

## The Promise of the Triple-Play

Carrier success in delivering the ‘triple-play’ of voice, video, and data is not only dependent upon the proper choice of service and content partners but also on the right network infrastructure. This network infrastructure must be capable of evolving as business and consumer needs change, as new services and applications are introduced into the marketplace, and as bandwidth needs grow. We are on the cusp of the next great evolution in subscriber connectivity – from dial, to low-speed broadband delivered over copper or MSO HFC (sometimes referred to as ‘midband’, or ‘advanced services’ and ‘first-generation broadband’ by the US Federal Communications Commission), and now to fiber, or ‘advanced broadband’, offering true broadband connectivity capable of supporting HDTV, multiple lines of voice, and Internet access bursting to 10 Mbps and beyond.

Although fiber deployment to the end-user or to the neighborhood is still in its infancy, with less than 1% of North American households and 10% of businesses connected directly, carriers are already looking to define their service architectures. With the regulatory environment clarified, LECs and municipalities are working to increase this number. The result – an advanced broadband services architecture similar to the ones already providing Europe and Asia with economic advantages (Figure 1).

These carriers are now in the planning phases for their triple-play deployments. They are selecting their access and aggregation network technologies, their video and voice server infrastructures, and their back-office systems. Ethernet is playing a major role in this selection process – a technology no longer confined to the campus of a small subset of carrier metro services. Ethernet is quickly taking the lead as the infrastructure technology of choice for these next-generation networks, in the same way that ATM and Frame Relay paved the way for the great carrier buildouts of the last decade.

This paper describes the Ethernet infrastructure in support of these next generation service networks. It looks at different last mile technologies, including active Ethernet and new forms of DSL, and contrasts them to other technologies which may offer less scalability or longevity. It presents a future-proof network core based on MPLS and VPLS. And, it looks at two actual deployments:

- UTOPIA network in Utah, an example of active Ethernet to the subscriber
- Telefonica Imagenio network, where fiber extends to the neighborhood and next-gen DSL provides advanced services to the end users.

Providing meaningful access to advanced telecommunications for all our citizens may also spell the difference between stagnation and economic revitalization. Building this infrastructure is important for all communities in this country  
**- FCC Commissioner Michael J. Copps, August 2003**




Figure 1: North America poised to follow Asia and EMEA Buildouts

Active Ethernet provides dedicated bandwidth to each end node, while PON is like a shared media network where multiple users share the same bandwidth.

## Alternatives: Active Fiber, Passive Fiber & DSL

There are multiple approaches to delivering triple play services to the end user, by fiber. One is an active last-mile architecture, sometimes called active Ethernet, providing each customer with a dedicated fiber connection to a switch at a neighborhood aggregation point. The second is based on a passive last-mile architecture, commonly referred to as a Passive Optical Network, or PON. The PON consists of a powered Optical Line Terminal (OLT) and subscriber units known as Optical Network Units (ONUs). Passive splitters distribute traffic from a OLT port to multiple downstream ONUs. Yet another option is to run fiber to the neighborhood serving node, where it then services copper-connected end users via DSL. Figure 2 presents an overview of the various last-mile alternatives, and how they meet the needs of typical triple-play advanced broadband services.




Figure 2: Service Requirements and Technology Capabilities

Conventional wisdom holds that if a single fiber strand is installed, and shared by many subscribers, it must be less expensive. In fact, the reality is that after the price of splitters (required for PON) and the price of splicing is taken into account, the cost of the outside fiber plant may actually be greater with PON than with direct fiber runs.


## Comparing last mile fiber approaches: Active Ethernet vs PON

At a high-level, the best way to look at active Ethernet and PON is to draw an analogy to a LAN. Active Ethernet provides dedicated bandwidth to each end node, while PON is like a shared media network where multiple users share the same bandwidth. In addition, there are misconceptions about the cost of fiber and the cost of hardware deployment. By lifting these misconceptions and shedding light on the limitations of PON, we show below that an active architecture is a more attractive and futureproof investment for the carrier.

Conventional wisdom holds that if a single fiber strand is installed, and shared by many subscribers, it must be less expensive. In fact, the reality is that after the price of splitters (required for PON) and the price of splicing is taken into account, the cost of the outside fiber plant may actually be greater with PON than with direct fiber runs <sup>(1)</sup>. These direct fiber runs also use newer 100BaseBX optics that require only a single fiber strand per subscriber, resulting in further savings.

The total cost of deploying and operating the Ethernet aggregation equipment is actually less than that of the PON hardware, and is more service rich in supporting the various Layer 2 and Layer 3 business and consumer services.

