Getting A Grip On Wide Area Ethernet

Michael Finneran

The technology is poised to replace some private lines and frame relay—over the long term.

After years of anticipation, Ethernet is making its move into the wide area market. Metro-area Ethernet services were pioneered by startups like Yipes and Telseon during the late-'90s network boom, and after a slow start, service revenues are now growing at double-digit rates.

LAN switching over point-to-point trunks eliminates the need to limit Ethernet's transmission range, and carriers are now offering both metro-area and wide-area services; hence we will use the term wide area Ethernet (WAE) rather than the more familiar "metro area Ethernet (MAE)." While startups introduced the concept, the incumbent local exchange carriers (ILECs) like SBC and Verizon are now in the best position to capitalize on the market.

At a high level, the use of Ethernet as a wide area transmission technology has tremendous appeal both in terms of economics and simplicity. Ethernet switches are far cheaper than SONET transport and earlier-generation packet forwarding technologies. From a simplicity standpoint, better than 95 percent of data traffic begins and/or ends on an Ethernet interface, so why use a router to generate a frame relay, PPP or ATM interface? The mantra for network simplicity has become: "IP on Ethernet on wavelengths." Those economics are also reflected in the carriers' opex budgets.

However, wide area Ethernet is also throwing more confusion into the already muddled services market. Depending on how it's configured, WAE can be viewed as a competitor to the entire acronym alphabet from ADSL to WiMAX! The difficulty is that WAE can be used as an access technology, a wide area service or both.

Michael Finneran is president of dBrn Associates, Inc., an independent consulting firm in Hewlett Neck, NY, specializing in the design and installation of domestic and international networks. He is the instructor for BCR*'s data communications courses on "Wireless Networks For Voice And Data" and "Wireless LANs."*

Configurations And Applications

The Metro Ethernet Forum (MEF) defines two different Ethernet services: E-LAN and E-Line. As shown in Figure 1, E-LAN is a switched service that uses carrier-based Ethernet switching equipment. E-Line, on the other hand, is essentially a high-capacity (e.g. 1 Gbps), point-to-point private line connection between two LAN switches, supporting an Ethernet interface at each end.

WAE has three major applications, according to Rich Klapman, group manager of converged access services for AT&T:

- **Internet Access:** In this configuration, Ethernet competes with traditional fiber- or copper-based private line, DSL, cable modem or broadband wireless access. According to Klapman, this is the major application AT&T has found for MAE. Carriers support access rates between 2 Mbps and 1 Gbps, though 50 Mbps to 100 Mbps seems to be the "sweet spot" today. AT&T customers will soon be able to use the Ethernet service to access any virtual private network (VPN) services, including frame relay and ATM.
- **Point-to-Point**: Either E-Line or E-LAN service could support a high-capacity connection between two customer Ethernet networks. While the Layer 1 E-Line service is more typical today, those point-to-point connections could also be provisioned through the carrier's switched Ethernet infrastructure. The obvious advantage of E-LAN over E-Line in point-to-point applications is the ability to statistically multiplex multiple users over the same shared transmission channels.
- **Multipoint LAN/WAN**: Wide area Ethernet can also be used to provide connectivity among geographically dispersed user sites within the same city (metro area) or across the country (wide area). In that application, WAE would compete with wide area private line, frame relay, ATM and Internet-based VPNs.

Clearly, there are lots of overlapping areas. A customer could buy an Ethernet service for Internet access and share the connection between public Internet access and inter-site communications between facilities in the same city. The picture becomes even more muddled when the inter-site connection is made via an Internet VPN service: An organization might use Ethernet-based Internet access at their main location, while smaller sites might use traditional private line or DSL.

