Targeting the Small / Medium Business Market with Ethernet in the First Mile Access Network

The practical road to profitable broadband services deployment

It's hardly surprising that so many competitive carriers in the US have fallen by the wayside in recent months. You only have to look at the economic models and the business realities they've been dealing with, to realize that naive optimism, high costs of deployment and the general obstructionism of the incumbent providers have all contributed to the long list of obituaries.

Naïve optimism

Let's examine the broadband services revenue opportunity in the small and medium business (SMB) market. Equipment vendors and service providers alike have harped on and on about offering high margin value-added services on top of basic Internet access. But the grim reality is that the average take-up of services in 2000 by US SMBs was less than two services per customer. For example, the take-up rate for voice services was under 1.5%. In other words, only some SMB customers adopting broadband Internet access are also buying value added services such as email, web hosting, data backup, virtual private network (VPN) or voice services. Few are buying more than one additional service, and many, none at all.

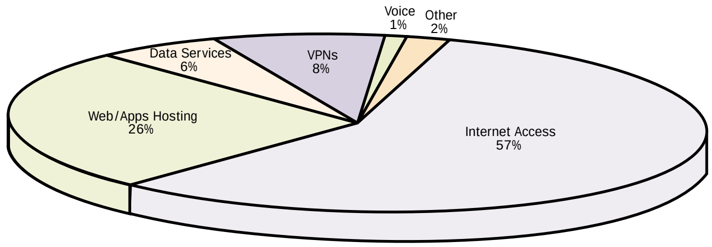

US ISP and CLEC (Tier 3) Broadband Services Revenue 2000

Source: Infonetics Research

It's still early in the game and those competitive carriers still in business remain optimistic. Their foot is in the door with Internet access and their customers will in time buy more value added services. And since the value added services accrue much higher margins than Internet access, they should be in good shape, if they can stay the course long enough. The point is, service providers should not bank on the futures of high margin services they may never get. Instead, they must develop business models that deliver profitability with mostly Internet access and low take-up of value added services.

High deployment costs

The problem stems from the high cost of service deployment. In terms of access technologies, service providers have a wide array choices at their disposal. These include anything from dedicated T1 to digital subscriber line (xDSL) and fixed wireless. Each technology brings with it inherent cost, performance and deployment trade-offs.

Although T1s are widely available and offer the best geographic reach, a service provider's cost to lease unbundled T1 loops makes the resulting retail price far too expensive for most SMBs to consider. While symmetric flavors of xDSL such as high speed DSL (HDSL2) and symmetric DSL (SDSL) can offer equivalent or better performance at a lower retail price point, they introduce difficulties of their own – the need to own and co-locate digital subscriber line access multiplexor (DSLAM) equipment, line qualification, long lead times from incumbent providers and "truck roll" logistics issues.

CO-based DSL deployment versus dedicated T1 cost structure

The lower retail pricing of xDSL creates an illusion that the cost structure is more favorable than dedicated T1s. In fact the two are much closer than people realize with xDSL requiring a higher upfront investment in capital equipment.

With dedicated T1 access, service providers can aggregate T1 local loops served from multiple central offices (COs) into a single point of presence (POP). At each CO the T1 circuits can be groomed into channelized T3s that are presented to the service provider's POP. The access equipment requirements are minimal - a T3 aggregation switch at the service provider's POP and a T1 router or integrated access device (IAD) at each subscriber's customer premises.

