

LIGHTWAVE®

■ TECHNOLOGY

'EtherOptics' optimizes intelligent IP/Ethernet transport

By RAYMOND XIE

Few people disagree that the convergence of Ethernet and optics (EtherOptics) is one of the cornerstones of next-generation broadband networks. This convergence offers fixed-line and mobile network operators a powerful range of benefits, including optimization of packet transport and cost-efficient deployment of new revenue-generating services.

Although the industry has reached consensus that Ethernet will play a fundamental role in the evolution from circuit- to packet-based services, the real-world implementation is not always as simple as it seems. Ethernet will continue to be a key technology in access networks, but the existing core/regional service provider network infrastructure is still largely TDM centric. The benefits of Ethernet simplicity and low cost are unquestionably compelling; however, service providers managing end-to-end regional and core service networks are faced with other economic and practical considerations that weigh heavily as they cost-effectively balance TDM-to-Ethernet convergence and support for traditional revenue-generating SONET/SDH services. Key implementation considerations include carrier-class quality of service (QoS); bandwidth optimization; standards-based interoperability; operations, administration, maintenance, and provisioning (OAM&P) robustness; and guaranteed service-level agreement (SLA) performance as the network scales.

Ethernet owes its phenomenal success in no small part to simplicity because it was designed for LAN connectivity in well-controlled private network environments supporting a finite number of users. The emergence of EtherOptics was born out of the desire to leverage Ethernet cost, simplicity, and efficiencies and to extend these attributes to metro and wide-area (or long-haul) networks. Extending the LAN-oriented feature set of Ethernet to the WAN, the core of which is an optical

infrastructure, poses more than a few interconnection challenges.

Skimming the EtherOptic cat

There is clearly more than one approach to extending Ethernet into a carrier's optical infrastructure. While each approach has its advantages and disadvantages, ultimately the suitability of a specific approach will depend on, among other things, the types of services offered, the service provider's network infrastructure and operations, the deployment strategy (overlay versus greenfield), the existing equipment, and the service provider's target customers. The following discussion highlights some of the more common approaches that carriers can consider today for extending Ethernet services to the MAN/WAN environment.

Native Ethernet (Ethernet over Ethernet). In this approach, Ethernet interfaces on routers or switches are simply interconnected over dark fiber networks. Since no additional protocol or format conversion is required, this is the most straightforward method for extending Ethernet to the MAN.

The primary drawbacks of this method are the same as the limitations of Ethernet for MAN/WAN applications outlined previously: carrier-class QoS, OAM&P robustness, network resiliency, interoperability, etc. In addition, this approach is difficult to scale efficiently, since new optical interfaces and trunks need to be deployed as new nodes are added, even if the existing interfaces and trunks are not fully utilized.

Ethernet over WDM. This approach is similar to native Ethernet, except that the use of WDM provides additional reach and capacity in comparison to the standard Ethernet PHY. Only a simple wavelength conversion from the Ethernet PHY wavelength to the appropriate WDM wavelength takes place with any protocol change.

Depending on the number of channels and distance required, either CWDM or DWDM can be deployed. WDM can be used with both Gigabit Ethernet (GbE) and 10GbE. Optical layer protection and performance monitoring can be used to overcome the rudimentary, inherent limitations of Ethernet.




Figure 1. MEF service definitions have defined two broad categories of service: point-to-point E-Line services (shown here) and multipoint E-LAN. Both take advantage of ITU-T and OIF standards.

Ethernet over next-generation SONET/SDH (EoS). In this approach, Ethernet is carried over a next-generation SONET/SDH network using standardized techniques such as Generic Framing Procedure (GFP), Virtual Concatenation (VC), and Link Capacity Adjustment Scheme (LCAS). VC provides a method for rightsizing the transport pipe, and LCAS provides a mechanism for in-service increase and decrease of transport capacity.

Widely deployed and proven, the primary advantages of EoS are its ability to leverage the existing MAN/WAN transport infrastructure and the superior and well-known protection, fault isolation, and OAM&P features of SONET/SDH.

Ethernet over MPLS (EoMPLS). Ethernet over MPLS is similar to Ethernet over SONET/SDH, except that it uses MPLS label switched paths (LSPs) instead of virtually concatenated circuits. EoMPLS is advantageous in converged multiservice networks. EoMPLS overcomes the scalability constraint of the 12-bit VLAN identifier field from native Ethernet, because the MPLS label is a 20-bit field that allows 1 million identifiers. MPLS also supports a fast reroute capability for network restoration after a failure.