Use BCR's Acronym Directory at www.bcr.com/bcrmag

FIGURE 1 MEF-Defined Service Configurations

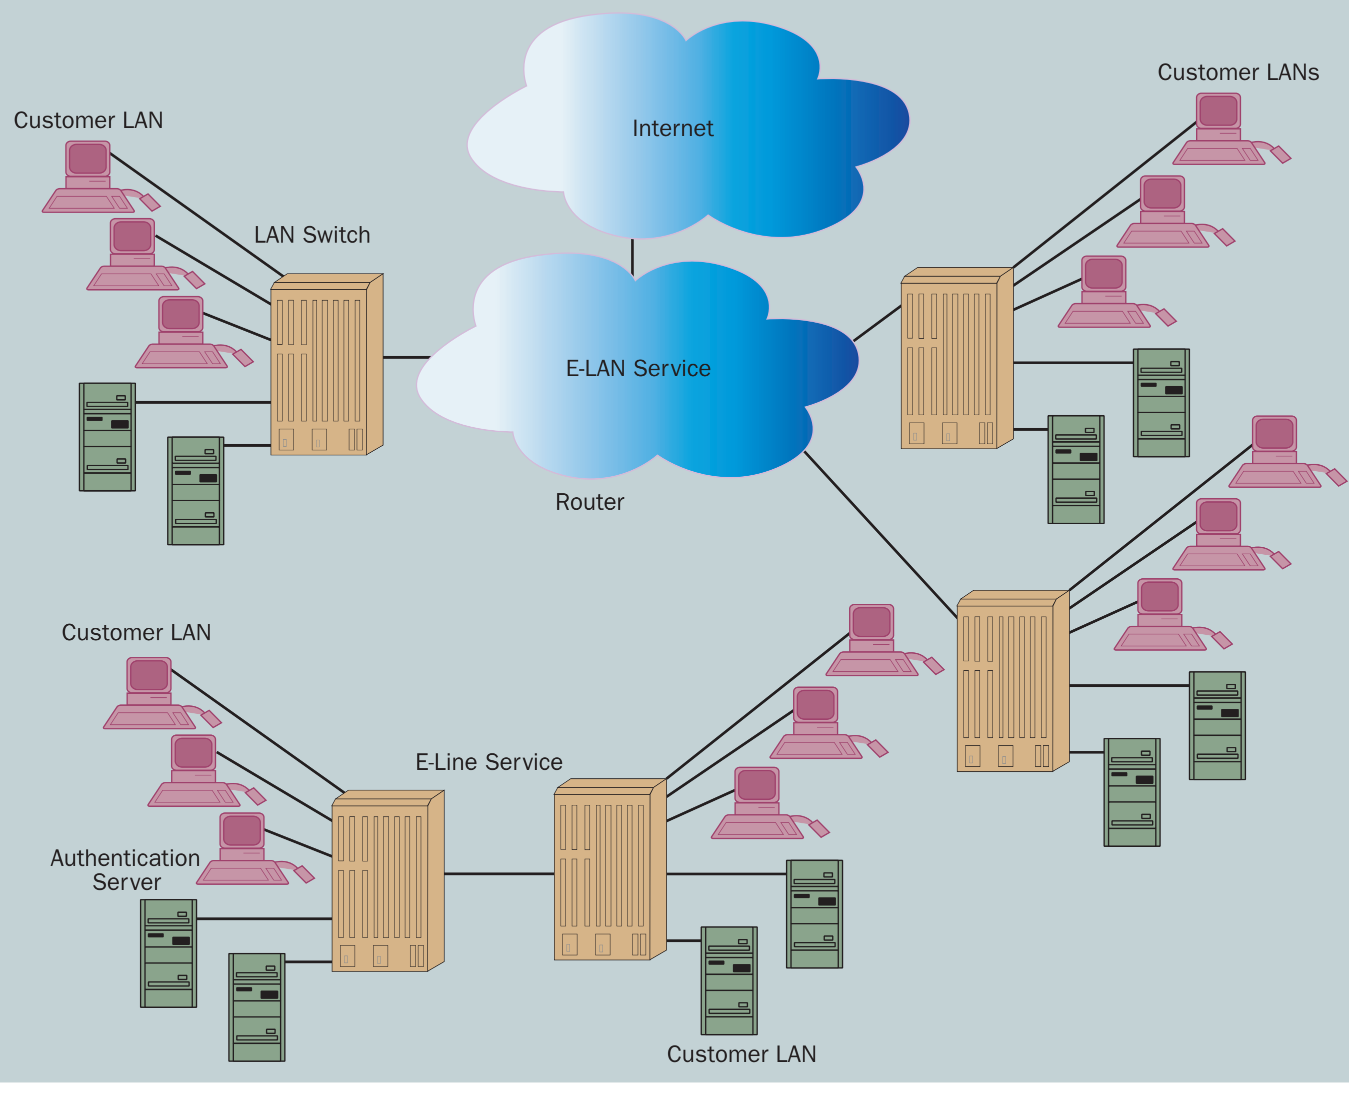

In short, it's tough to categorize exactly where Ethernet services will fit in relation to other services, and depending on the overall network configuration, they could be viewed as either complements or competitors.

Sizing Up Options And Prospects

In the shifting picture of wide area data services, it is difficult to get a clear focus on what customers are buying. Despite AT&T's view that the market is heavily weighted toward Internet access, market researchers see the majority of the network connections supporting LAN/WAN connectivity. Erin Dunne, director of research services for Vertical Systems Group, estimates the U.S. wide area Ethernet market for 2004 at $518 million, with about 60 percent of the ports supporting LAN/WAN applications. That puts Ethernet revenues at a fraction of the $9.8 billion frame relay or $12 billion private line markets. However while frame relay is basically flat and is expected to decline, Dunne expects the WAE market to grow at roughly 30 percent per year through 2008. AT&T has seen its Ethernet service revenues grow 100 percent in the past year.

In the wide area services mix, the major challenge to frame relay and private lines will come from Internet services—with or without Ethernet access. The generic term "VPN" is used to describe site-to-site or LAN/WAN communications through the Internet, but two vastly different service configurations are used. The carriers are pushing IP-based services using Multi-Protocol Label Switching (MPLS). These managed services can provide full-mesh connectivity using virtual circuits called Label Switched Paths (LSPs). Besides the inherent security provided with a virtual circuit, the appeal of MPLS is that it will allow the carrier to provide quality of service (QOS) differentiation, so that voice and video applications can see shorter transit delays and fewer dropped packets.

