White Paper

Choosing the Right FTTx Architecture

Comparative Costing for GPON, Active Ethernet, and FTTN Approaches

Most followers of technology trends in the access network have come to understand that pushing optical fiber all the way out to customer premises is no longer an extravagant fantasy — it is becoming reality today in a rapidly growing number of networks all across the world. With the advent of fiber in the access network has come some confusion, however, as there are many different approaches being used by operators and evangelized heavily by vendors, both too often laced with bias based on what they've done in their own network or what they have in their product portfolio, respectively. Sorting out the pros and cons of these different approaches can be difficult without some good, objective guidance.

The following material is designed to provide just that. We've taken an analytical look here at the drivers of demand for bandwidth, network cost pressure, the limitations of the most advanced copper-based technologies, the comparative costs of a handful of the most important architectures for fiber in the access network, and other considerations that can come into the decision-making process for network architecture choice.

We conclude with the same view on "what is the right answer?" that many analyses of service provider decision drivers have offered: it depends. The best network choice for you will depend on many factors, and contrary to what some vendors would like the world to believe, there is no single right answer for fiber architectures in access. This will be more clear after we've looked in more detail at the network choices and their attributes in the following sections.

A word on terminology — we will refer here to the general category of fiber in the access network with the common term FTTx, or Fiber to the "x", where the x can be the remote node, the curb, the basement, or the individual residence or business itself. As we'll see, what you put in the "x" is one of the key variables in determining the cost and utility of your fiber access approach in the immediate and longer term.

Drivers of Fiber Deployment in Access

Before assessing the comparative costs and merits of different architectures, it's worth understanding in general terms why fiber in the access network is being deployed in the first place. There are four primary reasons: an increase in demand for subscriber bandwidth, increasing network cost pressure, the limits of copper technology, and competition.

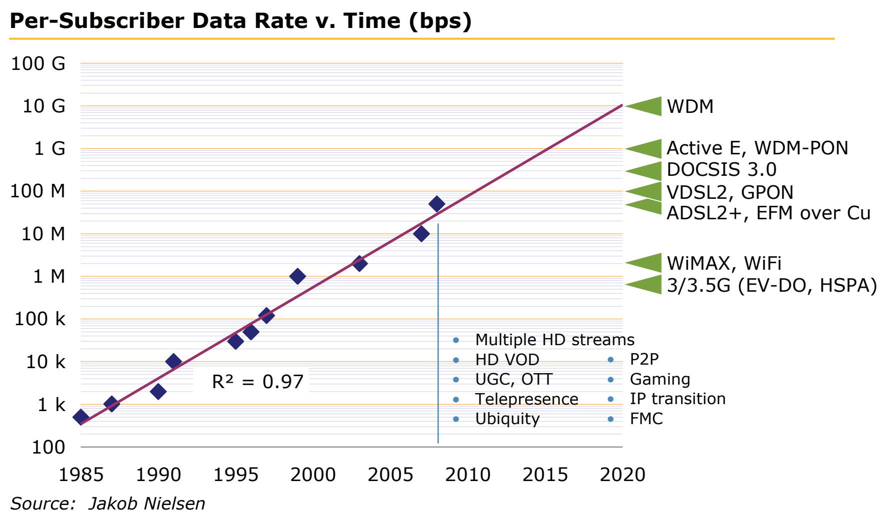

Figure 1:

Continued growth in subscriber demand for bandwidth is a key driver of the pursuit of alternative access technologies

Driver 1: Subscribers want ever more

The industry has seen a consistent pattern of growing subscriber demand for higher data rates — more bandwidth — ever since wide-scale computer networking first became possible. As we show in Figure 1, the pattern is consistent indeed — a regression on the log of modem data rates v. time for access technologies since 1985 yields an R^2 of 0.97, a remarkably tight correlation! Granted that correlation does not equate to causality, so one must exercise some caution not to simply extrapolate the trend blindly. Nonetheless, it does appear to fit intuitively with our experience of the history of access. Time and time again, the technologists in the back room would endeavor to create the next hot access technology, with an impressive boost to data rates, and when it was first introduced into the field, most subscribers would wonder, "What am I ever going to do with all THAT?" But before long the collective ingenuity of subscribers and application developers and service marketers would conspire to more than fill the pipe up again.

