LIGHTWAVE®

Technology

Advances in research, development, engineering, and standards

TECHNOLOGY

Packet optical networking aids metro network evolution

By SAM LISLE

Focus on

Network providers have relied on optical networks as scalable, low cost, general purpose aggregation and transport infrastructures. Optical elements provide significant economic value and deliver operational efficiency by embodying "low touch" operational capabilities including simple, in-service software upgrades, software backward compatibility, robust management interfaces, and network element fault tolerance. These operational capabilities have enabled the massive geographical scaling necessary to connect the tens of thousands of access points in a metro area with the handful of sophisticated service elements located in the metro core.

However, the emerging IP services environment places significant stresses on metro networks. The existing optical aggregation and transport infrastructure must evolve to address these stresses. In particular, three critical challenges arise:

- **Providing significantly more bandwidth at less cost.** Fueled by multimedia applications, network providers are delivering unprecedented amounts of bandwidth for residential, enterprise, and mobile services. By many estimates, bandwidth is projected to grow at more than 50% per year for the next several years.
- **Delivering a scalable Ethernet infrastructure.** Ethernet has considerable attractiveness as an inexpensive and flexible interface, and Ethernet-based services are growing rapidly. As Ethernet has evolved from a LAN to a WAN technology, service providers must build an infrastructure that not only transports Ethernet but also efficiently aggregates Ethernet traffic and hands it off to the IP service edge.
- **Managing the simultaneous growth of both TDM and Ethernet.** While Ethernet-based access to Ethernet and IP services continues to grow rapidly, TDM traffic—especially TDM circuits used as access for IP/MPLS VPN services—also continues to grow. Ethernet and TDM traffic both require large amounts of aggregation but have traditionally relied upon diverse technologies. Network providers can ill afford to deploy overlay aggregation networks.

Fortunately, new photonic, electronic, and software technologies have emerged that form the basis of the emerging class of optical networking equipment known as packet optical networking platforms (packet ONPs). This class of equipment is projected by analysts Heavy Reading and Infonetics to soon have an annual market size greater than $1 billion. Packet ONPs uniquely address the three key challenges IP services present while retaining the low-touch operational capabilities of traditional optical networking equipment. As a result, service providers can evolve their metro infrastructure to profitably deliver the new generation of IP services.

Distributed Ethernet aggregation using connection-oriented Ethernet transport tunnels

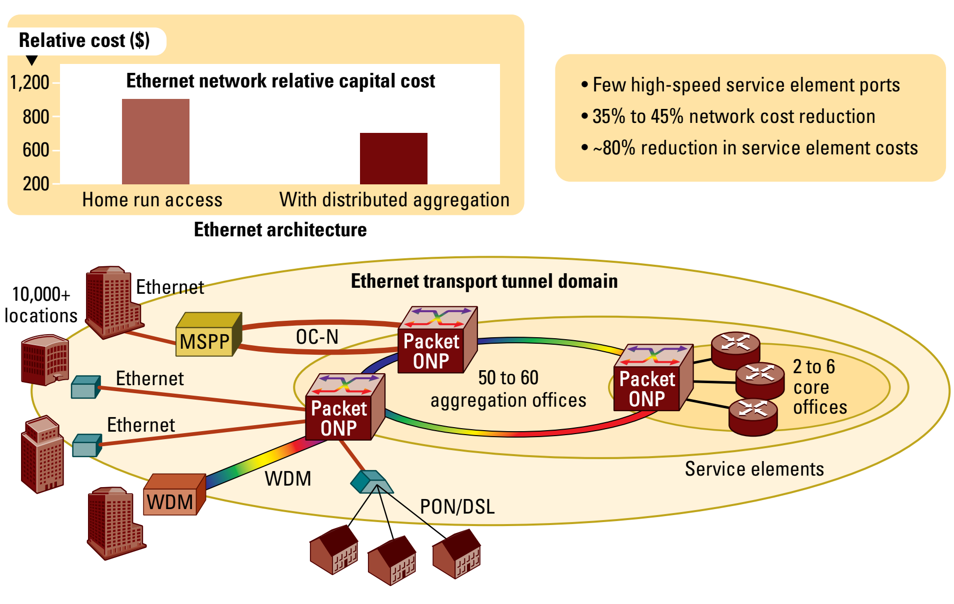

Figure 1. Connection-oriented Ethernet transport enables a distributed architecture that can reduce costs and provide service support.

