# Sky's the limit with Ethernet over fiber

*Why communities seeking a local broadband subscriber network should implement Ethernet over fiber.*

OCTAVIO MORALES AND CHRIS SETTY, World Wide Packets

By bringing together the limitless bandwidth capabilities of a fiber infrastructure with the affordability and performance of Gigabit Ethernet, "Ethernet-everywhere" networks will eliminate the constant bandwidth challenges now facing content providers. With technological barriers removed, providers will be able to focus on the development of advanced applications that will merge "wired" technology with everyday life.

The choice is inevitable. To deliver the broadband content and services that subscribers will demand, a fiber-to-the-subscriber (FTTS) infrastructure is needed. In fact, rural customers are driving the deployment of an all-fiber infrastructure. Garry MacCormick of Rye Telephone Co. (Colorado City, CO) points out that "it's a much more sensible business proposition to simply install fiber in the first place...the time will come when copper will no longer offer the benefits already available with an FTTH [fiber-to-the-home] installation (see *Lightwave*, June 2000, page 1)."

Service providers and infrastructure providers alike know they will eventually have to tear out their copper cable plants and replace them with fiber. As cost of installation becomes more competitive, it makes much more sense to install fiber from the beginning.

However, now that people are seeing the benefits of FTTS installations, some still debate which transport technology to use as the delivery vehicle. Just as MacCormick argues that it is much more sensible to use an infrastructure that is futureproof, the delivery technology should also be futureproof. Technologies such as passive optical networks (PONs), which have been proposed to increase Internet access speeds, may be completely inadequate for accommodating the future of broadband access. On the other hand, Gigabit Ethernet, in conjunction with fiber-optic media, offers the bandwidth needed to ensure an extended infrastructure lifecycle and provide content developers with a "sky's the limit" environment.