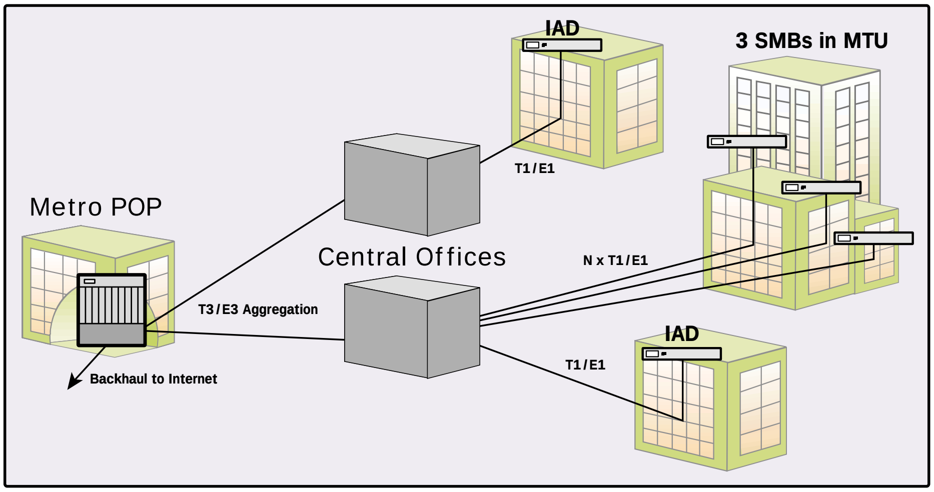

In contrast, when deploying DSL services, to enjoy a similar footprint, a service provider needs to co-locate DSLAMs in multiple COs, even before there are subscribers local to that CO, or at least when the first subscriber signs up. That's an initial capital investment of about $50K per CO for let's say 50 DSL ports.

There are an estimated 12,000 COs in the US and 1100 cities with a population over 25,000 people. A typical city with 25,000 to 50,000 people may have only two COs, while large cities such as Chicago or New York may have upward of 20 COs. To establish a footprint in an average sized metropolitan service area (MSA) a service provider may need to co-locate DSLAMs in as many as 10 or 12 COs. That's a half to three quarters of a million dollars upfront capital investment that is not required when deploying dedicated T1 service. On top of that, is monthly co-location facilities charges at roughly $1000 per month per CO - that's a burn rate of $10K-12K per month, whether you have customers or not.

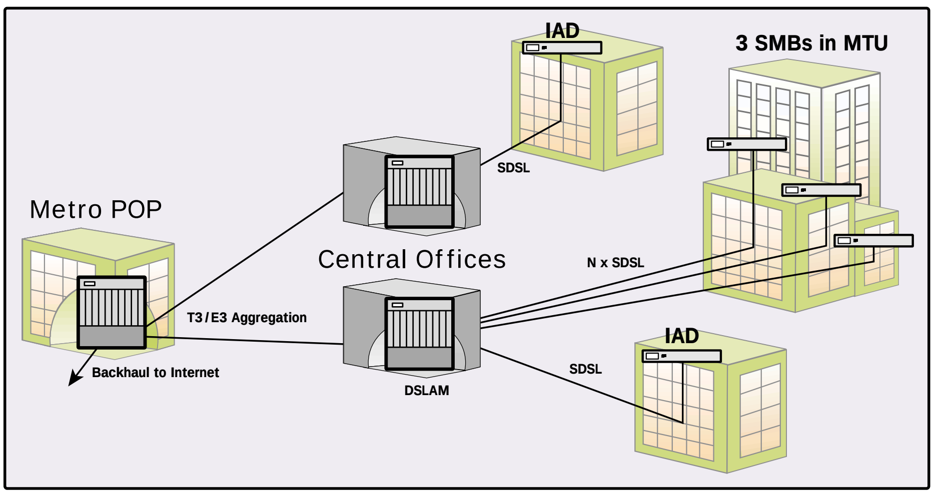

A similar infrastructure investment is required for local multipoint distribution services (LMDS) and multipoint microwave distribution services (MMDS) fixed wireless broadband access. Although, in its favor, since there's no interaction with the incumbent telco, service providers do have the advantage of being able to reduce the lead-time between commissioning a base station and turning on their first subscribers, thereby achieving a faster return on investment and better cash flow. However, initial rollouts of fixed wireless in 1999 and 2000 have not been as smooth as many service providers expected. A lack of network engineering experience, product immaturity and susceptibility to weather conditions has left many service providers and customers with a bad taste in their mouth.