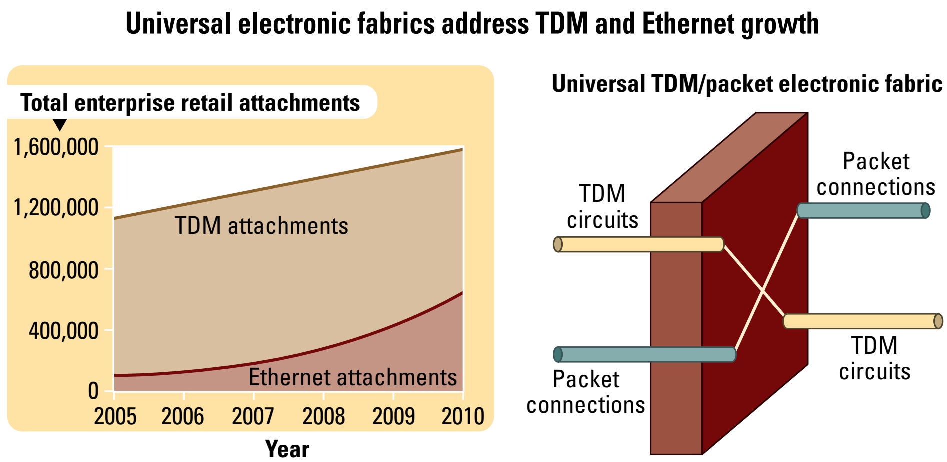

Figure 2. Universal electronic fabrics address TDM and Ethernet growth.

Enabling photonic technologies

Photonic technologies enable metro networks to scale in bandwidth and geography. A ROADM-based flexible photonic layer eliminates 90% of the capital cost of passing traffic through intermediate offices and enables service providers to deploy hundreds of gigabits per second of metro bandwidth in support of high-bandwidth services. ROADMs enable the service provider to deploy new wavelengths via remote provisioning rather than by dispatching a technician to each site along the network to manually patch the traffic. Innovations such as dynamic transient response amplifiers were a critical building block in establishing the photonic layer as the transport convergence layer for all services.

Standalone ROADMs currently must be stacked alongside standalone SONET/SDH and Ethernet elements to provide subwavelength grooming and aggregation and make efficient use of each wavelength. As bandwidth continues to grow, it becomes increasingly important to tightly integrate this ROADM-delivered photonic layer with TDM and Ethernet subwavelength grooming.

One critical enabling technology is the "pluggable ROADM." Ongoing density enhancements are now enabling the integration of amplifier units and wavelength-selective switch (WSS) based optical fabrics into smaller form factors. As these form factors improve, the ROADM hardware becomes a pluggable option in networking platforms that perform subwavelength grooming. This tight integration removes substantial capital cost and operations complexity. It is no longer necessary for network providers to deploy separate transponder hardware or other interfaces between standalone ROADM elements and subtended elements.

Pluggable ROADM hardware improves ROADM simplicity by integrating more functionality into fewer hardware units and greatly reduces the amount of installation fibering required. Further automation enhancements such as integrated power and connectivity tests and in situ span measurements can greatly reduce the operational cost associated with extending the photonic footprint.

Enabling electronic technologies

The challenge the simultaneous growth of TDM and Ethernet presents is unique in the optical era. Ethernet has attraction as a customer interface because it is elastic and enables the end user to increase bandwidth or add more connections without requiring the network provider to visit the site. Historically, switched Ethernet services networks have been deployed with little or no aggregation—with customer ports being directly tied to service element ports. Sometimes access is provided over a Layer 1 infrastructure that efficiently transports each customer port to the service element. This approach alleviates the significant fiber costs of a pure "home-run" architecture, but because Layer 1 infrastructures lack Ethernet aggregation, this results in "port overload" on IP and Ethernet service elements. Widely distributing service elements for aggregation and transport can result in high operational costs because these platforms do not provide the previously mentioned low-touch operational capabilities of optical elements.