The next cost element is the CPE device. With an active architecture, a very simple Ethernet-based device (aka Residential Gateway) may be deployed, integrating voice and video functionality where required. Given Ethernet's economies of scale, pricing for these devices is expected to continue to drop over time. The aggregation point under the switched architecture (the Remote Terminal in Figure 3) combines the



**Figure 3: Active Ethernet Architecture**

subscriber fibers into GE (or in the future, 10GE) uplinks, and serves anywhere from 100 to 1000 homes depending upon the density of the buildout. The number of homes served by this point, and thus the amount of investment required, is a linear function of the service uptake on the network. This is different from the step function (i.e., groups of 32 subscribers per OLT port) required within a PON deployment. The aggregation point is also the point where PON Optical Line Terminal equipment would have been installed. The selection of the aggregation point is critical, balancing the cost of fiber runs vs remote power requirements.

The need to power the Ethernet hardware is a point usually used against this architecture by PON proponents. This is of course an issue, but there are multiple solutions. For example, the carrier may have a local-loop architecture where there are environmental cabinets/vaults installed closer to the neighborhoods or office parks. In high-density areas, the basement of the MDU/MTU will suffice. In some areas, a business case may be made for homerun fiber from the Central Office (CO) directly to the subscribers, with fiber splitters co-located with existing copper splitters. Here, the CO acts as the aggregation point, given that the fiber has a 10Km reach, sufficient to reach almost every subscriber served from a given CO. There is therefore no one-size-fits-all architecture for where to deploy the aggregation point. In addition, the OLT within a PON architecture is by necessity a powered device, so there is still some requirement within a PON deployment for powered remote cabinets. Note that ONUs at the subscribers within a PON deployment are also powered, in addition to the Set Top Boxes (STBs). An active architecture requires only the residential gateway that may also serve as the STB.

An HDTV broadcast occupies approximately 20 Mbps. Using the current video over DSL baseline of 2-3 settops as a guide, video alone will require 40-60 Mbps. Add an additional 6-10 Mbps for data, and the 100 Mbps available via an active architecture no longer seems like a luxury.

Another element of reliability is the fiber itself. Given that a single fiber strand within PON serves up to 32 (or even 64) subscribers, a fiber cut is more likely to cause an outage for a larger subscriber base than with an active architecture. And, any changes to the splitter architecture will force off-line all subscribers connected to that strand.

Even with these considerations, the total cost of deploying and operating the Ethernet aggregation equipment is actually less than that of the PON hardware, and is more service rich in supporting the various Layer 2 and Layer 3 business and consumer services. It also closely integrates with the provider's IP backbone and provisioning systems, and offers greater interchangeability of components that in the long term will help reduce costs. For example, a carrier may use one vendor for the core, another for the aggregation point, and a third for the customer location. In addition, within the network core, the provider may deploy the same type of hardware as that used at the aggregation layer, resulting in additional operational savings. This is not an option with the PON alternative, requiring two different types of hardware.

## Futureproofing through Flexibility

In addition to financial considerations, the long term flexibility of a given architecture must be taken into consideration. As more services are added in the future, more capacity will be required and a dedicated fiber solution allows an easier capacity upgrade. For example, many active Ethernet connections are 100Mbps Ethernet. They can be easily upgraded to GigE connections by upgrading the network equipment on both ends of the fiber, a move that quickly pays for itself through new service revenues. There is no new fiber installation required.

With an APON or BPON deployment based on the FSAN standard, today, 622 Mbps is shared between 32 subscribers for an average of < 20 Mbps per customer. Older systems offering less bandwidth (i.e., 155 Mbps) result in even less throughput. Even with 1 Gbps Ethernet-based EPON or the 2.5 Gbps GFP-based GPON, the bandwidth is still shared. So when the time comes to offer services that need more bandwidth, PON deployments come up short since the fiber in the ground cannot provide the additional capacity required.

For example, an HDTV broadcast occupies approximately 20 Mbps. Using the current video over DSL baseline of 2-3 settops as a guide, video alone will require 40-60 Mbps. Add an additional 6-10 Mbps for data, and the 100 Mbps available via an active architecture no longer seems like a luxury. This is particularly true when you consider that HDTV is already well on the way to mainstream adoption and that by 2006 there should be no NTSC broadcast without an equivalent HDTV (ATSC) signal, at least in the United States. Even without multiple HDTV feeds, if one considers the available bandwidth of 100 Mbps, experience has shown that the service providers, content providers and application developers will create and customers buy into applications that make use of this bandwidth – the jump from 10-20 Mbps to 100 Mbps will unleash this next set of applications, as was the case with the jump from narrowband to DSL and cable.