As we indicate in Figure 1, there a sufficient backlog of bandwidth-hungry applications used heavily among early adopters and younger users today that will continue to drive this trend. It is likely safe to assume bandwidth demand growth over the coming 10 to 15 years will not deviate far from this trend.

Driver 2: Network cost pressures mount ever higher

While subscriber bandwidth demand rises inexorably, unfortunately their willingness to pay monthly fees does not. This creates downward pressure on the unit cost structure of carrier networks. To illustrate the magnitude of this pressure, we show in Figure 2 the typical monthly revenue for different service classes, but divided by the amount of capacity each consumes in that month. Note that these figures were calculated for commercially-available packages in the US market — prior analysis of other markets (both developed and developing) have shown that the relative prices and bandwidth consumption of these services is similar, so the graph ends up looking largely the same no matter which denomination of currency (and hence market conditions) you choose.

The important thing to note is that the most common directions of service portfolio extension for access service providers — to add cellular data to voice, for the wireless carriers, or to add video to a broadband plus voice package for wireline carriers — is driving their business to the left on this chart. We include the broadcast-model satellite and cable TV reference points here not because they are directly relevant for comparison, but because they represent (to an arguable degree) the future of an access world where over-the-top video played over a wired broadband connection becomes the functional substitute in the subscriber's life for what used to be broadcast TV. Clearly, networks with higher capacity per unit of capital and operating expense can create higher headroom and strategic advantage in this environment.

Willingness to Pay per Unit of Capacity Consumed

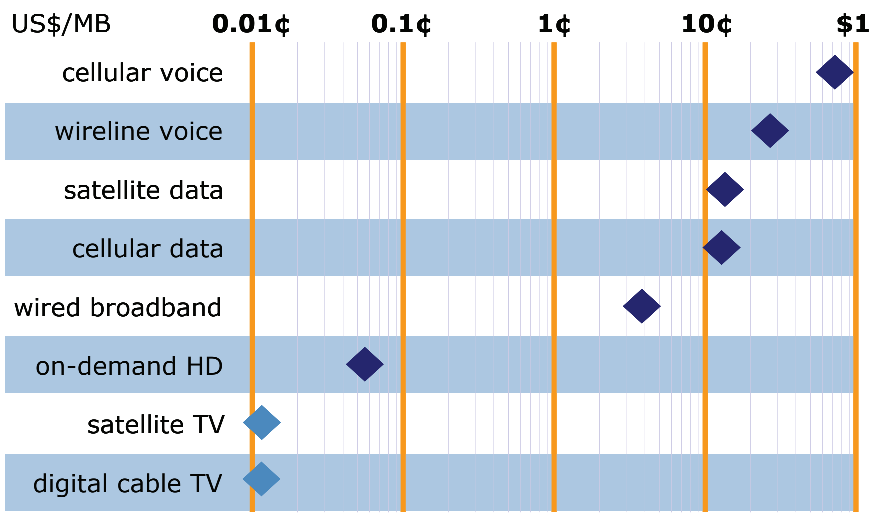

Sources: US commercial network offers circa January, 2008; Zhone analysis

Figure 2:

Service portfolio expansion is tending to create substantially greater network cost pressure, given falling average subscriber willingness to pay per unit of capacity consumed.

Driver 3: The limits of copper-based access technology.

While there is much excitement in the access industry about the latest developments in copper-based technologies such as ADSL2+ bonding, VDSL2, and EFM (802.3ah-standard Ethernet over copper), and these technologies certainly have many applications to which they are perfectly suited — it is important to understand their limitations in applicability to the general access network challenges indicated by increased bandwidth demands but falling relative costs.

