between any two PEs. The use of tunnels achieves at least two important objectives: first, it alleviates any scalability problems associated with realizing a large number of VCs within the network; second, it facilitates separation of control plane procedures used to establish tunnels and VCs. The second objective is important due to the fact that while tunnels are generic in nature, VCs as defined in this context are solely for the purpose of realizing a specific service, in this case VPLS. By separation of control plane procedures for tunnels and VCs, the provider can reuse existing tunneling mechanisms while building new service-specific VC establishment procedures. Such procedures will be implemented on all PE devices and be transparent to the other nodes within the provider network. Details of the VPLS model are provided in the next section.

Putting the Pieces Together: The Multiservice ADM

So how do all these building blocks come together to provide a data-aware and cost-effective transport architecture? A multiservice ADM that combines the benefits of a packet ADM for data traffic with support for traditional circuit-based services is clearly the natural next step in the development of SONET/SDH ADMs.

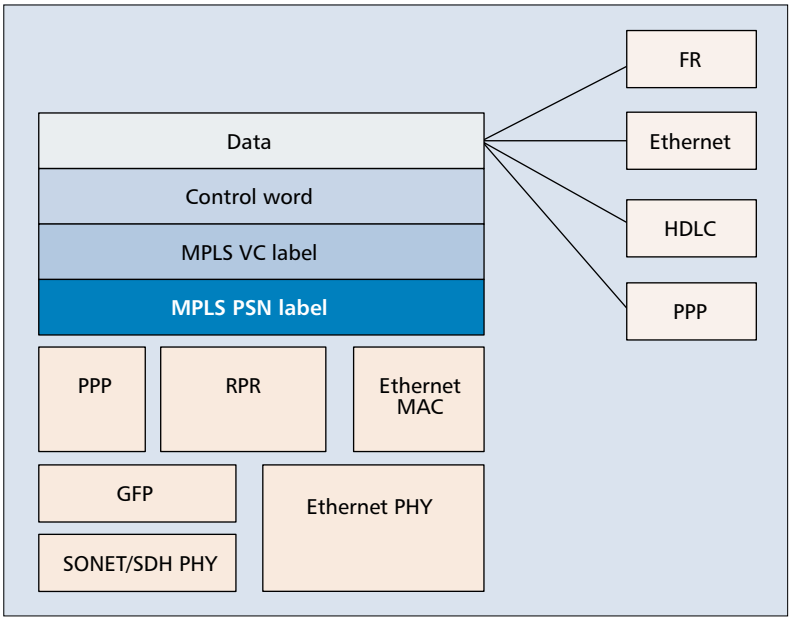

Figure 6. *Any layer 2 over MPLS model.*

Figure 8a illustrates the functional building blocks of a multiservice ADM; like a traditional SONET/SDH ADM, it has low-speed and high-speed interfaces groomed by a centralized TDM matrix, but like a packet ADM it recognizes that services on both low-speed and high-speed interfaces can benefit from statistical multiplexing, be assigned with different CoSs, and be switched on a per-packet basis. In some sense, the introduction of a packet ADM fabric to future SONET/SDH ADMs is similar to the introduction of a distributed crossconnect fabric to the second generation of SONET/SDH ADMs several years ago.

Layer Architecture

A multiservice ADM fuses SONET/SDH technologies such as GFP, VCAT, and LCAS together with packet technologies such as RPR, MPLS, and Ethernet to provide the most cost-effective data-aware transport solution in a way that seamlessly interworks with the existing SONET/SDH-centric paradigm. The packetized transport architecture is designed to carry all data services, including Ethernet services, as PWs over a packet-centric traffic management scheme. Figure 8b illustrates the layer architecture for the proposed packetized transport network

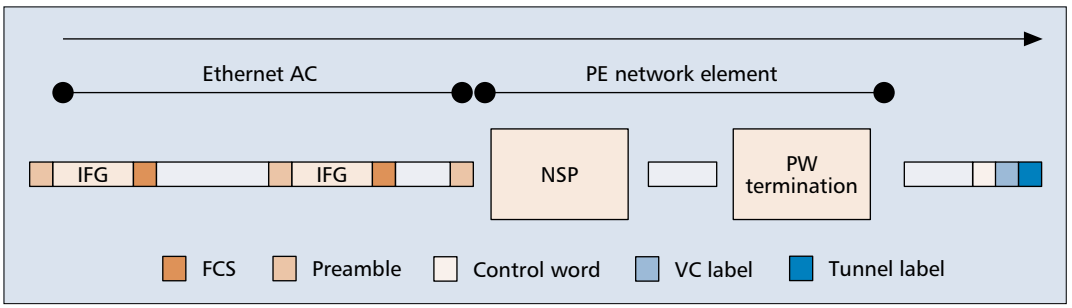

Figure 7. *The Martini encapsulation procedure for Ethernet frames.*

Either the IEEE-defined RPR or Ethernet MAC can be used to enhance the SONET/SDH-based multiplexing scheme for data services. Either approach can serve as the common per-CoS traffic management layer for various data services.

