Carrier Ethernet:
Not Just in the MAN

296 Concord Road
Suite 308
Billerica MA 01821-3487

978.715.2300 tel
978.715.2400 fax

www.soapstonenetworks.com

Ethernet is seen by many as a metro-area service architecture targeted at enterprise customers. While that is certainly a valuable mission for Carrier Ethernet, it is just one of many possible missions. Ethernet, like IP, is becoming a complete architecture for service deployment. Those providers who've become Provider Backbone Transport (PBT) aware have come to recognize this, seeing Ethernet and PBT as important tools in a multi-dimensional, multi-technology view of convergence.

Motivation for Change

Network operators are challenged by the radical shift in service revenue and technology that has occurred in the last two decades. In the late 1980s, nearly all service revenue was generated by wireline voice and leased-line services, based on traditional time-division multiplexing and SONET/SDH or “circuit-switched” network infrastructure. By the late 1990s, it was clear that the growth of the Internet and the shift of business buyers to packet-based services (including frame relay, ATM and IP services) would create a completely new revenue and service model. Today, faced with increased competition in the “bit-pipe” business, carriers are looking to drive new top-line growth from a growing array of value-added services as well as turning to smaller enterprise customers to fuel their financial future. In doing this, they are adding to another dominant trend - customer demand for orders of magnitude more bandwidth than consumed just a few years ago - and fueling the need for automated turnkey service offerings for small and medium enterprises outsourcing IT to the carrier. As a result carriers must find a way to satisfy customer demand for products by creating a portfolio that includes multiple and clearly differentiated services from the simplest bit-pipe through sophisticated applications while increasing the automation through the entire service lifecycle.

Although the transformation and convergence onto IP networks sounded good, carriers have found that legacy data network customers have not been willing to pay the same amount of money for an IP-based service. Customers are missing key elements of the traditional circuit-switched services that demonstrate the carrier’s deep understanding of how important its service is to them.

- Can customers choose a service that’s right for them?
- Can the carrier exceed their expectations for order fulfillment?
- Is the carrier watching over the SLA’s on their service better than they can?
- If anything isn’t met, is the customer’s bill adjusted without asking?
- If the service is down, does their carrier tell customers about it before they notice?
- If customers do put in a call, does the carrier know exactly what happened and how long it will be before service is restored?

Can carriers find a way to grow their top and bottom lines with new packet-based services while regaining the customer satisfaction they enjoyed during the era of circuit-switched services?

The development of NGN Frameworks

As demand for packet services has grown and surpassed the demand for voice and circuit services, traditional carriers have found themselves operating separate circuit- and packet-switched networks. The ITU has been developing a Next Generation Network (NGN) architecture to enable delivery of a wide range of telecommunication services over a single packet-based infrastructure. NGN uses multiple broadband, Quality of Service (QoS) enabled transport technologies. It makes service-related functions independent from the choice of underlying transport technologies. As a result, this architecture frees telephony providers from the constraints and costs of the stovepipe architecture common in PSTN networks in which each service would have its own dedicated or overlay network.

NGN's many benefits aside, packet services present special challenges to operators and customers used to circuit-based behavior. Circuit switching manages congestion purely through admission control since trunk bandwidth is allocated end to end. Packet services require traffic and congestion management that becomes more complex as the size of the network increases. Additionally, there is a constant risk of congestion-induced service failure where network load is high and the traffic is bursty. Circuits are "stateful" and set up along specific paths. Knowing those paths makes it easier for a carrier to protect against and quickly react to and repair node and link failures. Packet networks provide adaptive behavior to increase resiliency, but operators cannot easily determine the paths taken by their customers' critical data.

In many ways, IP exacerbates these problems. IP networks typically support Internet traffic as their major load. Consumer traffic variations can impact network load and performance even for business services. Carriers use Multi-protocol Label Switching (MPLS) to direct the flow of traffic in their networks. However, MPLS operation is intertwined with IP and consequently may inherit much of the adaptive behavior issues, congestion and security problems associated with IP.

It all comes down to the fact that networks don't have one topology - they have three: the "logical topology" of a service which the endpoints can address each other; the "traffic topology" of a network showing the actual path that the traffic follows between those endpoints; and, the "physical topology" of a network which is critical for availability management and recovery from failures.