TDM continues to find favor because of its high reliability, ubiquitous availability, and a clean carrier-to-carrier handoff. Circuits and revenue for TDM-based access to IP/MPLS VPN services in particular continue to grow.

With both TDM and Ethernet network attachments growing significantly, network providers must either deploy two aggregation networks using existing disparate technologies or look for alternatives that enable a single infrastructure that economically aggregates Ethernet and TDM traffic. Two key technologies address this difficult challenge: connection-oriented Ethernet transport and the universal electronic switching fabric.

Connection-oriented Ethernet transport tunnels. Ethernet as a connectionless technology cannot by itself deliver the connection performance and reliability needed to deliver a full range of services. Ethernet must rely on an underlying connection-oriented transport tunnel (Fig. 1).

Layer 1 transport of Ethernet is widely deployed and meets connection performance requirements but does not provide Ethernet aggregation and present an efficient Ethernet handoff to the IP service edge. Packet service platforms provide Ethernet aggregation but are complex elements with high-touch operational profiles that do not lend themselves to highly distributed aggregation infrastructure deployment and are not optimized for tunnel transport.

Several Layer 2 transport tunnel technologies exist or are under development, including routed MPLS with Martini encapsulation, Provider Backbone Bridging with Traffic Engineering (PBB-TE), and Transport MPLS (T-MPLS.). Nonrouted circuit-based approaches such as PBB-TE and T-MPLS minimize the software complexity of aggregation elements and fit within the current circuit-based operational environment of the aggregation network.

Connection-oriented Ethernet transport technologies can realize a distributed aggregation and transport network for metro Ethernet that reduces network costs while achieving the necessary quality of service and circuit protection that is required for a general purpose scalable infrastructure.

Managing TDM and Ethernet growth with universal electronic switching fabrics. Universal electronic switching fabrics that switch both circuit and packet traffic in their native formats are vital to delivering significant TDM and Ethernet circuit growth on a single converged network. Universal switching fabric technology enables a level of network element integration and convergence that has not been possible in previous generations of equipment (Fig. 2).

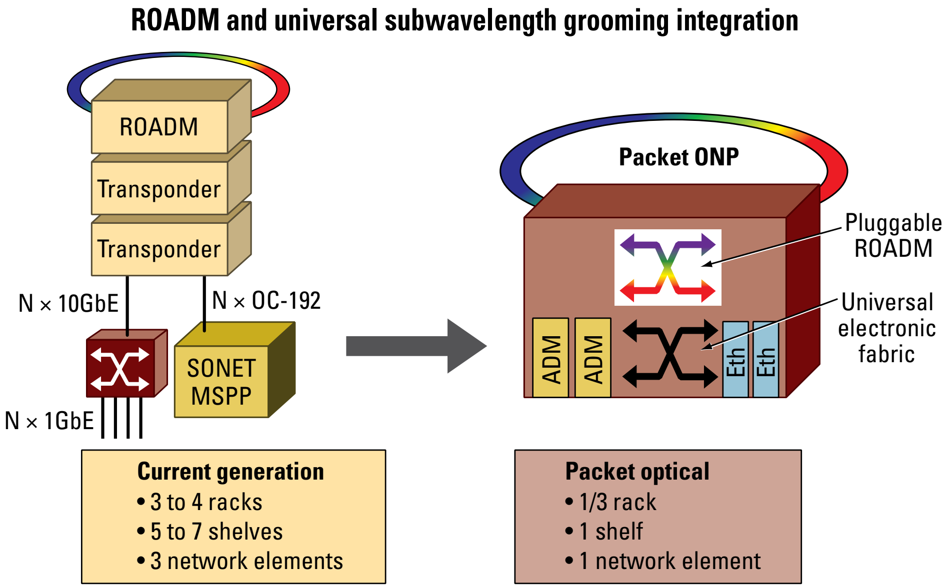

Figure 3. Packet ONPs integrate ROADM networking with universal subwavelength grooming.