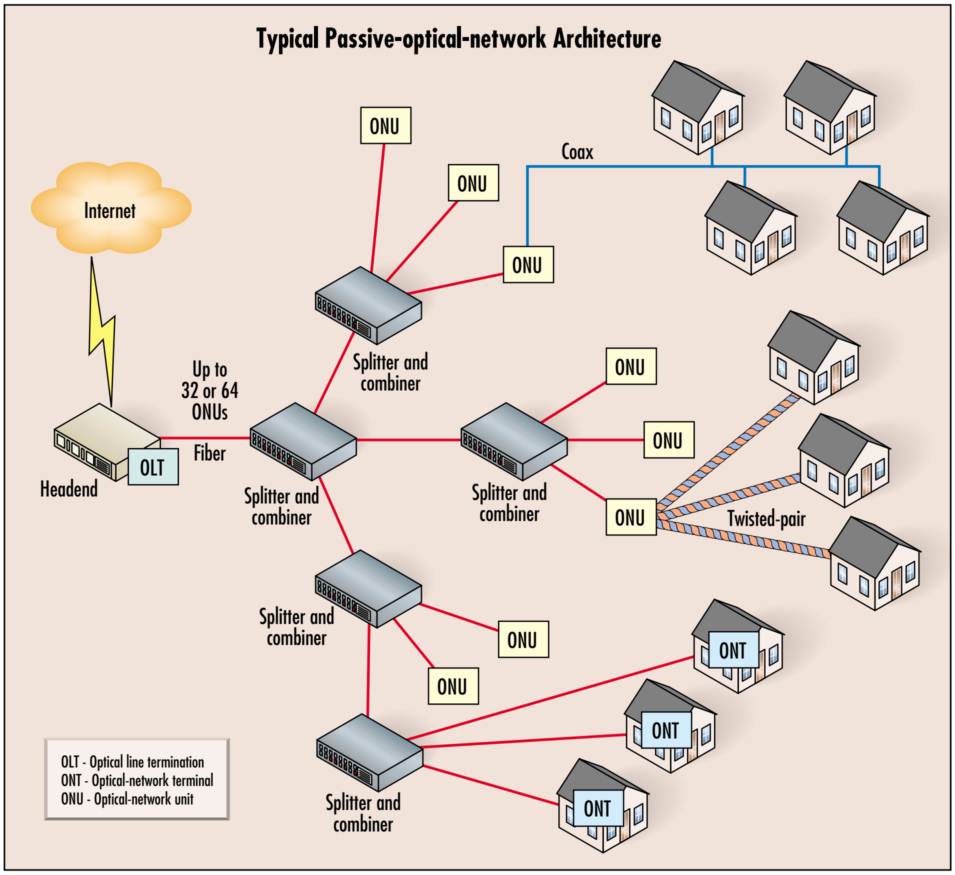


**Figure 1.** A typical passive-optical-network architecture supports multiple community distribution points where multiple lines are fanned out to subscribers. This architecture not only adds to network complexity but also effectively reduces the bandwidth available to each subscriber.

With anticipated residential broadband applications, the most significant of which is bandwidth-hungry, entertainment content, currently deployed technologies cannot deliver and will be quickly outgrown, just as T1s in the enterprise market have been.

In an effort to provide additional value and leverage their copper cable plants, telephone companies and cable-TV providers have suggested a number of "high-speed" access solutions. Integrated Services Digital Network was offered in the mid-1990s, but the telephone companies soon found their switched voice networks saturated with data traffic. That opened the door for DSL, which essentially modulates binary data into sonic frequencies above what humans can detect, allowing data to "ride" the phone lines alongside an active voice transmission undetected by the caller. But there are limitations, most notably, distance from the telephone company's central office (CO).

Faced with Federal Communications Commission-mandated competition as well as an increasing number of customers switching to digital satellite service, cable multiple service operators (MSOs) quickly adopted value-added services such as broadband Internet access to capture and retain their customer base. These systems consist of a headend device located at the MSO and a cable modem located on the customer premises. The cable modem provides an Ethernet port for connectivity to the subscriber's PC or network. While cable modems are faster than typical dial-up access speeds, the total available bandwidth is shared among all users on the cable network. Many times, users on the network number 1,000 or more. As the total number of cable modem users on a network increases, the per-subscriber bandwidth narrows drastically, rendering interactive applications such as gaming and video conferencing unusable.

These solutions represent copper-based, "band-aid" fixes. Simply stated, copper cable plants will present a brick wall. With new Internet appliances, video on demand (VOD), and other communication services under development, using these "high-speed" technologies would be like trying to suck a thick, rich, chocolate shake through a coffee stirrer. The new breed of packetized media—standard and high-definition television (HDTV), voice over Internet Protocol (VoIP), VOD, video telecommunications, etc.—will need to go well beyond what today's "broadband" technologies can provide.

But where does this leave other contemporary delivery technologies?

### Wireless local area networks

Wireless LANs represent the growing popularity of using radio frequencies to deploy a non-terrestrial network. Like many other technologies, the wireless LAN, developed by the IEEE 802.11 committee, has undergone some bandwidth evolutions: from 1 Mbit/sec to 2 Mbits/sec and then to 11 Mbits/sec. However, available bandwidth is determined by distance from the access point (the wireless hub) and the number of users on a particular frequency.

Since this technology allows large amounts of users to access resources without being plugged into a wired network, obvious applications include buildings and areas where wiring or retrofitting is not a possibility or where it is not economically feasible. It is certainly possible that a community in which fiber cannot be laid to each home could implement a wireless community network. Wireless access-point channels could also be spaced to allow for wireless roaming within an area.

In addition to communication in and out of the home, 802.11 wireless LANs can also be an inexpensive means of connecting in-home resources such as computers, printers, Internet terminals, and IP appliances. In fact, Apple Computer has even included 802.11 slots in several of its products, which allow consumers to dial the Internet or access the network from anywhere in the house at 11 Mbits/sec. Using wireless technology offers an option to retrofitting an existing unwired home with Category 5+ wiring.