The loss of independent control of these three topologies isn't an academic issue; it's a major problem and not just in the metro network. Senior executives from one of the world's leading communications solutions and services providers have been quoted often on the problems of converging legacy frame/cell services onto IP and specifically bemoaned the fact that premium services wouldn't "stay where they were put." Anyone who has tried to write a very meticulous SLA for IP services (i.e., the type of SLA routinely provided for frame relay) knows that the lack of stateful resource assignment in IP makes it very difficult to obtain.

Can carriers regain full control of their networks using packet technology in which adaptive behavior is not an issue?

The movement toward Carrier Ethernet

There is a real value in having an independent Level 2 infrastructure that can create "routes" that are associated with major traffic patterns and can be managed on how the physical infrastructure is used. This mission was supported by ATM in the early broadband deployment period. Now, providers are looking to Ethernet for this critical role. The major reason for that is PBT (Provider Backbone Transport), which is also being standardized by the IEEE as PBB-TE (Provider Backbone Bridge – Traffic Engineering).

Ethernet interfaces have been used on routers and other packet equipment for years, but there is a major difference between a point-to-point Ethernet mission and a "traffic topology" mission. Ethernet was developed as a standard for local-area networking and, in its LAN manifestation, presented major issues in terms of scalability and suitability for carrier operations. Ethernet's limits on the number of VLANs that can be present in a single network made it difficult for carriers to support as many customers as they'd like without resorting to non-standard means. In addition, as an enterprise technology, Ethernet did not include OA&M functions found in carrier technologies such as SONET, nor did it provide SONET's rapid fail-over capability. Finally, while plug-and-play operation is often desirable for enterprise operation, it is traffic engineering – the easier and more automated the better – that is notably more important to carriers, especially when they have SLAs to meet.

There are a number of standards groups working on Carrier Ethernet issues, including the IEEE, ITU and Metro Ethernet Forum (MEF). Formed in 2001, the MEF first addressed the issues of Ethernet services as a means of augmenting enterprise LANs in creating metro networks among major corporate sites. The activity has expanded as service provider interest in Ethernet evolved from a simple service interface to an architecture for a Carrier Ethernet ecosystem.

The Goals of Carrier Ethernet

The MEF developed five critical attributes for a Carrier Ethernet architecture and has worked to develop complete solutions in each area. The primary requirements are:

- Standardized services through the use of technology-neutral abstractions of service behavior.
- Scalability to metro, national and even global deployment to insure providers there is no risk of outgrowing a Carrier Ethernet investment.
- Reliability to insure an Ethernet layer would contribute to greater network availability.
- Quality of service to insure Carrier Ethernet could support the full range of service offerings currently available, many of which demand managed QoS and stringent service-level agreements.
- Service management to insure Carrier Ethernet could be linked to service and operation processes, thus supporting effective and efficient OAM&P.

The standards work from IEEE, ITU, IETF, and MEF has created a structure like that shown in Figure 1. A provider network of optical and Ethernet devices create a physical infrastructure in the metro and long-haul. Over this, a transport abstraction hides the differences in underlying technologies and facilitates the creation of point-to-point, multipoint and one-to-many services.

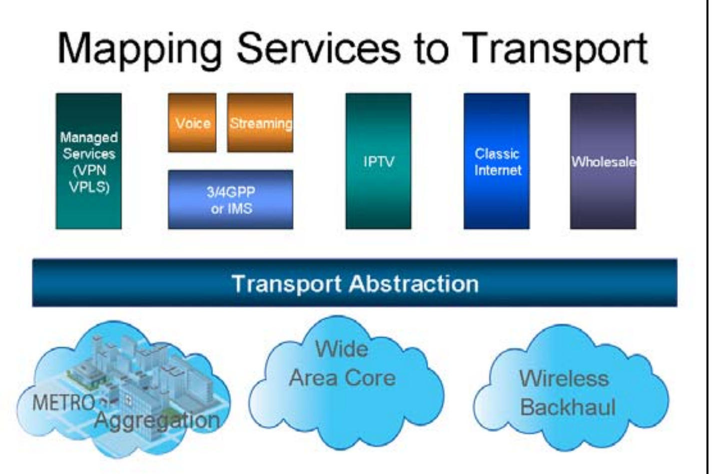

Figure 1

Some of these services are directly sold to end users while others form the foundation for "service infrastructure" supporting IPTV, broadband Internet and other applications.