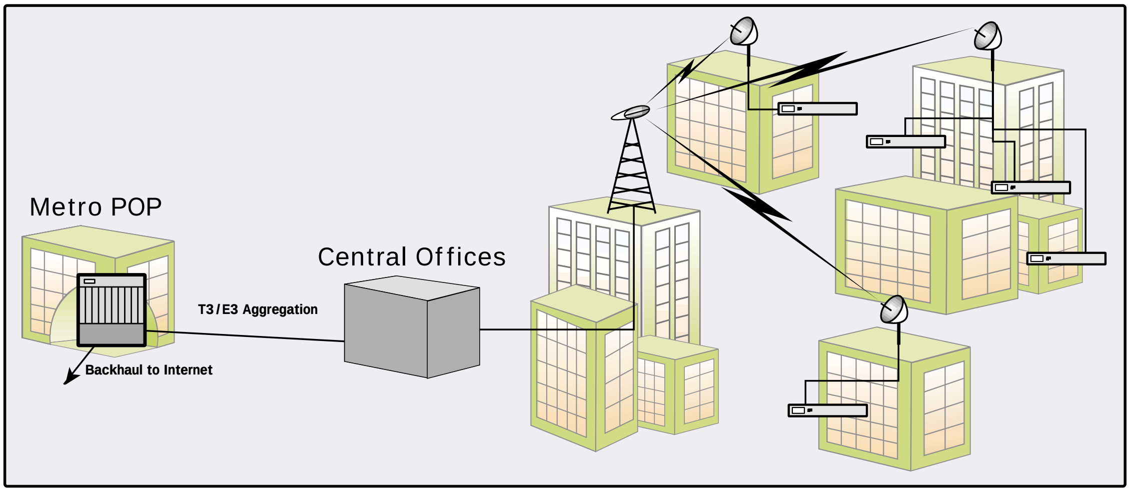

This is all well and good, but "I can't sell dedicated T1/E1 services to my customers – it's too expensive", I hear you say. This is true. Customers have been tainted with the lure of cheap DSL. But how do you give them DSL without falling prey to the incumbent? It goes without saying that the more a competitive service provider has to deal with an incumbent provider, the higher his costs, and the lower his ability to differentiate. CO-based DSL deployment epitomizes this dilemma, with a competitive carrier requiring multiple complex interactions with the incumbent in order to turn on each and every subscriber. However, as many fixed wireless casualties have observed, total avoidance of the incumbent provider with an unstable technology isn't the right answer either.

The bottom line: SMBs can't afford dedicated T1/E1 service, and competitive service providers can't afford CO-based DSL service deployment – so what is the alternative, the answer lies not only in technology choices, but also in market focus.

SMB Market Opportunity and Demographics

In the US, it is estimated that over 65% of small and medium businesses reside in multi-tenant unit office buildings (MTU). In Europe and Asia Pacific the percentage is even higher. According to Dun & Bradstreet, there are approximately 1.3 million small and midsize businesses in the US – that's about 840,000 located in MTU office buildings. In total, they spent approximately $47 billion on voice communications services and over $14 billion for data and Internet services in 1998. By 2002, these figures are projected to reach $53 billion and $40 billion, respectively. These numbers illustrate a key trend in the MTU market: a real and growing demand for communications services in a market segment that is among the least well served by the telcos. What's more, SMBs with fewer that 25 voice lines are often under the incumbent telco's radar screen, so you won't see them putting up much of a fight if you start prying these customers away from them.

Targeting MTUs with a "building point-of-presence"

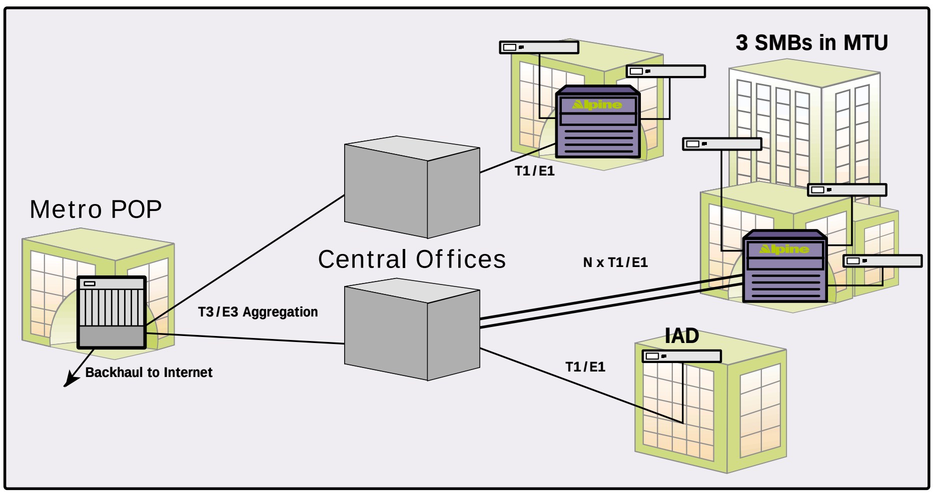

To target the MxU opportunity, service providers must first negotiate "right of entry" (ROE) agreements with the property owner or REIT (real estate investment trust). An ROE agreement with a single property owner can give you access to hundreds of buildings. In return for a 3-7% revenue share, these ROE agreements then allow the service provider to install equipment and cabling in the building – in essence, a "building point-of-presence", and to market services to its tenants. In keeping with the 1996