Many users have found a more cost effective solution using basic Internet service and secure

E-LAN transmission rates range from 2 Mbps to 1 Gbps

tunnels. The main drawback is that basic Internet service is "best effort," so it does not support QOS to prioritize delay-sensitive voice and video traffic. However, the quality of basic Internet service has been improving, and ironically, as the carriers improve their basic Internet service, they are giving the user less incentive to invest in the higher-priced MPLS option!

Ethernet service could be used to provide a high-capacity access to either of those Internet-based service configurations. Further, E-LAN services use virtual LANs to separate customer networks to ensure security, and most support QOS capabilities similar to MPLS.

Anatomy Of A Wide Area Ethernet

So what are you going to see when you sign up for a WAE service? Well, you will see a 10 Mbps to 1 Gbps Ethernet connection, but beyond that, it's anyone's guess. Unlike frame relay or private line, where the service configuration has long been standardized, WAE is a moving target. The vast majority of customer access connections today are fiber-based, but any number of carrier-provided equipment configurations can sit between your Ethernet interface and the carrier's fiber.

For a Layer 1 E-Line service, a SONET add-drop multiplexer (ADM) is typically used to transport the channel to the carrier's central office. That point-to-point connection could also be connected directly to a channel on a wave division multiplexer or a wavelength channel split out using an optical ADM (OADM). The carriers seem to prefer the redundancy offered by a SONET ring, particularly the "Next Generation" SONET equipment that supports Generic Framing Procedure