The key standard in the Carrier Ethernet architecture is the IEEE 802.1ah or Provider Backbone Bridge (PBB) standard. This standard incorporates what is often called "M-in-M" or "MAC-in-MAC" encapsulation to expand the number of "service VLANs" that can be supported on the network. When used in conjunction with other IEEE, IETF, and ITU standards the PBB standard creates most - but not all - of the features needed to meet the five requirements.

Predictability is achieved by PBT by ceding control of the "routing" algorithm to a separate control plane that determines the path that traffic takes through the network. Centralized provisioning gives carriers two major advantages. The first is exact knowledge of the path each customer's traffic takes. The second is precise and efficient allocation of network resources.

Building on this technology, carriers will be able to meet several important goals:

- offering and delivering precise SLAs
- achieving higher availability while lowering operating costs
- profitably expanding their target markets to include small businesses

These goals are made attainable by two capabilities unique to Carrier Ethernet in the packet-switching world: carrier-class OAM&P (including 50-millisecond failover) and finely tuned, automated traffic engineering enabled by an external control plane.

Carrier Ethernet inherits the full range of Ethernet OAM&P enhancements, including the IEEE 802.1ag OAM features and Ethernet Protection Switching (G.8031). The former provides frame loss, delay and jitter information as well as service availability and utilization information. Unlike MPLS LSPs, which are difficult to trace internally, Carrier-Ethernet paths can be traced for troubleshooting and are not impacted by adaptive routing changes. This allows fail-over paths to be pre-calculated if desired. These capabilities - along with PBT - make it possible to write exceptionally stringent SLAs for Carrier Ethernet services - resolving one of the major provider problems with enterprise services based on IP and MPLS. Similarly, Carrier Ethernet's immediate, automated fault detection means carriers are notified about network problems as soon as they occur.

The purpose of a service management framework is to build the link between the service experiences being offered by the service provider and the resources that sustain those experiences. In the early days of data networking (when data services were sold to enterprises for tens of thousands of dollars per month) the service-to-resource linkage was created through manual provisioning.

Although carriers still occasionally have to resort to manual provisioning, as a general practice, manual processes have fallen out of favor because the average revenue per user (ARPU) of broadband and data services is declining. Instead operators are looking for some form of process automation to create and maintain service-to-network linkages. For this to be effective, the connection and performance goals of the service must be translated into resource commitments automatically. This means converting the **abstract** vision of a service as a set of behaviors into the **provisioned** version set of resource commitments. Element/device management systems, network management systems and policy management systems are all useful for the provisioning portion of this process, but less so for the conversion from abstraction to provisioning.

This example illustrates the problem: A six-site, two-city Ethernet VLAN service can be conceptualized as a simple six-port LAN. However, to provision this service, it will be necessary to assign a VLAN ID, configure LAN switches, assign service quality levels, establish OAM processes, etc. If "automating" these tasks means nothing more than giving a human operator the ability to enter parameters through an attractive GUI, little has been accomplished in reducing operations costs. Compound this scenario with the certainty that infrastructures consist of multiple device types and multiple equipment vendors. What must be done is relate the properties of that abstract but simple six-port LAN to the complex configuration of a metro Ethernet infrastructure.

How can a management system know to set a specific parameter on a specific port of a specific switch?

What is Missing from Carrier Ethernet Today?

Carrier Ethernet certainly provides many benefits to carriers – and with a few advances it would further its adoption.

Incorporating Carrier Ethernet into the NGN architecture would provide detailed architectural guidance to carriers who would like to build Carrier Ethernet into their infrastructure. Additionally, it would provide more detailed technical requirements to vendors who would like to offer other Carrier Ethernet products that interwork with the elements of an NGN network.