## Video over PON: A Step Backward

Separate from any bandwidth considerations, the current PON architecture is almost a step back from the road to services convergence, in that the video signal is carried as an RF signal 'out-of-band' (over a separate wavelength, much like an MSO combining video and data on different frequencies over the same coax) from the IP traffic. This is due to the lack of bandwidth available within the data signal, as described above. It also mandates conditional access in hardware, when more sophisticated software-based techniques are available. It also mandates powered amplifiers along the path, contradicting some of the reputed advantages of the PON. This divergence is counter to where the industry is heading, and creates additional

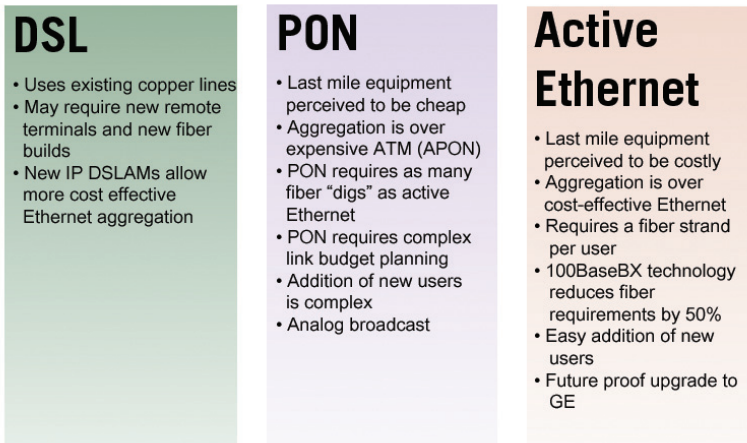
In the end, the comparison reduces to which solution is more suitable for these converged data, VoIP, and video services. With multiple services, especially in a wholesale environment where service isolation is critical, only VPLS, based on MPLS, is a viable option.

protocol and applications handoffs at both the headend and in the home, where consumers are converging on IP. It also leads to an architecture that will have to be reexamined in the future. In fact, this RF distribution is actually more complex than a pure IP network since if the carrier wishes to offer any on-demand services, control traffic will need to be converted to IP in the upstream direction at the STB.

## Copper based last mile: DSL

Where subscriber density or the funds available for investment do not justify a fiber overbuild, an upgrade of the existing DSL infrastructure is a good compromise even if it does not provide the headroom of a 100 Mbps dedicated Ethernet service. Here, the existing ADSL ATM-based DSLAMs are replaced by IP DSLAMs with GE uplinks. These new DSLAMs support ADSL2+, VDSL, or even VDSL+, offering downstream bandwidths from 12 to even 100 Mbps depending upon distance. What is more important is that these DSLAMs are deployed much closer to the end users, significantly shortening the copper loop, pushing bandwidth to the higher end of this continuum. In addition, the ATM aggregation network is replaced with Ethernet routing, now providing an infrastructure capable of supporting the bandwidth and QoS requirements of triple-play. Both SBC and BellSouth have announced plans along these lines while Verizon has announced PON plans. Figure 4 summarizes the approaches described above – Active Ethernet, PON, and DSL.

A Layer 2 architecture insulates the provider from customer complexities, and the content providers are free to manage their own IP addressing.



- DSL is a good choice for delivering triple-play without last mile changes
- Active Ethernet is better for enhanced triple play
- FTTH deployments like Utopia found PON networks the same cost as Active Ethernet, but with far less bandwidth

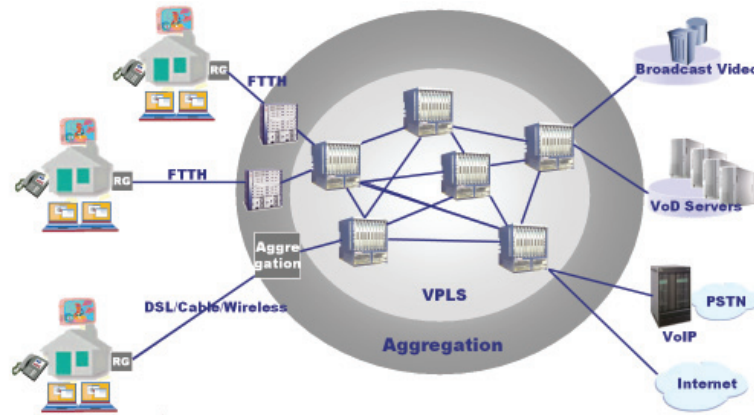
Figure 4: Last Mile Solution Comparisons