With all the positive features of 802.11, there is one major downside: bandwidth scalability. While 11 Mbits/sec might be considered enough bandwidth for a considerable amount of time, keep in mind that this bandwidth is shared among all users on a particular access point. Wireless signals are received and handed off to a local Ethernet network. Competing service providers presenting choice and cost options to subscribers can also populate the local network, reaching any subscriber on the network. This model, however, is very complex to deploy in a shared bandwidth environment. Wireless LAN will likely become a support technology to fill in the gaps and could be viewed as a complementary solution to FTTS deployments.

### Asynchronous Transfer Mode

ATM emerged in the early-to-mid-1990s as a telephone-company-grown technology that was being pushed to the enterprise LAN. Simply explained, ATM is a Layer 2 technology that establishes dedicated, connection-oriented, virtual circuits (VCs) across the network. VCs can be manually configured using permanent virtual circuits (PVCs) or setup and torn down dynamically, as needed, using switched virtual circuits (SVCs). Once a connection is established, data packets are segmented into 53-byte cells. These cells are transmitted across the VC to the egress ATM switch, reassembled into the original packet and delivered to the intended destination. This process is called segmentation and reassembly (SAR).

Since ATM is a protocol unto itself and switches at Layer 2, it is unbiased as to what can be sent over it. Data, voice, video, or virtually anything else can be sent across an ATM network. All data packets are converted into ATM cells for transport, then reassembled for end-node delivery. Bandwidth ranges from 25

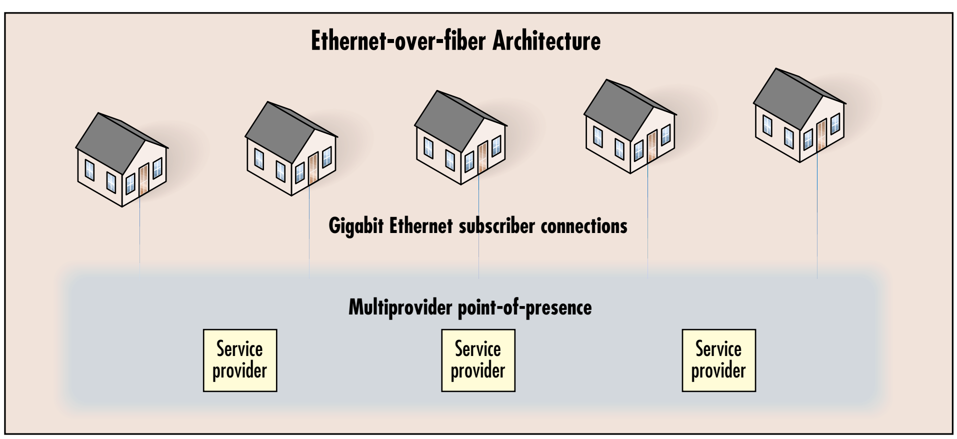


**Figure 2.** A Gigabit Ethernet-over-fiber solution provides dedicated upstream and downstream connections to each subscriber—up to 1 Gbit/sec to each premises.
The connections are therefore scalable for multiple services and future technologies.

Mbits/sec to 2.48 Gbits/sec, with backbones now being planned or implemented using OC-192 (10 Gbits/sec).

While ATM has found a niche in the telecommunications carrier space, the enterprise has refused to embrace it. One of ATM's best arguments is its ability to implement end-to-end quality of service (QoS) throughout the network. The benefits are simply not significant enough to balance out ATM's faults. While QoS is important, other technologies such as Ethernet also make provisions for QoS.

High price tags, difficult deployments, the need for sophisticated management applications, and a steep learning curve for implementation have plagued ATM. By necessity, as a technology and its implementation are brought closer to the subscriber, complexity must decrease. ATM is therefore very unsuitable to customer-premises connections.

