Can T-MPLS Finally Unhook SONET/SDH Transport?

Scott Clavenna

Ethernet hasn't cracked the transport market yet, but a new version of MPLS is promising to help.

Today there is little disagreement that Ethernet has become the dominant driver of growth in carrier data services, leaving frame relay, ATM, and even TDM private lines in its wake. Enterprises continue adopting Ethernet, and carriers are rolling out a diversity of Ethernet options, from simple point-to-point private lines to more complex multipoint services that leverage MPLS for scale and quality of service (QOS). Ethernet is also moving rapidly into carrier networks as an aggregation and transport infrastructure, supporting lower-cost-per-bit backhaul of consumer broadband services, cellular data services and business data services.

For years, this twofold ascendance of Ethernet has been expected to spell the end of SONET as the nearly ubiquitous transport layer for network operators, but each year SONET hangs on, thanks to its robustness, well-established standards and massive installed base. Recently, however, even SONET's advocates have begun to describe, and standardize, its departure from the stage. What will stand in its place is not necessarily today's "Carrier Ethernet," but what is more generally being called "packet transport."

For enterprise users, packet transport promises a scalable, resilient network infrastructure that pairs the reliability of SONET with the flexibility of Ethernet and IP networking. For network operators, packet transport means an end to the debate between SONET and Ethernet, and an Ethernet-based technology that is "packet friendly" and cost-effective while preserving the deployment, operations and management techniques of SONET. Not surprisingly, the key ingredient in this vision is Multi-Protocol Label Switching (MPLS).

Scott Clavenna is chief analyst at Heavy Reading, specializing in broadband and optical networks market research. He can be reached at Clavenna@heavyreading.com.

The Inevitability Of Ethernet

The interest in migrating to Ethernet-based transport is not new. To be sure, carriers have, for at least seven years, sought ways to exploit Ethernet's low cost, scalable interface and overall packet-friendliness—first in the metro for private line replacement and lower-cost LAN services, then in the access and core networks.

In the late 1990s, Ethernet played a key role in the arrival of the metro networking market, in which enterprise customers of many sizes and verticals could move away from TDM, SONET and ATM to an infrastructure that closely matched their LAN and campus networks. In well-contained metro networks, Ethernet served corporate data networking demands well, particularly for transparent LAN and high-speed Internet access.

Moving Ethernet into the mainstream to carry all services end to end, however, has presented formidable challenges. Major service provider transport engineers and network planners often reveal a strong bias towards SONET as a more reliable, robust infrastructure than Ethernet, yet they typically also favor a longer-term transition to Ethernet because of its superior economics and flexibility in an IP world. The tension between these desires has led to a fragmented market in somewhat of a stalemate over the past five years, looking for a common path forward.

In the past few years, many smaller network operators have found Ethernet to be a key to competitive differentiation, and have overhauled their SONET networks in favor of pure Ethernet infrastructure. The larger incumbent network operators have had a harder time of it—their SONET networks are just too vast in scope, and many of their large enterprise customers are too risk-averse to jump onto a network that is perceived to be "best effort" or only slightly better. But in the past two years, the realities of services transition have taken hold for operators of all sizes.

In Figure 1, the results of a recent survey conducted by Heavy Reading (the market research arm of Light Reading), illustrate the extent to which service providers see their future fortunes in Ethernet and IP-VPNs. Predictions for growth in traditional Layer 2 services and TDM private lines are modest at best, and most operators who do see growth there find it most often to be organic growth generated by existing customers, not new accounts.