## The Core

Regardless of whether an active fiber or DSL approach is chosen, VPLS interconnects the core nodes, fanning out to VLANs within the aggregation layer. Within the industry, there is a great deal of discussion on the most suitable services architecture – VPLS, VLANs, or QinQ - for an Ethernet backbone capable of supporting converged services. These discussions focus on a number of critical areas, including services support, network scalability and reliability, standards support, application QoS, the requirement for CPE configuration, network management, and both entry and upgrade costs. In the end, the comparison reduces to which solution is more suitable for these converged data, VoIP, and video services. If the carrier just wants

to offer a basic retail Internet access service, all architectures are suitable. However, with multiple services, especially in a wholesale environment where service isolation is critical, only VPLS, based on MPLS, is a viable option.

The traffic engineering available within MPLS provides the robust end-to-end QoS that these services require. This contrasts with the VLAN approach where the three priority bits within 802.1p are just not up to this. The network is also called upon to support real-time traffic. MPLS Fast Reroute and Backup LSPs are critical, as the tens of seconds that the VLAN or QinQ Spanning Tree Protocol requires for failover is unacceptable, which is not surprising as STP wasn't designed to span the hundreds of miles required for a triple play network buildout. Figure 5 depicts the resulting reference architecture – active Ethernet or DSL to the subscriber, and an MPLS/VPLS core.



**Figure 5: Reference Architecture**

## Footnotes

(1)

- PON Splitters needed to connect OLT to the home cost around \$40/Home.
- UTOPIA cabinet's cost about \$25K or \$25 / home
- Other costs include the cost of splicing downstream, estimated by DynamicCity to be \$53 / Home
- DynamicCity sees no major difference in the cost of actual fiber of PON vs Active
- Therefore the cost of outside fiber plant is actually higher for PON vs Active

## Appendix A – UTOPIA Case Study

As part of Dynamic City's UTOPIA project in Utah, the assumptions above were translated into an actual business model. The goal was to provide almost 140 thousand subscribers with a combination of voice, video, and data services. Surveys taken in advance of the project as well as evidence from other metro fiber projects determined that almost 50% of the residential subscribers would sign up to all three services. 92% of the consumers would sign up for voice, 61% for video, and 77% for data. About half of the business customers would opt for voice, while over ¾ would sign up for data. As was expected, with little experience as to advantages of a business video service, uptake is low. However, this is expected to grow over time.

A VPLS based Layer 2 VPN solution was selected since the network is designed to provide open access to any mix of content and application providers.

The projected makeup of the fiber plant was pretty typical, with about ½ aerial and ½ buried. Aerial construction would cost \$34K per mile, and buried \$102K per mile. Almost all subscribers would be served by newly laid fiber. The topology of the project- residential and medium-density cities – lends itself to these numbers. Other environments, such as dense city cores, would result in higher construction costs, but an option here would be to lease fiber capacity already in the ground, avoiding the need for construction. For this project, Dynamic City determined that the average density for the fiber plant was 96 homes passed per mile, ranging from 29 to 137. The financial viability of the project is very dependent upon the mix of aerial vs buried and the density. For example, an area with a 50/50 mix of fiber but medium density may be marginal, whereas a low-density area with all aerial buildout may be quite appealing. In the final analysis, the cost of the fiber and hardware infrastructure ranged from \$809 to \$3682 per subscriber, with an average of \$1171. The subscriber-dependent cost from the curb to house or business was \$1107, an investment only made if the customer signed on for the service.

The next step focused on the network hardware, given that the fiber plant is a long-term investment and will not need replacement. An analysis of each network element that took into consideration increases in required network capacity, general hardware obsolescence and failures was performed. The core equipment was rated with an MTBF of 12+ years, a 7-11 year obsolescence, and a planned replacement date of 7 years. The aggregation layer was rated for a 15+ year MTBF, a 7-9 year obsolescence, and the same 7 year replacement interval. Finally, CPE was rated for a 7-10 year MTBF, but only a 3-5 year obsolescence, resulting in an expected replacement at 5 years. As mentioned earlier, if subscriber services and required bandwidth grow to an extent that new hardware is required, the additional revenues would finance earlier replacement.