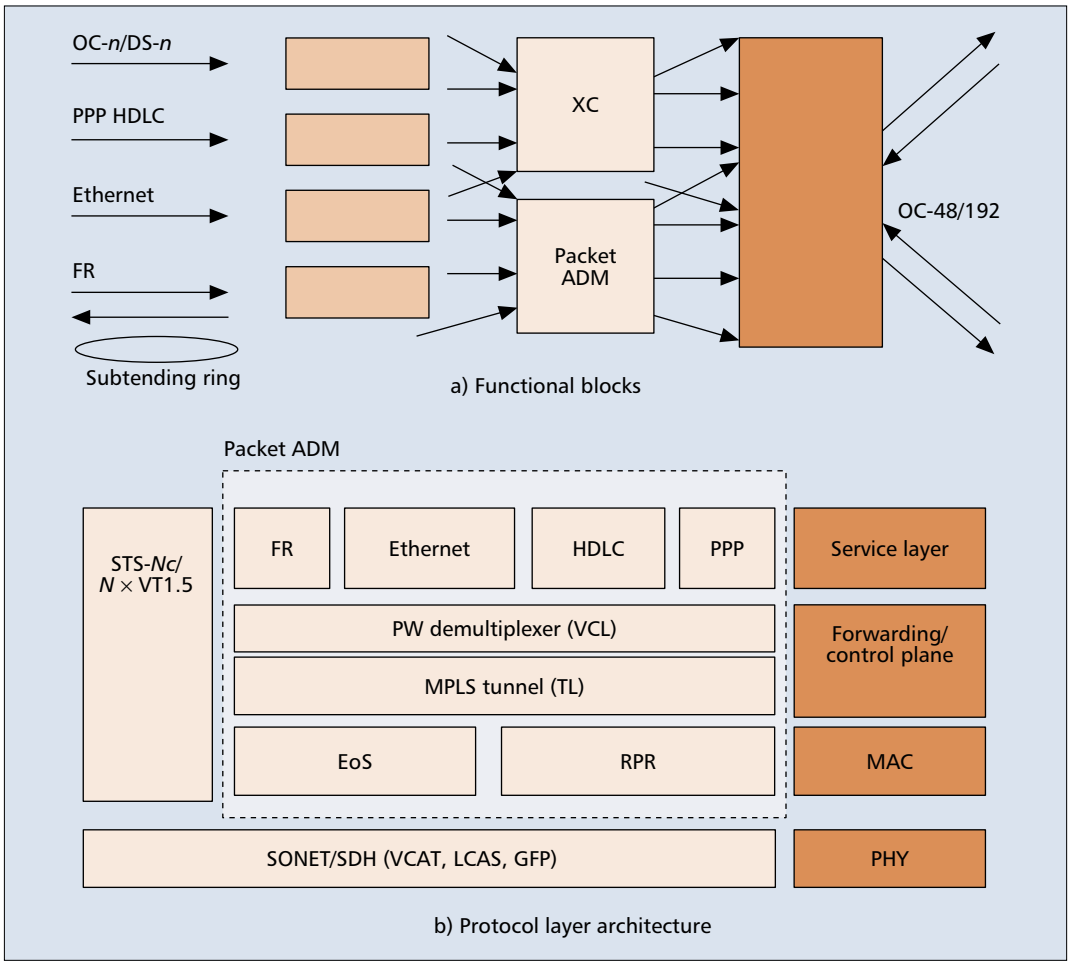

Figure 8. *A multiservice ADM: a) functional blocks; b) protocol layer architecture.*

Traffic Management

Traffic management for TDM circuit-based services is still performed using the traditional hierarchical multiplexing scheme, together with a scalable grooming capability and new technologies to enable finer granularity for these services.

Either the IEEE-defined RPR or Ethernet MAC can be used to enhance the SONET/SDH-based multiplexing scheme for data services. Either approach can serve as the common per-CoS traffic management layer for various data services. RPR is used to create a shared ring for multiple nodes, users, and services that enables optimal and fair use of bandwidth for bursty services by employing highly efficient statistical multiplexing, overbooking, and spatial reuse. EoS may be used to create an equivalent mesh-based logical topology over an existing ring or mesh SONET/SDH transport infrastructure. Standards-based MACs can be used to deliver data-aware transport services, provide service guarantees for traffic that requires dynamic allocation of bandwidth, and maintain strict SLAs for jitter-sensitive services, as well as for guaranteed services with committed and peak information rates.

Path Management and Control Plane

The IETF-defined PW scheme serves as the common path management and control plane for all data services carried over the data-aware transport network. Path management and control for TDM circuit-based services is still performed using the traditional SONET/SDH scheme. Circuits are provisioned between path termination equipment.