What is remarkable about these figures is the

FIGURE 1 Service Providers See Growth In Ethernet and IP VPNs

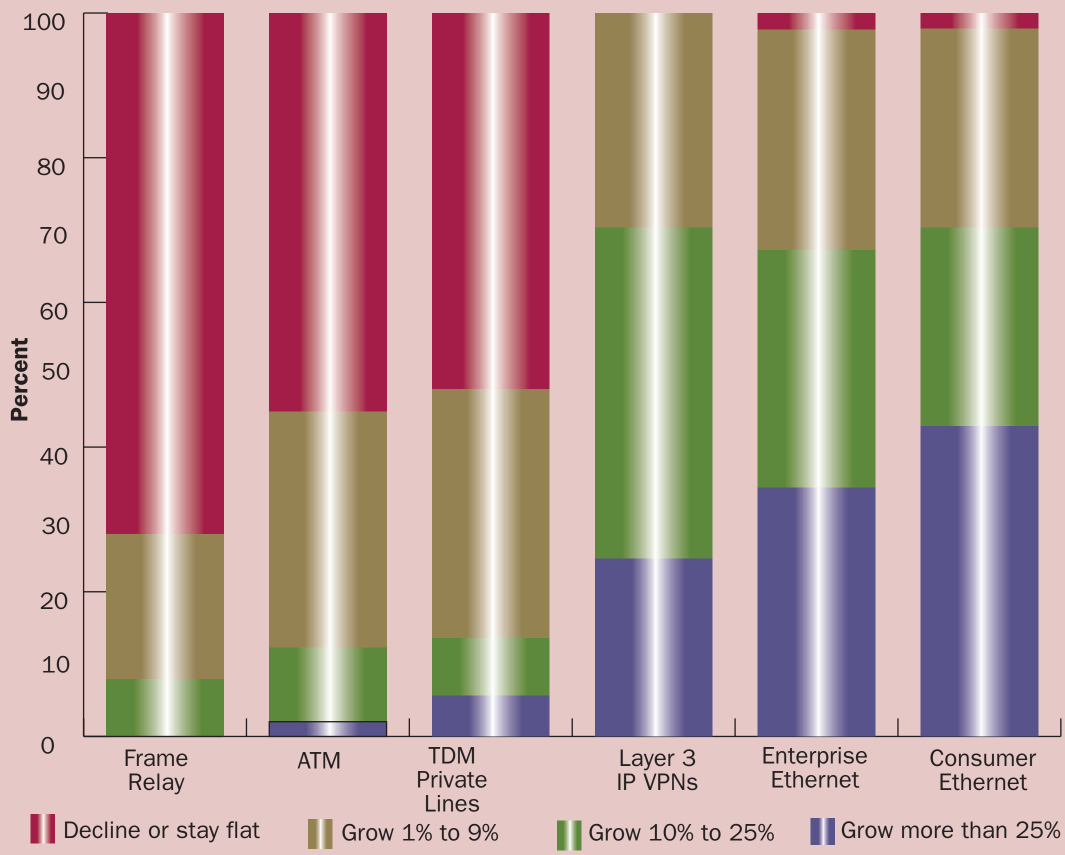

Source: Heavy Reading 2006 Survey of Ethernet Service Providers, 151 respondents

Carriers seem to want the virtues of both SONET and Ethernet

extent to which Ethernet is predicted to grow in new consumer rollouts of enhanced broadband DSL, cable modem and even fiber to the premises. With this in mind, service providers are beginning to see the value, even the necessity, of migrating their transport infrastructure to Ethernet.

Figure 2, p. 48 provides details on carrier service trends that offer the opportunity for packet transport. Business Ethernet services were the first such market opportunity. In practice, the largest concentration is currently being carried on SONET/SDH infrastructure, although a significant number are carried on best-effort Ethernet switches, true carrier Ethernet gear and some WDM (wave division multiplexing) for high-speed connectivity.

In the wireless world, a cost-conscious urgency has developed around the move to Ethernet transport. In current cellular network architectures, E1s and T1s are commonly used in the radio access network to provide backhaul from cell towers. SONET is typically used in the metro transport network, most often at STM-1 or OC-3 rates, carrying groomed TDM circuits and ATM inverse mux (ATM IMA) traffic. This traffic comprises data, ranging from text messages to Web browsing and streaming video.

In the core transport network, the trunk rates are higher, but generally one finds a mix of SONET and ATM. This architecture can become quite costly as these mobile operators adopt 3G technology, and as new classes of data and multimedia services take hold that have much lower revenue per bit than voice services.

This creates demand for more backhaul capacity on fewer individual leased lines in the radio access network. In addition, as most new 3G services are based on IP, there is a requirement to migrate to an end-to-end IP-based network, with simple traffic aggregation and limited QOS differentiation.

Ethernet's role may not extend all the way to cell towers initially, as there are legitimate concerns around frequency stability and reliable circuit emulation, but Ethernet will play an increasing role within the aggregation and metro transport network. Again, this transition to Ethernet-based access remains a compelling driver for packet transport in the metro and core networks and will continue to pressure SONET's dominance in these segments.