ROADM and SONET elements alike have included interface cards to adapt Ethernet to the SONET or ROADM environment. This has been termed card-level integration. Card-level integration is an excellent fit for spoke site applications, but because there is no centralized packet fabric there can be no Ethernet traffic aggregation or grooming across the interface cards. Therefore networks based on this equipment still require subtended Ethernet elements with large fabrics to provide the necessary aggregation at hub sites.

By contrast, a single universal electronic switching fabric offers an elegant solution to this problem. With recent advances in ASIC technology it is now possible to develop an integrated fabric unit that 1) identifies input traffic as either circuit or packet, 2) switches the incoming traffic natively, and 3) assigns the switched traffic onto an outgoing set of interfaces. Since this type of fabric is built from highly integrated ASIC technology, the cost can be low. The fabric design results in a modular system that does not penalize the network provider for unwanted functionality or force the network provider into predicting the ratio of Ethernet to TDM traffic. A universal fabric enables a network element to become technology neutral, empowering a single infrastructure to manage both TDM and Ethernet growth.

Tightly integrating the universal electronic fabric with the pluggable ROADM creates a single chassis that not only provides ROADM-based photonic transport but also eliminates the need for subtended standalone SONET/SDH and Ethernet aggregation equipment (Fig. 3).

Enabling software technologies

Central to managing and operating a converged element that supports both Layer 1 and Layer 2 technologies is the software concept of operational virtualization. Operational virtualization enables network providers with organizational and operational support system (OSS) boundaries between Layer 1 and Layer 2 management to simultaneously and cooperatively access and manage an integrated aggregation and transport infrastructure. There are three keys to such an approach: network management protocol support, granular object model, and a role-based security scheme.

A role-based security scheme grants management users authority based on the technology for which they are responsible, making it possible to support a range of network management applications that account for diversity in network provider operations environments. For example, users in a Layer 2 provisioning center can be granted provisioning privileges for Ethernet transport while having only read-only alarm-monitoring privileges for the underlying Layer 1 entities.

Packet ONPs

Packet ONPs are the newest class of optical networking equipment. These products are a logical evolution of both ROADM elements and SONET/SDH MSPP equipment and employ the technologies described here to meet the emerging challenges of delivering cost-effective bandwidth, building a scalable Ethernet infrastructure, and delivering universal aggregation for both TDM and Ethernet.

Universal electronic switching fabrics enable packet ONPs to groom TDM circuits and Ethernet connections with equal facility, thereby enabling network providers to economically deploy a "neutral" transport infrastructure that can be grown to address TDM or Ethernet needs.

Connection-oriented Ethernet transport enables packet ONPs to provide distributed Ethernet aggregation that greatly decreases the cost of Ethernet services networks. In particular, the nonrouted circuit-based approaches (PBB-TE and T-MPLS) match the existing low-cost operational approach used in the metro network. Connection-oriented tunnels can enjoy dedicated 50-msec protection and can provide jitter, latency, and packet loss performance comparable to TDM private lines.

To economically increase Ethernet and TDM bandwidth, packet ONPs embrace pluggable ROADM capabilities that map both traffic types onto the flexible photonic network, thereby scaling network links toward terabit capacity while simplifying network operations.

Because packet ONPs tightly integrate the pluggable ROADM and universal electronic fabrics, they make highly efficient use of the metro wavelengths and considerably reduce capital and operational expenses by collapsing ROADM, SONET/SDH, and Ethernet transport functions into a single chassis.

Most importantly, packet ONPs enable service providers to leverage and evolve the enormous existing optical network base while using familiar operational procedures. A core network composed of packet ONP elements is served by a mix of standalone SONET/SDH, WDM, and native Ethernet elements in the access network. With emerging packet ONPs, network providers will create cost-effective, low-touch universal aggregation and transport infrastructures that meet the needs of all residential, enterprise, mobile, and wholesale services. LW

Sam Lisle *is market development director at Fujitsu Network Communications (http://us.fujitsu.com/telecom).*