However, EoMPLS is incompatible with traditional SONET/SDH networks, and is architecturally more expensive than the EoS approach—interfaces on label switched routers (LSRs) are much more costly than equivalent interfaces on optical switches.

PBB/PBT. Provider Backbone Bridging (PBB)/Provider Backbone Transport (PBT) are extensions to native Ethernet technology and are initiatives designed to provide carrier-class Ethernet transport capabilities (e.g., determinism, QoS, scale, 50-msec resiliency, carrier-class OAM&P tools, etc.). While relatively new to the EtherOptics landscape, PBB/PBT adds to the existing set of options available to network operators. Limited from a deployment perspective, PBT is essentially designed for manually provisioned point-to-point Ethernet (transport) connections. PBB, on the other hand, allows Ethernet bridging and therefore enables automated establishment of Ethernet connectivity, but thus far has not been deployed.

Recently, there has been intensive debate in the industry about the cost savings, OAM&P capabilities, and interoperability issues between PBB/PBT and the incumbent

MPLS transport technology. With carrier reaction mixed to date, and the initiatives tainted to a certain degree from excessive hype, it is premature at this time to fully assess the market impact these new transport options will make. Nevertheless, on the access side, PBB/PBT may eventually evolve as a leading Ethernet access technology, while serving to complement the various core transport network architectures and technologies required for end-to-end service performance.

Embracing EtherOptic standards

Progress in the standards organizations and Ethernet-focused industry forums has played a critical role in advancing EtherOptic applications and services. ITU standards—for example, GFP, VC, and LCAS—have been instrumental in enabling operators to transform existing SONET/SDH networks into a flexible and efficient transport vehicle for the IP traffic driving network growth. Meanwhile, since there is no universally accepted meaning for “Ethernet services,” the Metro Ethernet Forum (MEF) has led a commendable effort to standardize service definitions to clearly communicate the features of various Ethernet services among service providers and end customers.

Based on the MEF service definitions, there are two broad categories of Carrier Ethernet services: point-to-point E-Line services (Fig. 1), and multipoint E-LAN services. These services are accompanied by five attributes that define carrier-class service: standardized services, scalability, reliability, QoS, and service management. More and more Ethernet access and transport devices have become MEF certified in the last couple of years.

On the transport side, the ITU-T Optical Transport Network (OTN) defines the next-generation framing protocol, G.709, to enable efficient transportation of 10GbE services. The 10.7-Gbit/sec throughput provided by G.709 overcomes the 9.9-Gbit/sec SONET/SDH throughput limitation and is designed to provide transparent transport for 10GbE services.

On the Ethernet standard front, a newly formed task force of the Institute of Electrical and Electronics Engineers (IEEE) has begun to review technical proposals for a new standard for 40GbE and 100GbE. The IEEE established 100GbE as its next step in 2006 but later decided to include 40GbE in

the same standard.

The need for 40G/100G Ethernet is growing as IP video and transaction-intensive web applications are exploding across the Internet. Companies such as YouTube and Yahoo regularly add 10-Gbit/sec service pipes to meet growing demand, and 40/100GbE will truly unlock the potential of the transport network and enable service providers to aggregate 10GbE services much more efficiently.

On the control plane side, the Optical Interworking Forum (OIF) is finalizing the User Network Interface (UNI) 2.0 implementation agreement to allow Carrier Ethernet devices to request connections from transport networks in a dial-up fashion using the standard control plane capability. Recently UNI 2.0 functionality was successfully demonstrated during the 2007 OIF Worldwide Interoperability Demo.

EtherOptics and the intelligent optical layer

An intelligent optical control plane is the key to realizing the full potential of Carrier Ethernet services across multiple service provider networks that may comprise a variety of legacy transport equipment, heterogeneous administrative domains, and different levels of survivability mechanisms. The ITU-T ASON framework, embodied by OIF UNI/E-NNI, IETF GMPLS, and vendor-specific I-NNI signaling and routing protocols, provides the unified control plane capability to enable provisioning and operation of Carrier Ethernet services in a true multicarrier, multivendor environment.