Making The Transition To Packet Transport

For a time, there was a feeling in the industry that operators would eventually wean themselves off SONET and move over to Ethernet. Carrier

FIGURE 2 Service Trends Driving The Transition To Ethernet

Service	Network Today	Ethernet Opportunity
Triple play	High-speed DSL (ADSL2+), FTTC, FTTN (VDSL), and FTTH	Ethernet-based PON to the home Ethernet backhaul from remote terminals, IP DSLAMs Metro Ethernet transport/backhaul to major service nodes
3G mobile services	One to four T1s or E1s from cell towers to aggregation points in the radio access network, with higher-speed links from aggregation points to larger offices	Backhaul from base stations using Ethernet over copper, fiber or micro-MSPP SONET/SDH platforms
Business Ethernet services	Ethernet over Sonet/SDH MSPP Ethernet over dedicated fiber Ethernet over WDM Ethernet over IP/MPLS (PWE3/VPLS)	Adoption of carrier Ethernet and packet transport platforms, moving away from SONET-based MSPPs
Enhanced consumer broadband	ATM and IP DSLAMs	Ethernet backhaul from IP DSLAMs
SAN extension	Dedicated fiber facilities or WDM facilities for Escon/Ficon, Fibre Channel, GigE and 10-GigE	Migration away from traditional SAN and mainframe protocols to 10-GigE over WDM
Content delivery	Typically leased facilities (dark fiber, OC-n, DS3, etc.) from wholesale operators	ASPs moving toward private ownership of optical networks Ethernet-over-optical architecture for high-speed content delivery

Source: Heavy Reading, 2006

Ethernet equipment suppliers used this as an effective marketing tool, and a number of operators deployed pockets of carrier Ethernet gear to support particular service initiatives.

Yet along the way, operators found that carrier Ethernet gear was not equivalent to SONET, particularly in terms of its robustness and its operations and management capabilities. In addition, the carrier Ethernet market is quite fragmented, with some gear derived from edge routers, others derived from large-scale enterprise LAN switches—so a common definition of a carrier Ethernet network element was hard to find. The result is that in nearly every carrier network today, you find a mix of technologies supporting Ethernet services, not a single multipurpose, Ethernet-optimized solution.

In Figure 3, also from the Heavy Reading survey of network operators, Layer 3 routers and carrier Ethernet switch routers (CESR) are the products that operators are deploying or trialing most commonly today, but their lead is slight over rival platforms, including the SONET-based multiservice provisioning platforms (MSPPs). Most large operators today deploy a broad mix of platforms to support Ethernet services, depending on such variables as geography, customer service requirements, topology and institutional bias.

This method has served the carrier community well enough to date as they make their initial Ethernet transport purchases. Looking ahead, however, it will be in the carriers' interests to embrace a standards-based packet transport platform, one that can be deployed widely throughout their networks for all their aggregation and transport applications. This should simplify the addition of new services, and reduce the capital and operations costs associated with operating multiple platforms for multiple services.

This brings us to the carriers' dilemma in the marketplace today: They own billions of dollars' worth of functional and fully depreciated SONET transport, yet future growth and profitability will be tied to IP and Ethernet services.

In a perfect world, a carrier would prefer true convergence between SONET and Ethernet, blending all the positives of each and leaving behind Ethernet's limited OAM, protection, and performance management as well as SONET's costly packet inefficiencies and rigid TDM-based