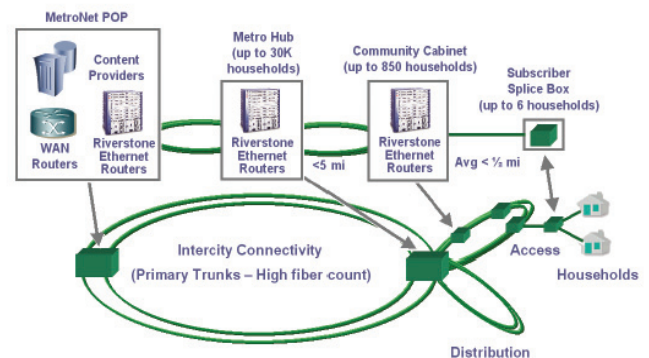


Figure 6: UTOPIA Architecture

The network is divided into five layers: Provider Access, Core, Distribution, Access, and Subscriber. The Provider Access and Core hardware is at the MetroNet (UTOPIA) POP. Here, other data, voice, and video providers connect to the network to provide services to the end users. Ethernet routers in the core each support up to 20 Metro Hubs. See Figure 6, above.

The POP then connects to the Metro Hub which in turn distributes the traffic to the Community Cabinets in the access layer. Each hub supports anywhere from 10-30 thousand households, while the community cabinets provide services for about 850 households. The routers within the hubs connect to the core and to the hubs via GE, and routers in the hubs connect to the subscribers via 100Base-BX providing range of up to 10Km over single strand fiber. Core nodes interconnect to each other and to the POPs via multiples of GE, though 10GE is possible in the future.

From the cabinet, the fiber runs aerially or underground to splices that service 4-6 households. Dedicated fiber runs from the splices to the subscriber, and connects to an access portal (residential gateway) that supports 3xFE, 2xPOTS, and VoIP. This then connects to the user's PCs and TVs. The device also supports VLANs, ACLs, and rate limiting, all useful features when mixing multiple services that require

different QoS. The video STB is separate from the gateway, though other devices do combine the functions, and the choice of gateway is usually separate from the aggregation and core decision due to the number of vendors in this segment of the market. For business subscribers, nothing would preclude the use of CPE that includes T1/E1.

At the protocol layer, content and application providers connect to the core using VLANs (representing services), which are then tunneled into VPLS instances for transport to the edge of the network. Here, VPLS terminates and the VLANs provide services to the individual subscribers. Video distribution relies on IP multicast.

A VPLS based Layer 2 VPN solution was selected since the network is designed to provide open access to any mix of content and application providers. For this reason, a Layer 3 solution where the provider has to manage the subscriber and content provider IP address space is unacceptable (i.e., multiple application providers with the same IP addressing). A Layer 2 VPN insulates the provider from these complexities, and the content providers are free to manage their own IP addressing. VPLS in the core is also preferred due to its resiliency, scalability, and QoS support vs VLAN-based solutions. Video services are delivered via an IP multicast overlay to the VPLS backbone.

Dynamic City operates the network as a utility, much like a power or road network, and subscribers pay the individual content providers for services. A small percentage of these fees are transferred to Dynamic City as the network operator.

## Appendix B – Next-Generation DSL Deployment – Telefonica Imagineo

Telefonica de Espana's Imagineo is a multimedia network for consumers combining Ethernet-enabled DSLAMs with Ethernet routing in the aggregation network as depicted in Figure 7. In some locations, remote DSLAMs connect via Ethernet-enabled SDH systems into routed Ethernet core. As with the UTOPIA deployment, the Ethernet core relies on MPLS/VPLS.

Initial services across the network will consist of high-speed Internet access, video-on-demand, and 30-40 channels of broadcast TV. Future services could include online gaming, interactive TV, and VoIP.

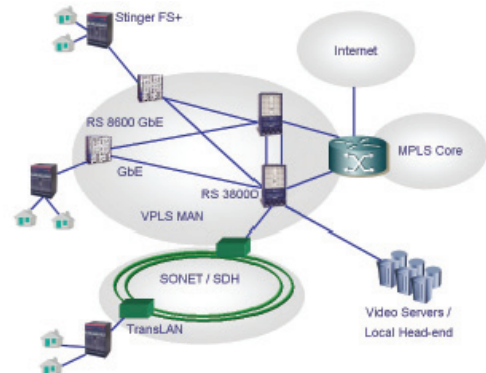


Figure 7: Telefonica Imagineo Network

Individual services such as video and Internet access map into individual VLANs. These then map into VPLS instances (FECs) across the core network. Video traffic is carried via unicast or multicast depending upon whether it is VoD or broadcast. Protocols in use across the MPLS backbone include link aggregation, fast reroute, and .1p to EXP bit mapping. Edge protocols include MAC move detection and loop detection.

### References:

<http://www.utopianet.org/news/file-11.pdf>

<http://www.convergedigest.com/blueprints/ttp03/z2utopia1.asp?ID=37&ctgy=Loop>