Telecommunications act, the FCC has recently ruled that US property owners may not have exclusive agreements with service providers. In reality, most property owners already embrace more than one provider, and are objecting to the US governments' bullying tactics of "forced entry". While there is clearly an advantage in being a first mover, a second or third provider can get a good crack at the whip. In any case, the good thing is property owners get it!

What is a REIT?

Real estate investment trusts are legally required to pay virtually all of their taxable income (90%) to shareholders every year as dividends, thereby avoiding paying corporate taxes. There are over 300 REITs operating in the US today, with assets of over $300B. About two thirds of these are publicly traded. There are several types of REITs: Equity REITs own and operate income-producing real estate; Mortgage REITs lend money directly to real estate owners and operators; and Hybrid REITs own and operate income-producing real estate and lend money. Equity REITs often specialize in specific types of properties buildings e.g. shopping centers, apartments, warehouses, office buildings, hotels, healthcare facilities, etc. US REITs are estimated to account for 70-80% of US commercial building stock. Meanwhile the REIT concept is spreading abroad. Japan and Mexico are considering REIT-like entities, and Turkey, Belgium and the Netherlands already do.

Property owners see broadband services as means to increase the attractiveness of their property and reduce customer churn. With compelling statistics like this… A Jan 2001 survey of 454 telecommunications decision makers of business tenants in US MTU buildings found that: "Four-out-of-ten businesses will consider moving elsewhere at lease renewal if their important telecommunications needs aren't met at their current location" … They may even help market your services!

Business economics of a "building point-of-presence"

Delivering IP services into the commercial MTU market has tremendous appeal. It is a vast, largely untapped market worth $20-30B in annual service revenues even with today's low take-up rates for value-added services. But more importantly, it's a market in which lower operational costs make profitability within two years an achievable objective. The cost savings come in many forms. First, by installing equipment in the basement of the building, service providers avoid the astronomic co-location charges incurred with CO-based DSL deployments. Second, by eliminating all the interactions with the Telco required to qualify each subscriber line, there is a dramatic saving in administrative costs and finally there are significant savings in customer acquisition costs – the customers are a captive audience – they are the companies in the buildings to which you have rights of entry. That means low cost targeted marketing via in-building posters, mail drops, even door to door selling (where it's still legal!) can be extremely effective. It is estimated that MTU customer acquisition can be reduced by 90% over traditional mass-market promotional strategies.

Unlike consumers, most businesses are under ever growing pressure to be connected, and have the means to pay for business-class Internet access and additional services. There is a wealth of market data to support a price tolerance of $150-400 for basic Internet access depending on the bandwidth offered (Probe, Infonetics, Yankee group). By deploying cost-effective broadband technologies service providers need only three subscribers per building averaging $300 service revenues per month per subscriber to reach profitability in two years. That is easily achieved with 384Kbps or 512Kbps Internet access and with just one or two subscribers getting one other service (web-hosting, data backup, voice etc).

Examining the building stock of commercial multi-tenant buildings shows that three subscribers per building is a reasonable target. In the 91,000 buildings with 6-10 tenants, that's an average penetration rate under 40%.

Tenants per Building	Total Buildings	Total Tenants
6-10	91,000	440,000
11-20	28,000	364,000
20+	31,000	682,000
TOTALS	**150,000**	**1,486,000**

Source: US Department of Energy 1995

In the US, there are also over 200,000 MTUs with fewer than six tenants. But going after these is probably not a good strategy. You need over 60% penetration unless these are SMBs or large enterprises with high numbers of employees that will drive bandwidth above the typical 384Kbps connection speeds. Even still, in such cases it may prove more economic to offer a dedicated T1/E1.