FIGURE 3 Network Operator Choices For Ethernet Services Support

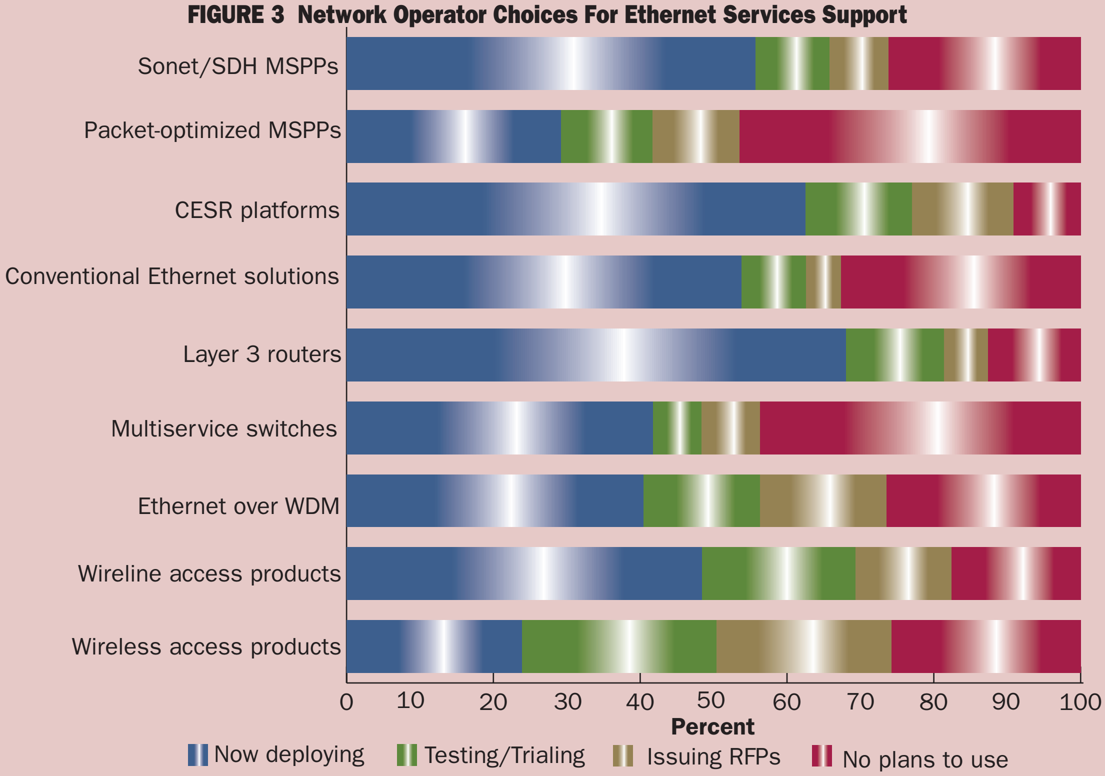

Source: Heavy Reading 2006 Survey Of Ethernet Service Providers, 151 respondents

multiplexing hierarchy. Next-gen SONET standards—for framing, concatenation and link capacity adjustment (G.7041, G.7042 and G.707)—and products, notably the MSPP, took important steps in that direction, but these do not completely overcome the limitations of TDM-based switch fabrics or hybrid switch fabrics. MSPPs simply do not offer enough transport efficiency for the future packet network, and they still require operators to force-fit packet flows into a circuit switched model.

The carrier Ethernet switch router (CESR), however, also has not quite delivered on its promise as a true replacement technology. Although CESRs are certainly optimized around Ethernet aggregation and packet scaling, they often have rudimentary or limited OAM (operations, administration and maintenance), protection and circuit emulation support. In a recent survey of transport network planners and engineers, Heavy Reading found many operators still not convinced Ethernet is on par with SONET, citing such factors as immature OAM, no end-to-end provisioning and visibility, immature circuit emulation technology (T1s and E1s won't go away overnight), and lack of SONET's dependable 50-ms protection (Figure 4).

Will T-MPLS Be The Key Ingredient?

The emerging consensus today is that MPLS will be necessary if Ethernet is to fulfill the promise of packet transport. MPLS first appeared on the transport networking scene in the form of GMPLS, a set of protocols or protocol extensions offered up by the IETF that could be implemented as a distributed control plane for the setup and tear-down of connections (circuits, wavelengths or whole fibers) in a transport network.

The success of GMPLS thus far has been modest at best, as it asks large operators with substantial core optical networks to migrate from their traditional centralized top-down management and provisioning model to one that is distributed, automated and dynamic in nature. In addition, GMPLS's focus is solely on improving the management of core optical resources, not on packet transport efficiency. Some carriers have taken advantage of GMPLS to support topology and neighbor discovery in order to improve provisioning of wavelengths or circuits in their networks, but in most cases this is limited to network segments based on a single vendor's equipment, typically core optical switches and large MSPPs.