Adoption into the NGN architecture aside, Carrier Ethernet still needs to turn one of its promises into reality: relief from the problems of spanning trees and traffic engineering. The basic spanning tree protocol (STP) of Ethernet had already been augmented to Rapid Spanning Tree Protocol and Multiple Spanning Tree Protocol (RSTP and MSTP), but these still create data plane and control plane constraints on the size and complexity of the Carrier Ethernet network. It is in this area that a debate among standards groups, providers, and equipment vendors has emerged. One side has promoted an extension of Carrier Ethernet called "Provider Backbone Transport" or PBT, and the other a variation on MPLS called "Transport MPLS" or T-MPLS. The debate is often characterized as an "Ethernet-versus-MPLS" dispute, but the issues are more subtle.

It is probably easiest to start with what the various approaches PBT and T-MPLS have in common. Both of the suggested enhancements to the Carrier Ethernet framework create an alternative to the "Provider Backbone Bridge." Both have a specific "edge" and form a network-within-a-network with special internal behavior. Both also create a more manageable control plane architecture based on something other than the usual topology update messages. Supporters of PBT and T-MPLS also agree on what control plane architecture would be used—Generalized MPLS or GMPLS. And, they agree on one very important point, **that there must be a connection-oriented transport technology layer as Layer 2 of provider packet networks**.

GMPLS' role as a control plane for both architectures stems from its roots as a control plane for optical routing, which was initially called it "Generalized Multi-Protocol *Lambda* Switching" (to reflect its role in routing optical connections). Since optical devices didn't exchange topology information, GMPLS assumes this information can be gathered through a separate control plane layer using standard discovery protocols to learn topology. The results are then "fed downward" to lower layers to manipulate connection behavior allowing considerable control over the conditions that result in updating the topology. It is also possible to gain topology awareness from a higher-layer control plane exchange such as the one already used for IP/MPLS. The key requirement for GMPLS is the correspondence between the nodes and trunks in the control plane and the data plane so the topologies below can be correctly represented. **The PBT/TMPLS debate may be an indication that the issue of topology and configuration management by an independent higher-level control plane is of general importance to operators**.

The technical capabilities of PBT and T-MPLS are substantially the same, but PBT is a simple evolution from other Carrier Ethernet technology and is easily introduced into many Carrier Ethernet products. PBT builds upon IEEE 802.1 to support the use of provisioned point-to-point tunnels to link elements of L2 VPNs linking the pieces of E-LINE, E-LAN, or E-TREE services. In PBT the generation of spanning tree updates is suppressed. As a result, there is no control plane behavior to build a bridging table. Instead, bridging tables are built using an external management system. T-MPLS is a simple extension to router technology, just as PBT is an extension to Ethernet technology. Since routers are typically more expensive than Ethernet switches, PBT will be lower in cost unless the network is already supporting routing.

Delivering on Carrier Ethernet's Promises

Carrier Ethernet is an opportunity for carriers to position their services to their greatest advantage. To the retail customer, Ethernet is a convenient, familiar and cost-effective packet-transport technology. To the carrier serving that customer, Carrier Ethernet is an opportunity to offer a simple and useful baseline upon which tiers of value-added services, ranging from IP-services to VoIP, can be built.

Carrier Ethernet has similar value in wholesale services. There the tiers of service begin dark fiber, possibly followed by lambdas and then perhaps a framed layer 1 service such as SONET. Carrier Ethernet would be an excellent choice for a lightweight packet-switched service that could be offered on top of such an infrastructure as a wholesale service offering. Because Carrier Ethernet is transport agnostic, the carrier has complete freedom to evolve or build out its layer 0/1 infrastructure in whatever way is best. Likewise a single Ethernet Virtual Circuit can span a carrier's metropolitan and long-haul infrastructure regardless of what combination of layer 0/1 technologies may be in use. These characteristics and Carrier Ethernet's support for precise SLAs makes it an excellent basis for wireless backhaul and similar applications.

A carrier's choice of control plane technology may turn out to be as important as the make and model of the switches it buys. Switch vendors will implement Carrier Ethernet in compliance with the IEEE and ITU specifications and recommendations. Consequently, the switches' data forwarding capabilities should be reasonably similar (switch capacity aside). And, we can expect that support for external control planes will become a new point of switch vendor differentiation. The choice of an external control plane solution will make the biggest difference in the value the carrier gets from its Carrier Ethernet infrastructure because that is where we can expect the greatest amount of vendor innovation with the broadest range of product capabilities.