As we show in Figure 3, these copper-based technologies can provide subscriber data rates well in excess of 100 Mbps — but only over fairly short range. Achieving these higher data rates requires the use of more complex waveforms, which require higher signal-to-noise (SNR) ratios for intelligible receipt of the signals. Over real-world copper loops, these high SNR values cannot be sustained for more than a few hundred feet. As a result, copper loops can no longer be counted on to play an exclusive role in the provision of adequate bandwidth in access networks over the lon-

Rate/Reach for Current Cu-based Broadband Access Technologies

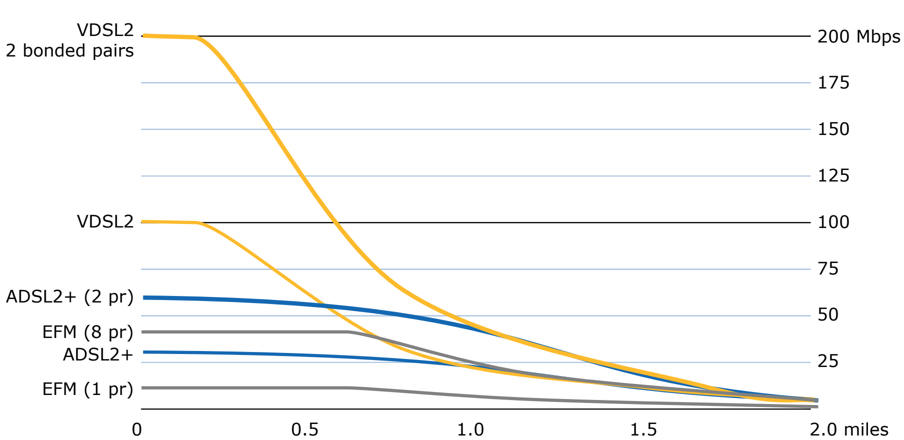

Sources: DSL Forum; Zhone testing.

Figure 3:

Limited range for the most advanced copper-based access technologies reduces their long-term utility in the access network given growing subscriber demand for bandwidth.

Rate/Reach for Fiber-based Broadband Access Technologies

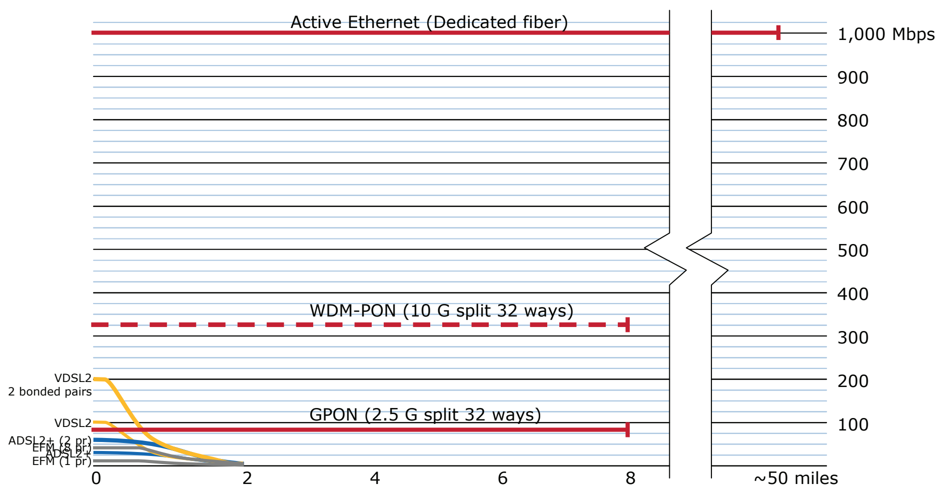

Sources: DSL Forum; Zhone testing.

Figure 4:

Fiber provides more viable alternatives for the long term.

ger term. They can provide a medium-term solution in general application (especially EFM and ADSL2+ bonding), and a longer-term solution for in-building or in-neighborhood distribution (VDSL2), which we'll discuss further in a moment.

Figure 4 shows the rate and reach of the copper technologies in comparison to what is possible with fiber. The rate and reach advantages of the fiber approaches are clear, considering the longer-term demands for bandwidth. It's worth noting that the passive optical network (PON) approaches, whether single frequency (GPON) or multiple (WDM-PON) are more limited in range than the point-to-point or Active Ethernet approach. The passive splitters involved in a PON architecture reduce the signal strength from the downlink transmit laser, so their practical limit is 12 km, or about 8 miles, on average across all endpoints.

Driver 4: Competition is on the rise.