Despite this slow start, a new version of MPLS is under development that will offer packet-based alternatives to SONET circuits. This new version of MPLS for packet transport goes by the name Transport MPLS, or T-MPLS for short.

T-MPLS promises operators much greater flexibility in how packet traffic is transported through their metro and core optical networks. In T-MPLS, a new profile for MPLS is created so that MPLS label switched paths (LSPs) and pseudowires can be engineered to behave like TDM circuits or

Optical system vendors are embracing T-MPLS, but what about their carrier customers?

Layer 2 virtual connections, in effect replacing SONET section, line and path. SONET will remain, but only as a physical interface (PHY) choice, along with Gigabit and 10-Gigabit Ethernet.

T-MPLS is a work in progress by the ITU-T in concert with the IETF, bringing the forces of the major telecom operators worldwide together with the IP-centric talent within the data community. As envisioned today, T-MPLS includes these key features:

FIGURE 4 Perceived Impediments To Ethernet Transport

Rank	Impediment
1	Immature OAM
	Lack of fiber to the customer premises
3	No end-to-end provisioning
4	Immature circuit-emulation technology
5	Lack of SONET/SDH-like protection
6	Cost of transition
7	Installed base of CPE does not support Ethernet handoff
8	Lack of demand

Source: Heavy Reading, 2006

- A connection-oriented networking model, which assumes these connections will have long holding times.
- Engineered point-to-point bidirectional MPLS label switched paths (LSPs) to mimic the current TDM transport network's circuit paradigm.
- Client/server layering and partitioning for network scalability.
- End-to-end LSP protection, with survivability methods mimicking the current linear, ring and shared mesh options.
- Per-layer monitoring, fault management and operations that leverage existing SONET OAM mechanisms.
- Strict data and control plane separation.

T-MPLS does not use Layer 3 and IP-related features, such as LSP merge, penultimate hop popping, and equal cost multiple path (ECMP), eliminating the need for packet transport network elements to process IP packets, while also improving OAM. For more information on these issues see *BCR*, December, 2006, pp. 8–9 and pp. 16–17.

T-MPLS: Another God-Box Or True Packet Transport?

Today, many of the largest optical systems vendors are embracing this T-MPLS concept, as it gives them a clear migration path for their SONET/SDH MSPPs, products that have shown tremendous growth over the past six years, but whose growth is forecast to flatten and even diminish as carriers adopt carrier Ethernet and WDM-based transport platforms in their place. Carriers are also actively participating in the standardization of T-MPLS, and see in this the potential to integrate their IP/MPLS core packet networks with their transport networks, fulfilling their next-generation network (NGN) migration goals.

Yet one must be skeptical of a new standard that claims it will once and for all enable packet-optical layer convergence. Is T-MPLS motivated by SONET/SDH vendors hoping to stay relevant in the face of carrier Ethernet's rise? Or does T-MPLS address shortcomings in carrier Ethernet and next-gen SONET/SDH that cannot be solved effectively by other means? The market will take some time to decide.

First, the vendors must deliver on the promise of T-MPLS. Currently, there is back-and-forth between T-MPLS advocates and those who would prefer to keep the networking layers as they are, with Ethernet and IP firmly rooted in the packet layer of the network and transport limited to SONET/SDH and WDM.

Network operators can't be blamed for feeling a little fatigued, having heard a great deal about the revolutionary capabilities of MPLS, yet finding implementation thus far (primarily in a cap-and-grow substitution for ATM) to be much more complicated than promised. For T-MPLS to succeed, it must ultimately address operator concerns about its interoperability with IP/MPLS, its support for multipoint-to-multipoint connections, its potential conflict (or alignment) with recent OAM standards work for Ethernet, and perhaps most importantly, the broader concern that the billions of dollars of investment required to migrate the current SONET transport network to T-MPLS will be worth it.

As has been typical through the history of carrier transport networking, a platform's success will not be defined by its promise to lower operations costs, but by its association with a popular service. Today, business Ethernet services, IP-VPNs, mobile data and consumer broadband video are the hottest offerings in the market. What T-MPLS and packet transport have in their favor is their professed ability to aggregate and transport all these services with the greatest packet efficiency possible, while preserving carrier networking and OAM models. In 2007 we'll begin to find out if T-MPLS can make the sale□