FIGURE 2a E-Line Service

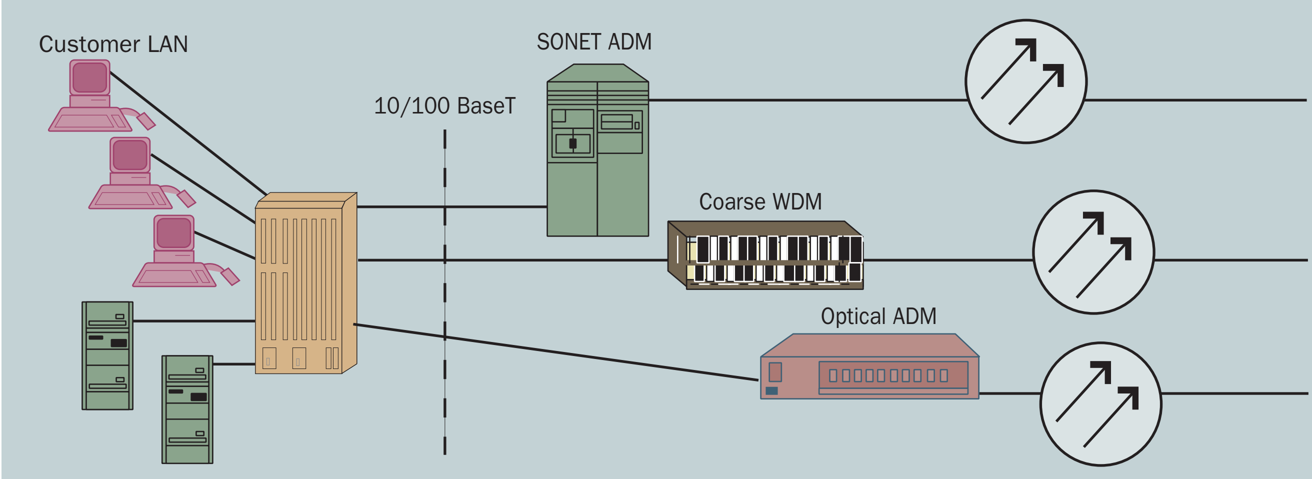

FIGURE 2b E-LAN Service

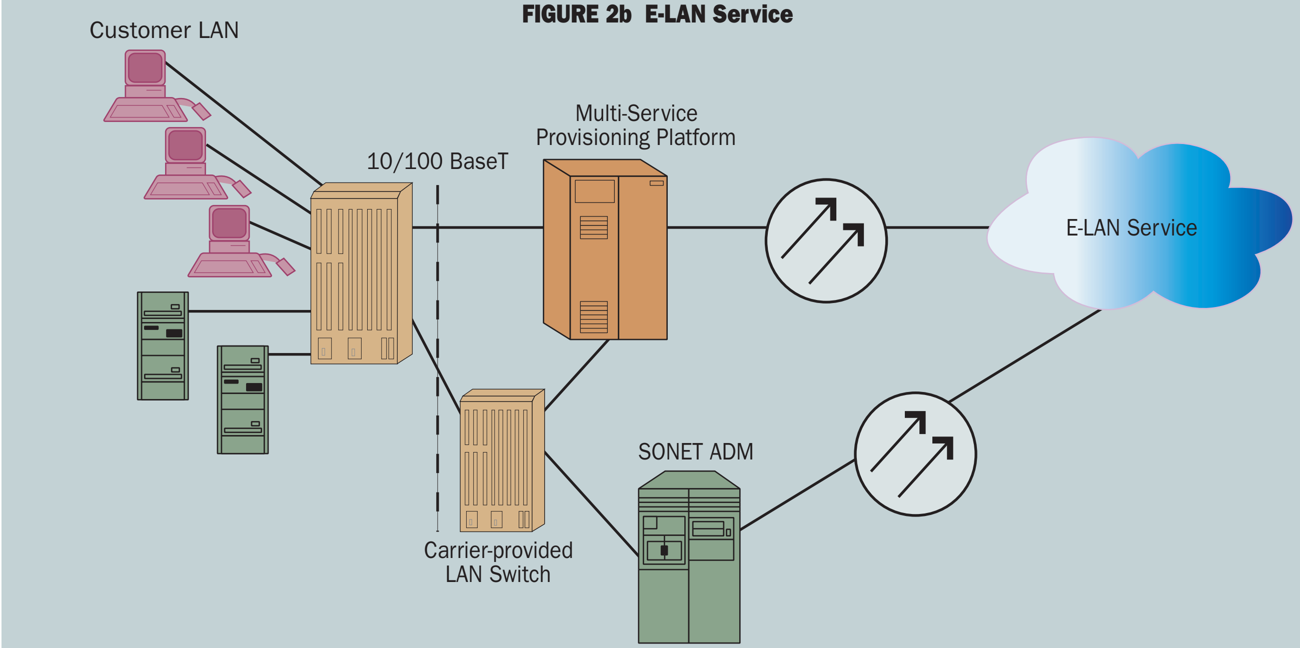

(GFP) and Virtual Concatenation (VCAT) for greater link efficiency (Figure 2a).

For E-LAN services, some type of Layer 2-aware device is needed, because the network capacity is metered. While the access connection may operate at 100 Mbps or even 1 Gbps, each carrier offers bit rates ranging from 2 Mbps to 1 Gbps. So if you are buying a 5-Mbps service, the connection may operate at 10 Mbps or even 100 Mbps, but your forwarding rate is limited to 5 Mbps. Those services require a CPE box called a multiservice provisioning platform, and in some cases the carrier will also provide a small LAN switch to sit between the customer's network and the carrier's service interface (Figure 2b)

Each carrier offers a range of network features that address reliability and QOS. While SONET can provide for failover reliability at Layer 1, several reliability mechanisms can be built into a Layer 2 service. Most WAE networks use Spanning Tree (IEEE 802.1d) or the accelerated Rapid Spanning Tree (IEEE 802.1w) to ensure backbone reliability. Carriers have found performance problems even with the Rapid Spanning Tree (see "New Battles to Fight").

The other reliability option starting to get some play is Resilient Packet Ring (RPR), as described in IEEE 802.17. RPR is a Layer 2 protocol that allows LAN switches to be arranged in a ring, with the ability to restore service within 50 msec of a fiber break. Users evaluating wide area Ethernet must determine what recovery mechanisms will be used, what parts of the service are protected (e.g. access, backbone, inter-carrier connections), and what type of restoration interval they can expect for each type of failure.

As in all network arrangements, it is important to look at features like facility diversity to protect against single points of failure (i.e., "backhoe cable fade").