In addition, since the majority of home and business networks are, and will continue to be, built using Ethernet, complicated conversion technologies will be needed. ATM requires LAN segments such as Ethernet to be "emulated" over the ATM backbone, creating additional complexity, cost, and points of failure in the network. When the ratio of cost to speed and features is considered, ATM does not present a compelling solution in the deployment of local broadband services.

### Passive optical networks

PONs are access networks in which fiber trunks are fed toward end-points and split into multipoint trees along the way until reaching a termination of the fiber run. This termination point could be at a distribution node in a fiber-to-the-curb installation or at a home in an FTTH installation.

A PON network is deemed "passive" because the physical connections between devices consist only of passive components such as optical fibers, connectors, splitters, combiners, and splice points. It should be noted, however, that while a PON's passive nature sounds unique, in any FTTH model where fiber directly connects the customer premises to a headend system, the transport of the signal between both points is passive by the very nature of fiber-optic technology.

In a typical PON installation, one branch from the headend, or optical line termination, typically supports up to 32 community distribution points, or optical-network units (ONUs). From there, multiple lines are fanned out to subscribers' premises. Depending on the type of installation, these subscriber connections could be fiber or even copper runs (see Figure 1).

Where coaxial cable or twisted-pair wire is used to distribute services, the optical signal has to be converted into an electrical signal for delivery over the copper wire feeding the home. In addition to signal conversion, a multiplexing algorithm is also needed to transport multiple subscriber signals on a single fiber. All of this equals one thing: complexity that rapidly increases as the network scales.

Bandwidth in a PON system is typically OC-3 (155 Mbits/sec) or OC-12 (622 Mbits/sec). Each fiber trunk is split to service the community distribution points; it effectively reduces the available bandwidth per subscriber. Given an OC-12 trunk and an FTTH topology, a PON network with 32 community distribution points reduces the bandwidth available per subscriber location to approximately 19.5 Mbits/sec (622 Mbits/sec/32). In the case of fiber-to-the-curb (FTTC), the bandwidth per subscriber is further reduced.

Even if the bandwidth were scaled to OC-192, a cluster of 32 distribution points for a typical neighborhood would only have 312 Mbits/sec to distribute among 100 homes. In an FTTC environment, this amount would equal approximately 3.12 Mbits/sec of effective bandwidth to each subscriber. Obviously, as the number of homes increases, the effective bandwidth per subscriber decreases.

Even though PON seems to be generating a lot of support, the fact remains that with its blend of bandwidth limitations, expensive headend systems and overly complex operation, PONs simply don't allow for the kind of futureproof broadband access required to deliver the content of today, let alone tomorrow.

### Ethernet to the rescue

Ethernet has enjoyed phenomenal success in the enterprise LAN and has also emerged as the choice for metropolitan area networks (MANs). Today, well over 90% of deployed networks are based on this solid, standardized technology. It has grown from a shared, 10-Mbit/sec technology into a switched full-duplex technology, providing dedicated bandwidth to each subscriber of up to a full gigabit of throughput in each direction. Thousands of Ethernet devices are available to handle everything from small home-based networks to Fortune 500 backbones and regional MANs. Worldwide shipments of Ethernet devices measure in the tens of millions of interfaces each year.

Ethernet is ideal for community networks deployed over FTTS infrastructures. Since fiber reaches all the way to the subscriber, it is possible to provide a home with up to a full gigabit of bandwidth. While that may not be necessary today, it does provide a significant amount of scalability for the future. Ethernet can deliver on the promise of high bandwidth today without the costly headend and customer gear required in the installation of other technologies.

With a full gigabit of capacity available per subscriber, providers could even offer several service plans, each providing a higher bandwidth rate. Changing subscriber preferences could occur in seconds rather than days or weeks.