Besides multivendor interoperability, the intelligent optical control plane enables new revenue-generating services such as bandwidth-on-demand Ethernet. Traditionally, it takes weeks if not months for a service provider to turn up an end-to-end Ethernet connection, and these service contracts typically require 12-month commitments from end customers. Enabled by end-to-end automatic service provisioning and termination, the optical control plane changes the entire service paradigm by introducing new short-term, high-speed Ethernet services to meet the growing demand of video conferencing and IPTV distribution applications. These services can be offered on a daily or even hourly basis.

With the intelligent optical layer, Carrier




Figure 2. Ethernet packet intelligence extended into the optical domain enables a rich and flexible set of features that make service management more efficient and provisioning more reliable.

Ethernet services become far more dynamic, and richer in available service parameters, than previous static point-to-point Ethernet connections. The 2007 OIF Worldwide Interoperability Demonstration showcased end-to-end provisioning of dynamic switched Ethernet services over multiple, control-plane-enabled intelligent optical core networks through the use of OIF's implementation agreement of UNI 2.0 and E-NNI.

The adoption of ASON-empowered optical switched networks has also contributed to one of the most compelling applications for carrier-class EtherOptics. ASON-empowered optical switched networks incorporate automatic network topology discovery, the ability to dynamically provision end-to-end SONET/SDH circuits across a variety of topologies (ring, mesh, linear), and flexible ways to assign specific

protection/restoration characteristics (e.g., UPSR/SNCP, BLSR/MS-SPRing, dynamic path restoration, dynamic span restoration, etc.) to each constituent circuit.


These two features (flexible topology and tiered protection schemes) enable service providers to choose the transport mechanism best suited to the Ethernet service requirement based on the SLA. For example, a best effort service for Internet access can use a bandwidth-efficient mechanism, such as dynamic path restoration on a mesh topology, while a high-availability, delay guaranteed service can use an LSR protection ring circuit.

The combination of Ethernet transport protocols and an intelligent, dynamic optical layer provides a robust foundation for next-generation Ethernet services while lowering overall transmission costs and simplifying operations. Ethernet packet intelligence, for

example, tightly integrates the service and transport layers of the network, which simplifies EoS provisioning, increases security for network partitioning, and enhances the manageability and resiliency of Ethernet transport. High-density aggregation and intelligent packet-to-circuit mapping routes Ethernet traffic flows over SONET/SDH networks more efficiently, and provides a scalable, cost-effective alternative to packet-over-SONET/SDH (PoS) interfaces to accommodate ever-increasing data traffic.

Mapping Ethernet traffic to virtual concatenation groups (VCGs) based on Layer 2 label information (VLAN tag, MPLS label) extends the traffic engineering features of packet networks to the optical domain (Fig. 2). Combined with support for VCAT, GFP, and LCAS standards, and robust control plane and management software, enhanced provisioning, flow control, and protection mechanisms simplify the management and customization of Ethernet services.

Ethernet packet intelligence supports diverse applications, such as connecting remote offices to a centralized resource or larger corporate facility, disaster recovery, or any situation that requires fully distributed and protected services.

Backed by significant technology advances, carrier-class Ethernet services will continue to proliferate in quantity, range of data rates, and sophistication of defined service attributes. The critical relationship between Ethernet technology and the optical layer will continue to deepen and grow in importance as this market evolves. By combining the positive attributes of Ethernet with intelligent optical layer network elements, network operators will be well positioned to deliver a highly scalable, flexible, and carrier-class transport infrastructure in a cost-optimized manner. 

Raymond Xie, PhD, is senior product manager at Sycamore Networks Inc. (www.sycamorenet.com).

Sycamore Networks, Inc. • 220 Mill Road • Chelmsford, MA 01824-4144, USA
Phone: 978-250-2900 • Fax: 978-256-3434
www.sycamorenet.com

Sycamore Networks, Inc. (NASDAQ: SCMR) is a leading provider of intelligent bandwidth management solutions for fixed line and mobile network operators worldwide. From multiservice access networks to the optical core, Sycamore products enable network operators to lower overall network costs, increase operational efficiencies, and rapidly deploy new revenue-generating services.

Sycamore assumes no responsibility for the accuracy of the information presented, which is subject to change without notice. Sycamore and Sycamore Networks are trademarks or registered trademarks of Sycamore Networks, Inc. in the United States and/or other countries.

Copyright © 2008 Sycamore Networks, Inc. All rights Reserved.