Increased subscriber demand for bandwidth is not all by itself an adequate driver for new network investments in most circumstances. The final and essential ingredient in the mix is the presence of credible competition. The specifics vary by market, but whether it's unbundling of local loops creating opportunities for alternative carriers to deploy innovative architectures, municipal broadband projects, or deeper penetration of the fiber segment in cable television networks (along with the adoption of the next generation of the cable TV standard, DOCSIS 3.0), competition in access is on the rise all over the world. Increased carrier activity to deploy fiber in Europe, the Middle East, North America, and Asia is a direct result.

Modeling Considerations for Last-Mile Decisions

The general benefits of fiber-based architectures are clearly substantial. To examine the merits of different FTTx approaches in more detail, we must turn next to an analysis of their requirements in deployment and the associated costs. There are many, many variables in any reasonably realistic model of the up-front capital and ongoing operational costs of an access network deployment (fiber based or otherwise). Enough of them are market- and carrier-specific to make absolute accuracy in a general-purpose comparative model such as what is required here impossible. Our goal here is instead to provide a directionally-correct look at relative costs. This requires making reasonable assumptions about market and carrier-specific variables,

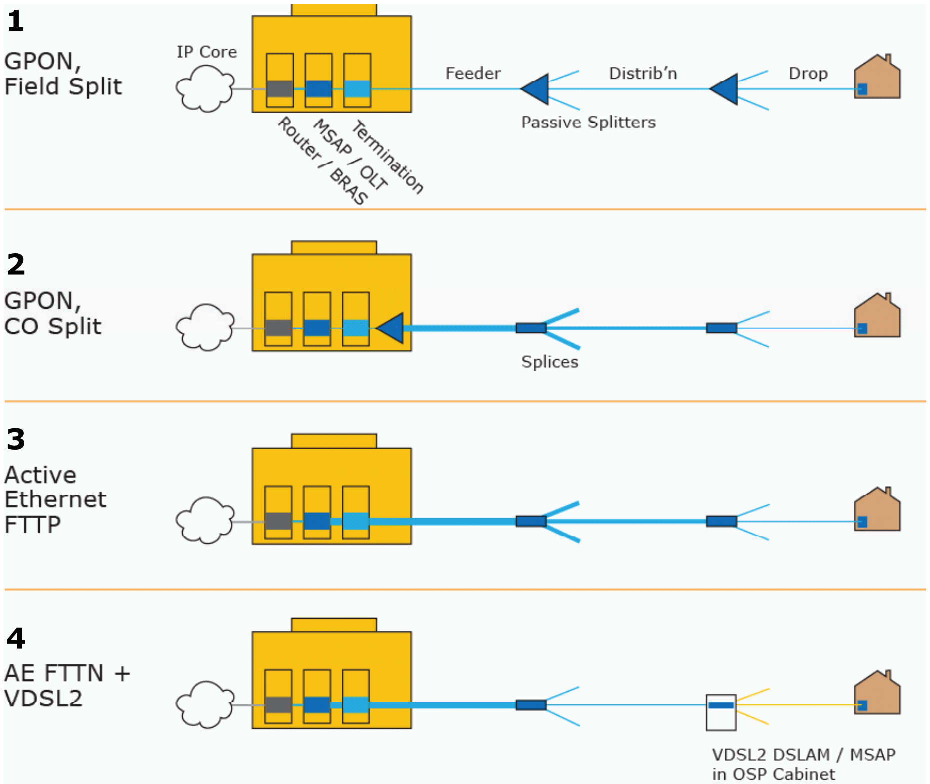

Figure 5:

Four basic fiber deployment options cover the relevant range of alternatives for comparative cost analysis.

excluding irrelevant costs, and calibrating the total costs to current reports from real-world deployments.

The architectures we assess here are illustrated in Figure 5.

Scenario 1: GPON, Field Split . . . . Whether used for Gigabit PON (GPON) or any of the other PON variants, this is the "classic" passive optical network architecture. This involves usually two levels of passive splitters in the network, a 4-way split between the feeder and distribution legs of the access network, and an 8-way split between the distribution and drop legs. As shown in the figure, the aggregation of these 32 endpoints back to the central office (CO) is carried on a single fiber. This fiber is terminated in the CO and aggregated with others in a multi-service access platform (MSAP) or dedicated optical line termination (OLT) unit.