As with MPLS-based VPN services, Ethernet carriers can provide QOS to address the requirements for voice and video services. Using the priority indication specified in IEEE 802.1p, traffic can be classified into two or more service levels (e.g., Gold, Silver, Bronze) and the carriers define different service level agreements (SLAs) regarding delay, loss and possibly jitter for each level. Of course, those SLAs are available only for Layer 2-based E-LAN services.

Ethernet carriers can provide QOS

New Battles To Fight

Nortel's John Hawkins, marketing manager-carrier Ethernet, has had a hands-on role in expanding Ethernet from a LAN technology to the WAN. While he sees access as a major impediment to the growth of Ethernet services, other elements will also have to fall into place. Fortunately, the wheels are already turning.

One area he cites is the need for an inter-carrier connection for Ethernet services. The Metro Ethernet Forum (MEF) is working on just such an Ethernet Network-to-Network Interface (ENNI). With an ENNI, a customer could use a metro area Ethernet service within a city, and use that same service to access a different wide area Ethernet carrier. Of course, we had the same requirement with frame relay, and the carrier NNI never materialized.

Other problems deal with scalability. Within a carrier's backbone, customer networks are separated by identifying each as a unique virtual LAN. The carrier's LAN switches will only deliver a customer's traffic to ports defined in the same VLAN.

Customers also use internal VLANs, so the first problem is to figure out how to use VLANs to separate customer networks and also maintain the customer's VLAN identifiers. This "Q in Q" problem (as in the 802.1q standard) was solved with the Provider Bridging standard (IEEE 802.1ad), which inserts a second VLAN identifier when the frame enters the carrier's network. However, there's another challenge: VLANs' 12-bit identifiers limit the carrier to supporting 4,096 customers.