Operations, administration, and maintenance (OAM) is traditionally done at all levels and layers of the network. A multiservice ADM naturally provides the traditional SONET/SDH-based OAM toolset, but it also provides a completely new toolset for each of its building blocks. The RPR MAC consists of OAM tools to detect faults on the RPR MAC, and MPLS/PWE3 introduces its own rich OAM mechanisms to detect control and forwarding plane faults.

Service Protection — When considering service protection, both link and node failures have to be considered. Link failure scenarios can be addressed by having both the edge and core network elements support link and facility protection, and rely on the protection features supported by the underlying transport mechanism. For IEEE 802.3-based interfaces, link protection can be achieved via IEEE 802.3ad-based link aggregation techniques.

Node failure scenarios can be addressed by having the edge network connected to multiple core NEs. This configuration can be implemented by either local MPLS fast reroute mechanisms or end-to-end MPLS LSP protection (per

[9]). As part of the configuration/negotiation of labels between label edge router (LER) and label switch router (LSR) devices to identify customer logical ports, backup LSPs can also be configured/negotiated. The LSR and LER devices must then support such backup mechanisms. Note that backup LSPs may be preconfigured or dynamically negotiated upon detection of failure on a currently active link.

Failure Detection, Verification, and Maintenance — Multiservice ADMs must implement mechanisms that enable failure detection and localization in an LSP path, including PW LSPs and tunnel LSPs. LSP resources are managed via the management plane or control plane protocols such as LDP and the 8Resource Reservation Protocol with Traffic Engineering (RSVP-TE) [10]. RSVP-TE messages that set up, tear down, and maintain MPLS LSPs and RSVP adjacencies, and LDP messages that control targeted LDP sessions used for setting PWs are generated and processed by corresponding processes running on a central route processing engine on a core switch/router. Failure of a routing engine or of the RSVP or LDP processes running on a routing engine should not result in declaring that tunnels or PWs are down, as long as the data plane is still alive. Control plane graceful restart capabilities [11] are also planned to reduce the impact of such control plane failure events.

MPLS OAM loopback capabilities [12] and/or MPLS-ping [13] can also be implemented to enable the diagnostics of an LSP path and determine points of failure when they exist.

CONCLUSION

A multiservice ADM offers a robust solution to deploy a fully data-aware transport solution to next-generation transport networks. The multiservice ADM concept is based on standards-based technologies that enable a service provider to distinguish between lucrative high-end services and best effort low-end services. Multiservice ADM solutions enable TDM/Ethernet/MPLS transport solutions over a public SONET/SDH transport infrastructure. The multiservice ADM approach enables network operators to incorporate basic connectivity services such as Ethernet/HDLC/IP/FR private/virtual lines, transparent LANs, and Internet traffic backhaul to customize carriers' evolving needs for profitable and economically viable data services.

REFERENCES

[1] ITU-T Rec. G.7041/Y.1303, "Generic Framing Procedure (GFP)," 2001.
[2] ITU-T Rec. G.707, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)," 1996.
[3] ITU-T Rec. G.7042/Y.1305, "Link Capacity Adjustment Scheme (LCAS)," 2001.
[5] IEEE P802.17 "Resilient Packet Ring," work in progress.
[4] IEEE Std. 802.1Q, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks," 1998.
[6] Bryant *et al.*, "PWE3 Architecture," draft-ietf-pwe3-arch-06.txt, work in progress, Oct. 2003.
[7] See archive from the IETF PWE3 working group.
[8] J. Martini *et al.*, "Pseudo-Wire Setup and Maintenance Using LDP," draft-ietf-pwe3-control-protocol-03.txt, work in progress, June 2003.
[9] ITU-T Rec. Y.1720, "Protection Switching for MPLS Networks," Aug. 2003.
[10] "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209, Dec. 2001.
[11] "LDP DoD Graceful Restart" RFC 2526," draft-ietf-mpls-ldp-dod-restart-00.txt, work in progress, Feb. 2003.
[12] ITU-T Rec. Y.1711, "OAM Mechanism for MPLS Networks."
[13] "Detecting MPLS Data Plane Failures" draft-ietf-mpls-lsp-ping-03.txt, work in progress, June 2003.

BIOGRAPHIES

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. He has worked in the research and development of high-speed data communications protocols and systems for over 15 years. He is a member of the Association for Computing Machinery and Sigma Xi.

GADY ROSENFELD [M] (gadyR@corrigent.com) has more than 12 years of experience in marketing, business development, and strategic planning. Currently he is responsible for defining and leading product marketing and planning strategies for Corrigent Systems. Prior to joining Corrigent, he served as a marketing and business development manager for Orckit Communications, where he led planning and market research initiatives. Prior to Orckit, he served in the Israeli Ministry of Defense as the head of a section responsible for strategic planning, operations research and systems analysis. He holds a B.Sc. in physics and mathematics from Hebrew University, Israel, as part of the prestigious "Talpiot" military program, and an M.Sc. in theoretical physics from Tel-Aviv University, Israel.

The multiservice ADM approach enables network operators to incorporate basic connectivity services to customize carriers' evolving needs for profitable and economically viable data services.