Scenario 2: GPON, CO Split . . . . A number of operators are taking advantage of shared optics on the CO side for their low cost initially but deploying fiber in the field that enables an easy transition to dedicated, point-to-point optics in the future. In this CO Split scenario, multi-fiber bundles are deployed in the feeder and distribution legs, with splices to a number of smaller fiber bundles at the same points in the network occupied by the splitters in Scenario 1. The passive splitters in this case are used only in the CO. A transition to Active Ethernet from this scenario can be made without touching the fiber plant in the field, a very useful option to have for the longer term. For comparison purposes our cost analysis includes a wave-division multiplexing (WDM-PON) variation on this scenario.

Scenario 3: Active Ethernet to the Premises . . . . This Active Ethernet (AE) architecture uses the same plant in the field as Scenario 2 to take so-called "home run" or dedicated fibers all the way from the CO to the customer premises — but in this case each dedicated fiber is terminated individually at the MSAP/OLT. This requires a dedicated laser on the CO side for each customer drop. As an aside here, it is worth noting that one of the criticisms of AE architectures sometimes voiced in the industry conversation about FTTx — that there are unacceptable costs and risks associated with putting active electronics in remote locations in

capex		opex
customer prem	central office	opex per sub per month
CPE	fiber termination per sub	OSP maintenance per sub-month
installation	subs per terminated fiber	OSP maintenance per sub, annual
hours	cost per fiber termination	hours/year per premise
cost per hour	splitters per terminated fiber	hours/year per handhole or splice
drop cable	fibers per patch panel rack	hours/year per cabinet
length in m, average	patch panel rack	hours/year per DSLAM
cable cost per m	installation cost per fiber	hours/year per cable km
install cost per m	hours per fiber	cost per hour
splice or connector	cost per hour	CO operation per sub-month
handhole + splitter cost per sub	connectorized jumper	people costs per sub-month
subs per handhole	aggregation per sub	people costs per sub per year
splitter	subs per port	subs per head
handhole unit cost	aggregation per port	CO operation hours per head
installation	chassis cost per port	cost per hour
hours	each chassis	facilities costs per sub-month
cost per hour	uplink	% of 1 rack per sub (term + agg)
node cost per sub	slots per chassis	cost per rack per month
subs per node	card cost per port	backbone connectivity per sub-month
handholes per node	card itself	busy-hour bandwidth per sub, kbps
hardware and install total	ports per card	backbone cost per Mbps
distribution cable (to handholes)	SFP per port	depreciation expense per sub-month
length in m, average	backbone connectivity per sub	network amortization timeline, years
cable cost per m	subs per backbone port (1 agg shelf)	
install cost per m	backbone router cost per port	
cabinet (or handhole), outfitted	electronics installation + turn-up per sub	
VDSL2 DSLAM	hours per shelf	
AE ONT (or uplink for DSLAM)	cost per hour	
splitter		
splices or connectors		
quantity		
price each	survey / site engineering costs	
installation	"tax" rate on total	
hours		
cost per hour		
feeder fiber cost per sub		
splice/connector		
total fiber cost		
fiber length, m		
cable cost per m		
install cost per m		

Figure 6:

Meaningful comparison of FTTx alternatives requires analyzing their relative costs in some detail.

the field — is not based on an adequate understanding of how these point-to-point architectures are actually deployed. The active electronics need appear only at the CO and customer premises ends of the fiber — the fiber and splices in between are completely passive.

Scenario 4: AE Fiber to the Node + VDSL2 . . . . Where copper in the last few hundred feet of the access network is available, it can be very attractive to consider using dedicated fiber in the feeder and distribution legs, but advanced copper-based technology (specifically VDSL2) in those last few hundred feet. The VDLS2 aggregation can be deployed in an outside-plant cabinet (a "remote node") in a neighborhood of single-tenant buildings, or for multi-tenant scenarios a variations referred to as Fiber to the Curb or Fiber to the Basement can be used. The costs of these latter two variations, relative to Scenarios 1–3, are reasonably similar to FTTN, so we've condensed the analysis to the FTTN approach.

