Pipeline Knowledge is Power.

Ethernet as a Carrier Service
By Barbara Lancatster

No wonder Ethernet is popular. It's easy to get along with. It talks to everything – Windows, Linux, Unix, Mac, FreeBSD. It's simple to implement: just about anyone can sling together an Ethernet LAN and it will work, mostly.

Ethernet is technically *forgiving*. Sketch a LAN on the back of an envelope, plug it all together and it will usually work. Yet, certainly for a large network, the designer should understand the main sources of traffic and segment the network to optimize flows. But even if you don't do that, the data will still get there. Thanks to Ethernet's forgivingness, LAN designers do not, generally, need to calculate exact bandwidth requirements in every segment. LAN administrators in many companies don't rely on complex management systems to tell them if there's a capacity problem rather they respond to user complaints of network sluggishness and address the congestion problem by throwing more hardware at it. You can throw in some class-of-service prioritization and improve quality for VoIP and video simply by upgrading the switches. And bit-rate conversion comes built-in, unlike in traditional phone networks. There is no thought required: people routinely interconnect devices running at a variety of data rates over an Ethernet LAN.

As a result of Ethernet's friendliness, every IT department in every enterprise depends on Ethernet. And due to its familiarity and ubiquity, costs of Ethernet switches and routers have steadily decreased; at the same time as performance and reliability have increased. Ethernet on the LAN is a success.

By contrast, ATM is hard work. ATM is a telco-oriented technology that was designed to do the same sort of job in the WAN as Ethernet does in the LAN, and at one time, ATM was the leading candidate for a unifying technology across WANs and LANs. In the world's public telecommunications networks, ATM was a success. But no one now is seriously talking about ATM migrating into the LAN. It is not that ATM is an inferior technology. Far from it – it is versatile and full of features and it enables management of QoS for multiple types of traffic. But ATM is a complex technology and network architects and engineers need to know what they are doing with ATM to make it work well. As a highly specialist technology aimed at telcos, ATM hardware prices have not experienced the big cost reductions achieved by the deployment of Ethernet as a general purpose technology across millions of corporations, small businesses, and homes.

The success of Ethernet in the LAN led people to think about ways of realizing the same sort of benefits (especially low prices) in the wider network. They understood that if Ethernet could be used across cheap high-capacity fiber that could be privately owned, or rented, or shared (as a municipal or condominium network) then these links could be easily integrated with their enterprise networks. At the same time this approach would reduce reliance on traditional carriers.

The implications of this were not lost on strategists in telcos. If anyone could deploy Ethernet, and Ethernet was everywhere, and any content or application (including voice) could be delivered via IP over Ethernet, then what role would there be for the "common carrier." If traditional telcos chose not to provide Ethernet-based services then they would eventually be marginalized. Add this to the potential of VoIP and IP in general for relegation of telcos to being simply bit-carriers and we realize that being a telco strategist must be a really interesting job these days.

So clearly network service providers had to start delivering Ethernet to their enterprise customers, and many of them are doing so, as described elsewhere in this edition of *Pipeline*.

When carriers started to think about deploying Ethernet in the wider network, LAN folks everywhere approved. Their familiar, forgiving and inexpensive friend would escape the walled garden of the LAN and get to play in the big world of the WAN. LAN people would no longer need to understand arcane technologies like SDH/SONET, Frame Relay and ATM. With Ethernet everywhere, and IP running over it, the global network would be elegantly simple and easy to manage. There would be one unifying Layer 2 technology spanning the globe.

At the same time, telco folks allowed themselves some skepticism. It's their job to be skeptical about the deployment of new technologies. They have a tradition of delivering their customers the highest possible standards of quality, reliability and security. Could a technology conceived and designed for the LAN environment possibly make it as a WAN technology? Ethernet needed to meet the WAN challenges of QoS, reliability, and security while being able to handle WAN volumes of traffic, and WAN geographical distances. Clearly, these are not trivial challenges. Nevertheless, driven by the prospect of even larger global markets for Ethernet technology, the network technology community has come up with some answers.

Ethernet uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Isn't this technique inherently unscalable to the sort of traffic levels used on carrier networks?

Apparently not. Ethernet networks at any level need to be designed with sufficient overhead to minimize congestion. However, higher bit rates and the ubiquitous use of Ethernet switches reduces the number of simultaneous messages on any one segment, and higher bit rates throughout enable more transmissions to take place without collisions. It's scalable.

Ethernet is intrinsically a short-distance protocol. LAN distances are trivial, but telcos need to think about distances in transcontinental or even global terms. How can Ethernet be adapted to long distances and still be Ethernet?

The distance challenge is not related to the Layer 2 data link design of Ethernet, but to the various choices that had been made for Ethernet Layer 1 physical connectivity, which traditionally looked for pragmatic solutions for LANs. However pragmatic solutions for metro-Ethernet and long-haul networks exist too, now that we've started

putting them together. In the WAN arena we can now use 1000BaseLX or LN GBICs (Gigabit Interface Cards) over certain legacy 1310nm single mode fiber to reach distances of up to 10km, further with regeneration. (The approach is similar to the LAN standard 1000BaseSX, in that it sends and receives Ethernet directly over dark fiber, but SX is distance-limited because it uses cheaper multi-mode fiber.) New implementations can benefit from new generations of fiber, for example 1000BaseLX can be made to work over 70km or more using 1510nm NZ-DSF single mode fiber. DWDM (Dense Wave Division Multiplexing) techniques can take us even further: Gig-E or even 10Gig-E can be transmitted up to 500km before regeneration is necessary. All of these developments provide a range of technologies suitable for deploying Ethernet at metro, regional, and national levels.