In fact, external control planes have the potential to add value to every switch in a carrier's network. The right control plane solution can reduce the cost and time to provision services, improve network availability and resiliency, and increase the network's effective capacity – no matter which vendor built the Carrier Ethernet switches.

While no switch vendor has announced plans to field its own external control plane offering, it is reasonable to expect they will, especially as Carrier Ethernet extends beyond the metro-area and evolves into a complete architecture. On the other hand, carriers who want to ensure they can operate a multi-vendor network would do well to seriously investigate third-party external control plane offerings, even if they have only one switch vendor in their network today.

Dynamic External Control Plane
Soapstone Networks™

As Carrier Ethernet continues to increase its influence on networks and industry standards, external control planes will become a key element in the deployment of Ethernet services. The separation of the control plane from the data plane decouples service provisioning from the underlying network technology, allowing carriers to deploy heterogeneous networks and to evolve their network and service offerings independently. An external control plane accomplishes all of this by taking a holistic view of the network and its resources as it computes the paths that each customer's traffic will take. However, merely computing paths is not enough. How do you provision effectively and efficiently manage your services and multiple resources?

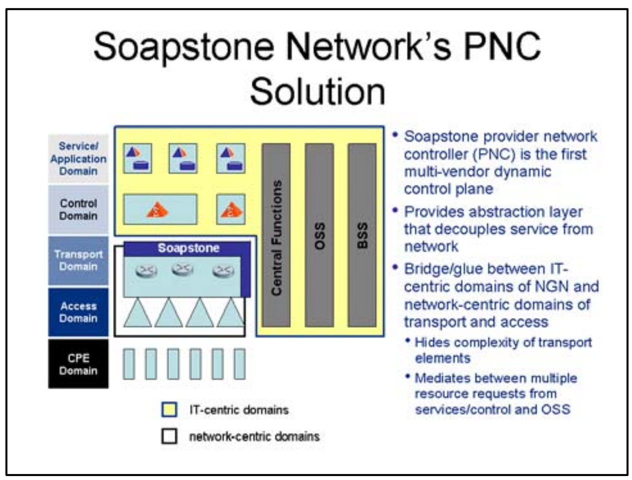

Addressing the challenges in building services across today's evolving network infrastructure, Soapstone Networks™ PNC (Provider Network Controller) implements service activation and layer 0-2 management tools for PBT, MPLS, T-MPLS and optical networking platforms. PNC is technology agnostic and designed for networks that encompass single or multiple switch technologies.

Unlike traditional, static service provisioning, PNC continually monitors and optimizes network resource commitments, providing a far higher level of network reliability and efficiency than has previously been possible with packet-based technology. The Soapstone Networks PNC mediates between multiple resource requests from the OSS services/control layer via a standard SOA based interface. It dynamically models the available and occupied capacity of the network, selecting the most effective path based on modeling customer and carrier QoS and policy requirements. Services on the network are dynamically migrated in the event of a failure, to restore normal primary and backup paths when the fault has been repaired, or to allow network equipment to be taken on/off line. The PNC continuously monitors for SLA compliance as well as network faults, enabling intelligent fault remediation, as well as fault reporting between network operation and provisioned services on a service correlated basis.

Carrier Ethernet is proving to be an important element in today's evolving carrier network infrastructure. Soapstone Networks, through its PNC product, accelerates service creation and deployment and increases network resiliency, allowing carriers to take full advantage of their investments in Carrier Ethernet by quickly, flexibly, and profitably deploying the service offerings that their customers demand.

soapstone NETWORKS

About Soapstone Networks:

Soapstone Networks is at the forefront of the movement to Carrier Ethernet by delivering resource control systems that realize NGN software-provisioned services in the new Carrier Ethernet transport network. Soapstone's common control framework decouples services from underlying network technologies. The Soapstone solution dynamically provisions precise, SLA-quality services, continuously optimizing utilization of network resources to bring orderly, predictable business-driven behavior to service provider networks. The future of Carrier Ethernet - www.soapstonenetworks.com

296 Concord Road
Suite 308
Billerica MA 01821-3487

978.715.2300 tel
978.715.2400 fax

www.soapstonenetworks.com