To address the scalability issue, the Provider Backbone Bridging initiative was initiated in November 2004. Rather than a "Q in Q" solution, the new standard will use "MAC-in-MAC." That is, another MAC address will be inserted in front of the frame, allowing millions of possible networks.

Another problem came from the Spanning Tree and Rapid Spanning Tree (RST) bridging protocols. While RST greatly reduced reconfiguration time in customer networks, it did not have the same effect in Ethernet carrier backbones. The problem was the time required to distribute VLAN maps to all network points.

The IEEE is now developing a Multiple Registration Protocol (MRP); the protocol is specified in IEEE 802.1ak. MRP allows participants to register attributes with other participants in a bridged network. Configurations are defined to pre-register VLAN associations. As the VLAN definitions are pre-registered, MVRP will allow the network to recover more rapidly, because the VLAN information will not have to be distributed after the failure.

Converting a LAN technology to wide area applications is not simply a matter of turning the key and driving it off the lot. However, it does appear that the EFM committees are coming up with solutions as they are needed□

Carriers aren't pushing Ethernet on copper as a widespread solution

There is no "standard" configuration for a wide area Ethernet service. The carriers are feeling their way along. This high degree of customization is an attribute of a "young" service, and will have to change if these services are to develop into multi-billion dollar revenue generators.

The Achilles Heel

All observers recognize access as the single greatest obstacle to the expansion of the WAE market; however, there is considerable difference of opinion regarding the appropriate solution. The vast majority of current services require fiber access to connect the customer to the network. According to Vertical Systems' Dunne, only 10 percent of large commercial buildings (i.e., BSBs or "Big Shiny Buildings") are fiber-connected today.

The ILECs own most of those fiber connections, which gives them the inside track on capitalizing on the Ethernet service expansion. The FCC's Triennial Review, announced in October 2004, stated that the ILECs are not required to share those fiber connections on the same basis as is required for copper loops. SBC referenced that ruling as one of the key factors in its decision to accelerate deployment plans for the Project Lightspeed fiber-to-the-neighborhood (FTTN) initiative that it recently announced.

While Dunne and many equipment manufacturers see a solution in copper access technologies, the carriers disagree. Bob Walter, SBC's executive director for metro data services, notes that the RBOC has had little trouble deploying fiber in most locations where Ethernet services were required. Currently, those connections are almost exclusively point-to-point fiber links carried over SONET rings, though SBC is looking into passive optical network (PON) solutions. According to Walter, SBC is currently looking at copper access, particularly VDSL, only for residential customers.

AT&T's Klapman tends to agree. AT&T's Local Network Services arm has fiber links into roughly 8,000 commercial buildings, and point-to-point fiber is clearly their access method of choice. Klapman cites the cost of leasing copper loops as a big part of that decision, though he notes they are monitoring developments on the copper front.

In the meantime, he points to AT&T's current WiMAX trials in Atlanta and Middletown, NJ. WiMAX's ability to deliver a shared high-capacity radio channel over a range up to 5 miles (or greater distances at lower data rates) could make it the second choice after fiber. Further, the WiMAX protocol includes QOS that could tie into the backbone's QOS capabilities.

Certainly copper access could expand availability to smaller customers, and could open the MAE market to carriers who lack the financial base to deploy fiber all the way to the customer. However, all the copper connections being described are point-to-point arrangements (see "Ethernet on Copper?"). With no redundant path to the network, copper access would be a far more vulnerable configuration than a ring-configured SONET access.

The other option would be to partner with other LECs to increase the network footprint. However, the ILECs have the lion's share of the fiber access connections, and as a rule they are not disposed to sharing.

Ethernet On Copper?

The Ethernet First Mile (EFM) committees are looking at three sets of options for Ethernet service access:

- **EFM Fiber**—Point-to-point fiber access
- **EFM PON**—Passive optical network access
- **EFM Copper**—Point-to-point wire pair access

Today, the vast majority of access connections use point-to-point fiber links, typically carried over SONET rings. PON technology is getting its first major test in conjunction with the ILECs' fiber-to-the-premises (FTTP) rollouts (see *BCR*, March 2005, pp. 32–37).