The elements of the model we've constructed for comparative assessment of the economics of these four scenarios are outlined in Figure 6. As we noted above, the goal was to include only relevant costs and provide adequate resolution to get the relative impact of the differences among these four scenarios directionally correct. Reasonable assumptions were made for the unit cost values of the 60+ variables in the model, based on information from a variety of public sources, vendor and operator presentations, and Zhone experience with customers and prior network cost modeling work. These assumptions were held constant across the various scenarios unless deployment specifics required variation, to maximize the "apples to apples" nature of the comparison. The aggregate results were calibrated with recent reports of operator costs in actual large-scale deployments.

Network Cost Comparisons

Looking at the results of the model (see Figure 7), one can clearly see why operators are making some of the choices they are in their networks. The GPON field split scenario provides the least expensive approach to running fiber all the way to the customer premises. The marginal cost of moving to a CO split model is relatively low (at a little over US$100 per subscriber), which many operators see as a modest price to pay for future flexibility while still tapping the lower cost of the GPON shared optics model. With all else held equal, the point to point AE model adds about US$500 per subscriber over the field-split GPON scenario, primarily on the CO side, for the additional optics and

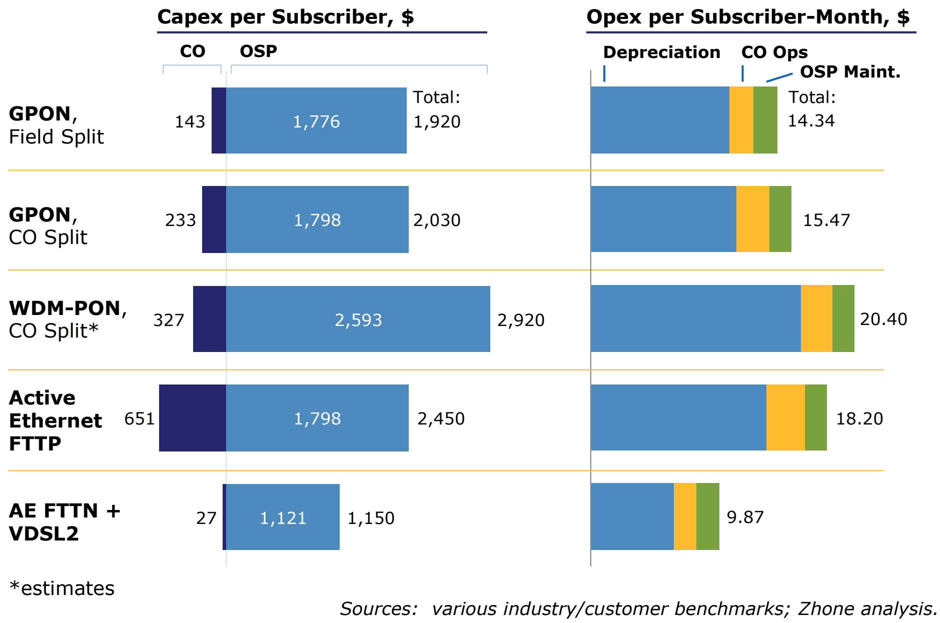

Figure 7:

Cash costs of the FTTx architectures considered vary most significantly in the up-front capital required for deployment.

aggregation electronics there. WDM-PON, with its much more expensive optics on both ends of the fiber, is unsurprisingly the highest cost pure-fiber model. Finally, the attractiveness of FTTN+VDSL2 is also quite clear, for situations where existing copper is available in the last few hundred feet.

The ongoing operational costs of these scenarios differ little in cash terms, although the large variation in capital costs required per subscriber creates similar variations in P&L expense terms.

Recalling the network cost structure driver of interest in fiber deployments, it's instructive to look at these network architectures on the basis of cost per unit of capacity provided, as show in Figure 8. Using this metric, the gigabit per second capacity of the AE approach offers a very clear long-term advantage. The