**Ethernet can deliver on the promise of high bandwidth today without the costly headend and customer gear required in the installation of other technologies.**

In an Ethernet FTTS network, a collocation of service providers can exist, providing specialized, high-quality service. Subscribers could even choose which provider will deliver each of the services they want. A subscriber could receive all their services from one provider or they might sign up for digital-video broadcasts from one provider and high-speed Internet access from another. While not the only possible model, this competitive service-provider scenario ensures that quality and cost will be positively influenced.

Gigabit Ethernet is capable of transporting virtually any application from a large number of end-stations. With 10-Gigabit Ethernet in the standards process now, capacity is guaranteed for the future as well as providing the ability to carry Ethernet frames from the subscriber to the MAN and out across the WAN—an end-to-end Ethernet solution. These advantages, coupled with practically universal support across the industry, make Ethernet a standout in community networking technology.

**But what about QoS?**

Ethernet's critics purport that the technology is not built with inherent features such as QoS, class of service (CoS), or redundancy mechanisms. The fact is standards that provide QoS have already been developed and are in use in today's Ethernet networks. With the large amount of industry support Ethernet enjoys comes an incredible selection of device-specific technologies that add these functions, and more, to Ethernet communications.

Many devices now incorporate comprehensive QoS measures that allow packets to be classified, prioritized, policed, queued, etc., then forwarded accordingly. That allows certain communications to be handled differently compared to others; for example, a packetized MPEG-2 stream can be forwarded along a higher-bandwidth, lower-latency link than typical Web-surfing traffic. Additionally, one of the inherent advantages of high-bandwidth Ethernet is that when the bandwidth pipeline is big enough, all traffic gets through on time. Complicated management is not needed once bottlenecks are removed.

The ideal way to extend a technology is to keep the technology itself simple, while introducing high-level features in the hardware that transport the data. Inundating the technology with complex protocols, algorithms, signaling and features only presents a difficult-to-manage, high-price-tag monster (did someone say ATM?).

Another popular anti-Ethernet rant is that Ethernet is a shared, broadcast-oriented, contention-based technology. While that was true a decade ago, with the advent of full-duplex communications, Ethernet has evolved into a dedicated-bandwidth technology. A single user can receive up to 1 Gbit/sec of full-duplex bandwidth for exclusive use. Faced with an unbeatable combination of unsurpassed performance, advanced features, simple deployments and affordable pricing, Ethernet's foes always fall back on the arguments of yesterday.

Ethernet's performance, scalability, acceptance and support, together with advances in hardware, make it a very viable community networking technology. Concerns about cost, interoperability, scalability, and ease of management simply aren't warranted in the Ethernet-everywhere scenario.

Many industry consultants echo this opinion. Several university research groups that advise local and national governments on their data infrastructures, give standard advice to their clients to abandon their entire existing copper infrastructure and build a new fiber infrastructure. They feel that Gigabit Ethernet has a great potential, especially for metropolitan area networks and the last mile.

As shown in Figure 2, a full 1-Gbit/sec access rate is provided to each premises, meaning services are not limited to merely high-speed Internet but could include voice, video, and other services as well. In addition to transporting Internet-bound and sourced data at rates up to one thousand times what DSL and cable modems can offer, Gigabit Ethernet is able to simultaneously deliver multiple HDTV channels and toll-quality voice calls to each subscriber. Gigabit Ethernet has enough bandwidth to accommodate and deploy future services and applications over the same infrastructure.

When all the technologies are compared across the board, Gigabit Ethernet over an FTTS network has distinct advantages for combining high bandwidth with low cost, simple installation and maintenance, widespread availability of devices and software, scalability, and new standards. Furthermore, it enables a future-proof network that can accommodate emerging technologies. ❑

*Octavio Morales is vice president of marketing and Chris Setty is product manager at World Wide Packets (Spokane, WA).*

Contact Information
World Wide Packets
115 N Sullivan Road
PO Box 950
Veradale WA 99037

Tel. 509.242.9000
Fax 509.242.9001
www.worldwidepackets.com