For copper access, the EFM committees are planning two Ethernet access connections: 10BaseTS and 2BaseTL. The 10BaseTS is a short-range interface based on very high speed DSL (VDSL). The goal is to deliver 50 Mbps over 1,000 feet or 10 Mbps over 3,000 feet. Like VDSL, 10BaseTS is targeted at distributing services from a fiber terminal located in the basement of a multi-tenant building.

As most customers will be farther than 3,000 feet from the central office, the 2BaseTL interface will be more critical. Based in the ITU's G.shdsl technology (G.991.2), 2BaseTL will use a trellis-coded pulse amplitude modulation (TC-PAM) system to deliver a symmetrical 2-Mbps data rate over 9,000 feet. Current products are achieving 5.7 Mbps over 7,000 feet, according to Matt Squire, CTO of Hatteras Networks, which makes Ethernet-over-copper equipment.

To support higher bit rates, the plan is to bond multiple 2BaseTL systems using inverse multiplexers. The inverse multiplexing will be designed so that if one link fails, the interface will continue to operate, though at a lower bit rate. The real vulnerability is that the copper interfaces are non-redundant and point-to-point. If continuous access is critical, the user may have to plan for multiple redundant connections, or install an ADSL service as backup□

Conclusion

Clearly the market for Ethernet services is growing substantially, though from a very small base. While current growth rates are impressive, many analysts and observers anticipate "hockey stick" growth curves like we saw with frame relay in the mid-1990s.

Clearly, Ethernet's gains will come at the expense of other technologies, particularly frame relay and ATM. Ethernet's role in the overall services mix is still being defined. Internet access and metro area site-to-site connectivity are the stars today, but the interplay with Internet VPN solutions remains the big question.

Local ATM should eventually fade out as Ethernet becomes the technology of choice. If the carriers had to do it over again, it is unlikely they would have built their ADSL networks on top of ATM. As Ethernet technology takes off in the local network, we can expect to see those ATM-based ADSL connections swept along in the Ethernet tide□

Ethernet's role in the overall service mix is still being defined

Companies Mentioned In This Article
AT&T (www.att.com)
Hatteras Networks (www.hatterasnetworks.com)
Nortel (www.nortel.com)
SBC (www.sbc.com)
Verizon (www.verizon.com)
Yipes (www.yipes.com)

Managing With Ethernet

Besides developing a wider range of access technologies, the Ethernet First Mile committees' other top priority is to develop management capabilities. Because Ethernet evolved in the LAN, virtually no management capabilities were built into the original interface. In a carrier environment however, we will need mechanisms to monitor and test customer services to meet SLA objectives. In EFM, those management capabilities are referred to under the general heading of *Operations, Administration and Maintenance* (OAM).

Hatteras Networks CTO Matt Squire has served on the OAM committee, and he sees loopback and alarm indication capabilities as the most important thus far. These would allow a carrier to recognize a failure and use the same type of circuit confirmation that we have traditionally employed in carrier environments. In an Ethernet interface however, the loopback is done at the frame level rather than the bit level.

The OAM standards define a range of monitoring capabilities as well, including: Errored Symbol Period (errored symbols per second); Errored Frames (errored frames per second encountered based on CRC failures); Errored Frame Period (errored frames per N frames); and Errored Frame Seconds Summary (errored secs per M seconds).

The management philosophy is different from traditional carrier approaches, where an error detection mechanism is built in at Layer 1 (e.g. the CRC-verification in Extended Superframe Format on DS1 channels or Bit Interleaved Parity on SONET systems). In Ethernet, the basic verification comes from re-computing the frame check sequence at the end of each Ethernet frame.

According to Squire, the initial focus with OAM has been to address a single link (e.g. the connection between the customer and the network). The next major step will be to expand that management view end-to-end□