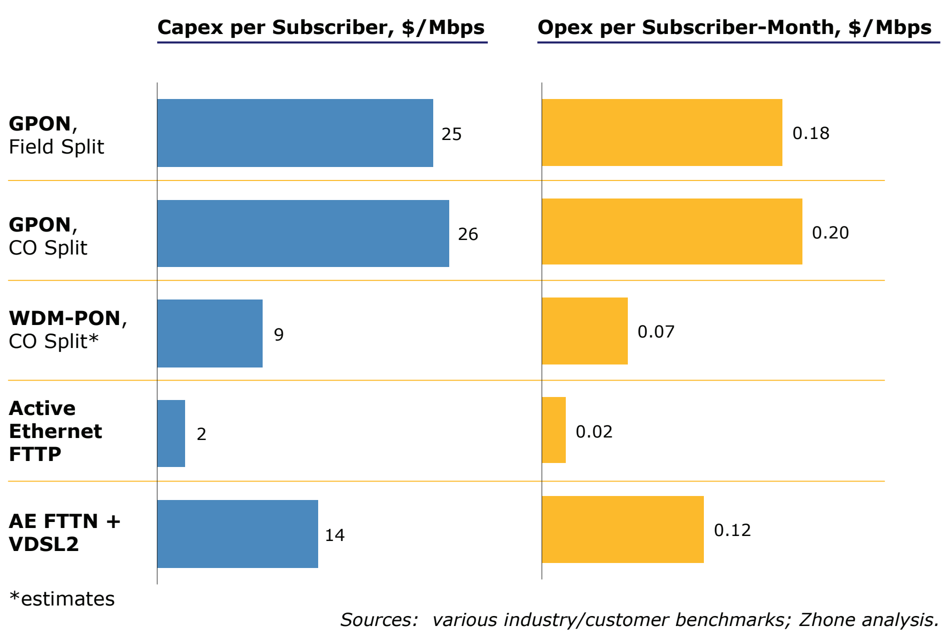

Figure 8:

Active Ethernet provides the most significant operator headroom in terms of cost per unit of capacity created.

bandwidth advantages WDM on both ends of a PON network appear to justify considering the extra costs as a possible upgrade path from GPON. As with the costs in absolute terms, the advantage of AE+VDSL2 is clear in this analysis as well.

Conclusions

So what's the "right" access technology to choose, in the end? As we said at the outset, it depends. The following table provides a very high-level view of the considerations that might drive selection of one architecture over another for a given portion of an operator's network:

Factor	GPON	AE FTTP	AE FTTN + VDSL2	ADSL2+ Bonding
Investment time horizon	Medium	Long	Medium	Short
Service definition	Tomorrow	Long Term	Tomorrow	Today
Customer segment needs	Residential	Residential + Business	Residential	Residential + Business
Copper assets	Limited or none	Limited or none	Plentiful	Plentiful
TV Approach	RF or IP integrated	IP or separate network	IP or separate	IP or separate

There are of course gray areas among all these choices and factors, but this provides a general sense of direction.

One of the most important factors influencing the decision, and one that can vary significantly across a wide spectrum between operators, markets, and even points in time, is the availability of suitable capital. While the creation of a new broadband infrastructure with the kinds of fundamental advantages fiber architectures can offer can certainly provide a step-function increase in competitiveness to an operator or community that deploys it, other constraints on capital may not allow it. This can drive a choice of GPON over AE, use of copper in the last few hundred feet via VDSL2, or even making the best of available copper exclusively, but moving to ADSL2+ bonding for higher bandwidth services.

Other factors of a given situation can tend to narrow the range of options as well. If a telco is already providing video service using an analog head end over RF (or if they are choosing to deploy RF video technology rather than IPTV because of concerns about cost, expertise, or reliability with HD content), the inherent capabilities of GPON to provide an RF overlay may tip the balance in favor of that architecture. Conversely, a focus on both business and residential subscribers may drive network designs toward the point-to-point Ethernet approach, given the much higher bandwidth available immediately with AE. In markets where the last few hundred feet of copper are unavailable, or where public policy dictates full FTTP architectures for new construction, the otherwise attractive AE+VDSL2 approach is not an option.

Ultimately, we've found that the right choice for an operator tends to be more than one of these network architectures — and in some cases, even for relatively small operators, we've seen all of them employed in the same network. Different segments have different requirements and different initial conditions (primarily in terms of network assets already in the ground or the building), so the ideal solution is actually a mix of these architectures.

This brings us to a final word of advice as you look at how you might leverage fiber in your access network, related to your choice not of the technology, but of your technology provider. There are many sources of access networking gear in the market today. Many tell a similar-sounding story about how their equipment supports multiple services and seamless transitions from copper to fiber, or from TDM to IP. How they provide a single management interface to disparate systems, making your life easier. Or how they can scale to meet your needs as service complexity and bandwidth grow. When it comes to true multi-service access capability, it pays to inspect these claims carefully. Most suppliers have "bolted on" fiber architectures to their copper-based access products, without the kind of integration required to make heterogeneous network architectures — like those we see operators implementing in the real world — easy to use and upgrade over time.