Can carrier Ethernet ever be as secure as traditional telecoms links?

Protection of private customer traffic can be enhanced by secure VLANs that enable multiple customers (or applications) to exist on one Ethernet network. This is now an established standard (IEEE 801.Q) for doing this. For large enterprises, for government, or for particularly sensitive applications, telcos can always provide physically separated capacity (separate fibers or wavelengths) which is how traditional networks kept customers apart: even if this uses up some capacity it may be the most cost-effective way to do it.

Can a LAN technology such as Ethernet be delivered with "carrier-class" availability and reliability?

Probably. It depends how much you want to pay, but there seems to be no technical barrier to delivering carrier-class Ethernet. One approach might be to use IEEE 802.17, more commonly referred to as Resilient Packet Ring (RPR) technology. RPR promises 50ms network restoration: if a link fails, the traffic can be rerouted in less than 50ms. To put this in perspective, continuous latency of less than 150ms in a voice call is considered acceptable: a 50ms break in transmission would hardly be noticed.

Can carriers deliver QoS based services over Ethernet? Can carriers deliver Ethernet as a network service with a guaranteed QoS?

The answer to both questions is "yes." Ethernet allows class-of service prioritization, ensuring that QoS-sensitive services (such as voice) are given priority over an Ethernet link, reducing levels of latency and jitter. Given adequate bandwidth, we can engineer any link to ensure a certain level of QoS for any type of traffic – but, it must be noted, this level of QoS is statistical rather than absolute. This is not necessarily a bad thing and reflects the way QoS has traditionally been delivered by telcos throughout the history of telecommunications. However, it may be reasonable for carriers to respond to the requests of some customers for tougher SLAs that contain *guarantees* of certain levels of quality. Some vendors now are able to deliver "hard QoS" for Ethernet which essentially reserves specific capacity for certain services

carried over the Ethernet (which is not quite the same as QoS for the Ethernet service itself).

However, there is a trade-off: hard QoS means softer QoS for customers who are not paying for rigorous SLAs. As for QoS for the Ethernet service itself, the answer is the same as the answer to the question of availability and reliability: pre-assigned dedicated bandwidth, available whether it is in use or not, costs money. What is important is that the ability to deliver QoS appears to no longer be a big barrier to deploying Ethernet in the WAN, and that carriers may be able to offer customers QoS choices depending on their needs and budgets.

A quick survey of today's state of the Ethernet art indicates that the technical barriers to Ethernet as a carrier service are being steadily knocked down. At this stage, it seems entirely reasonable for enterprises to expect carriers to deliver WAN Ethernet that works just as well as LAN Ethernet, and can interwork with it seamlessly. At the same time as this is a revenue opportunity for carriers, it's also something of a threat: ubiquitous IP over Ethernet permits enterprises and others to build geographically large networks that can carry services that can supplant more traditional carrier data services – specifically voice and dedicated data lines. Despite this dilemma, it is clear that eventually traditional carriers will not have much option, they will deliver WAN Ethernet services, or others will. If they do it, they need to be successful, and this may be achieved by not adhering too strictly to some traditional carrier approaches.

The explosion of Ethernet in the enterprise and business market has taught us that the "best" technology isn't always the theoretically most elegant technology (token ring) or the most versatile and capable (ATM). Perfection is not required, otherwise people would have lost patience with both Windows and Linux a long time ago. Enterprises have learned to deploy equipment quickly and effectively by cutting out some of the complexities we used to assume were necessary for successful deployments. They have been able to do so because bandwidth and Ethernet equipment have both become relatively inexpensive, while the cost of expertise, and the cost of delays associated with excessive thinking have both risen.

This success, and how it has been achieved, should be instructive for service providers as they roll out Ethernet WAN deployments for their enterprise customers. In the past, network service providers had the time and resources to aim to deploy technology that was closer to perfection than many enterprises needed. Today, we need to look for a good, stable and secure deployment, but not necessarily one that approaches perfection. Carriers may need to deliver reliability and availability through redundancy and over-provisioning rather than by complex and elegant design. Enterprise customers will value visibility of performance, some level of QoS, and the ability to manage bandwidth to some extent; but they will be even happier with inexpensive Ethernet pipes so big that QoS will not be an issue and so cost-effective that it will not be worth managing bandwidth up and down except to handle significant changes in actual needs.

For many companies today, if the service requires them to think about the service every day, then it's not the service they want or need. They need consistent "good enough" performance, and they want to mostly forget it's there.

In other words, as Ethernet is deployed as a carrier service, carriers should aim to make it as friendly and effortless as Ethernet in the LAN and good enough to minimize the need for thought and intervention on the part of the customer. People in corporations want to spend their time, money, and effort on running their core businesses, and are really quite happy to let the networks run themselves. Extending the Ethernet to the WAN and making it effortless is the big contribution traditional network service providers can bring to communications-intensive enterprises. If it turns out to be anything other than cheap, friendly and effortless, people will find ways of doing it themselves.