Zhone has been the industry's pioneer in multi-service access platforms since we invented the category more than 9 years ago. Today we have the industry's broadest portfolio of truly integrated access solutions (see Figure 9) for both copper- and fiber- based services — all implemented in a family of products that share a common single line, multi-service architecture. This tight integration across xDSL flavors, GPON, and Active Ethernet provides you the versatility and flexibility you need, across all the elements in your network from CO to the customer premises, to manage your network for maximum strategic advantage at the lowest cost over time. For more product and contact information, please visit www.zhone.com.

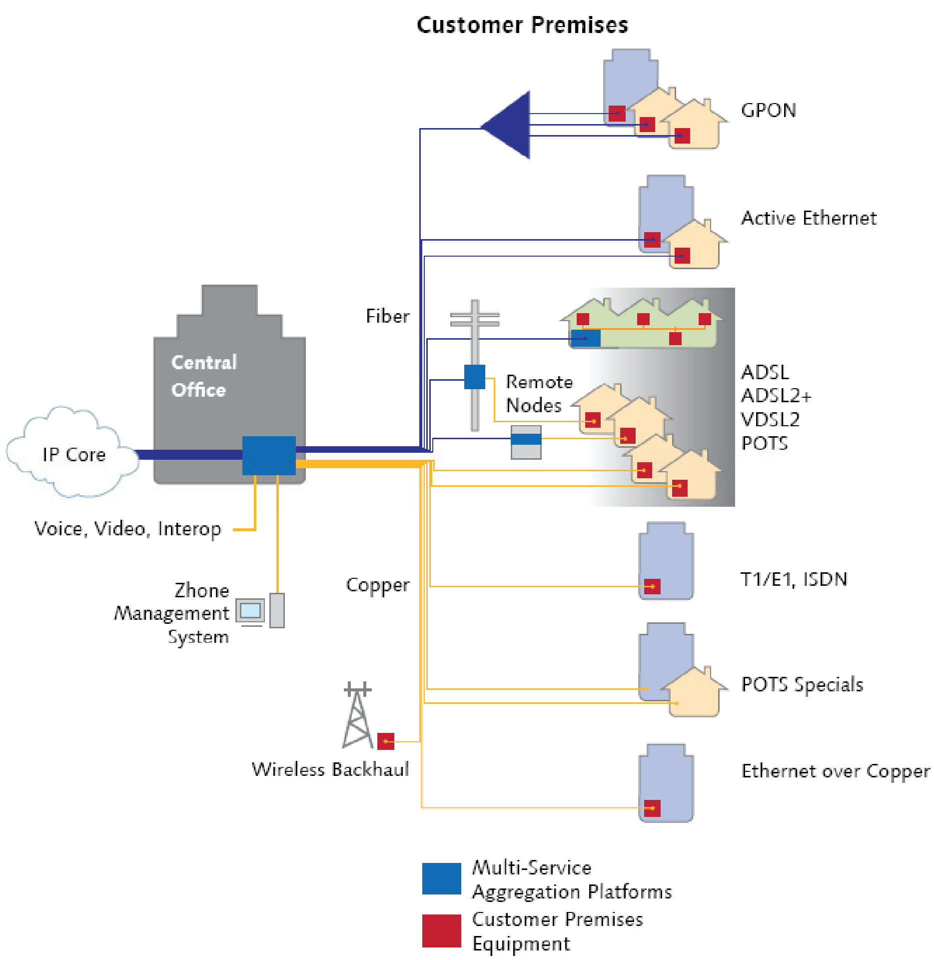

Figure 9:

Zhone provides the industry's broadest portfolio of current and legacy last-mile technologies, all within one common multi-service aggregation and premises architecture.

Zhone Technologies, Inc.
@ Zhone Way
7001 Oakport Street
Oakland, CA 94621
+1 510.777.7000 Tel.
www.zhone.com

For more information about Zhone and its products, please visit the Zhone Web site at www.zhone.com or e-mail info@zhone.com