Of the 4000 or so competitive carriers in the US, fewer than a hundred have focused exclusively on MTUs. The market is far from saturated, and those still struggling in the CO-based DSL space are beginning to reconsider their business strategies. Despite the improved economics over CO-based DSL deployment however, some notable MTU-focused providers still failed. The main causes were twofold: overly complex and costly service delivery models (e.g. ATM WAN), and the high investment costs of re-wiring buildings, before having paying customers. So beware, if you go after the MTU space, you still have to make architectural and technology choices that are practical to implement. The days of land-grab, and uncontrolled investment are over. To survive, you have to get paying customers, not just light up buildings. The technology options and ramifications of each are discussed in the next section.

Building and Access network requirements

To deliver broadband IP services to small and medium enterprises in multi-tenant commercial buildings, the building equipment must address many requirements:

At the building, service providers need a low-cost entry-level platform that has the port density and performance to scale as more subscribers sign up for services. Service providers also need to provide scalable per-subscriber bandwidth to meet growing demands as subscribers take up more value added services, or increase their Internet access traffic. Service providers should choose solutions that would allow video-conferencing or streaming video applications from the start. The last thing they need is another truck roll to upgrade equipment. And of course all subscribers need isolation and privacy from one another. Additionally, to avoid prohibitively expensive rewiring, the solution should leverage existing wiring within the building risers, to each tenant's suite.

On the wide area network (WAN) access side, in the absence of fiber to the curb in most areas, service providers must rely on traditional WAN links (T1/E1 or T3/E3) for backhaul to their metro POP. The fewer components required to achieve this, the better, and the simpler it is the better. Complex WAN configuration costs both in time and in requiring more experienced systems engineering resources. Since multiple subscribers will be sharing the backhaul bandwidth, so whichever WAN backhaulis selected, quality of service (QoS) and bandwidth management at both the building and the metro POP are critical for success.

These requirements, the currently available solutions, their shortcomings and Extreme's solution are described in detail in the following sections.

Challenge 1: Cabling versus bandwidth scalability dilemma

In order to minimize cost of entry, service providers want to leverage the existing in-building wiring infrastructure – rewiring a building with data-grade category 5 cable and fiber for runs over 100 meters costs about $500 per drop. That can add up to 10s of thousands of dollars in a large building, for example the Empire State Building cost about $1 million to rewire. And to maximize the long-term services revenue opportunity, they also need to be able to scale bandwidth to higher speeds on demand. Who knows what services they might be able to offer in years to come, voice services are an obvious target, while two-way and streaming video also offer high-margin revenue potential. Who knows how much bandwidth these will require and how soon it will be needed. A typical business-class video conferencing application requires about 400Kbps full duplex, while an MPEG2 video stream will take about 6Mbps. Of course, for business the bandwidth available needs to be symmetric because customers may be hosting web-servers or running communications applications. These requirements dictate a solution that is both past and future-proof – a tall order indeed!

These combined requirements create a major dilemma. The optimum solution for dealing with legacy in-building wiring is to leverage DSL technologies. But the architecture of most mini-DSLAMs is not a high-speed packet switching fabric, but a circuit switching multiplexor. In contrast, the optimum solution for providing low-cost scalable bandwidth is a multi-gigabit packet switching fabric. The decline of Token Ring and asyncryonus transfer mode (ATM) in the local area network (LAN) in favor of switched Ethernet are a testament to that. But Ethernet has limited range over copper, especially voice-grade copper and pulling fiber in the riser for each suite, is cost prohibitive.

	Subscriber Bandwidth	System Scalability	Traffic Management	Building Wiring
Ethernet Switches	10/100/1000 Mbps	Multi-gigabit Switch Fabric	Advanced L2/L3/L4	Cat 5 Fiber
Mini-DSLAMs	2-10 Mbps	< Gigabit Switch Fabric	Limited L3/L4	Cat 1, 2, 3

Ethernet switches scale well, but lack reach, Mini-DSLAMs have long reach but don't scale.

The Extreme Solution: Riser Independent Subscriber Provisioning

Extreme Networks'® solution combines Ethernet switching technologies and very high bit-rate DSL (VDSL) to elegantly address both of these requirements in the form of an Ethernet DSLAM. This provides a riser independent transport solution that leverages data-grade Category 5 cable and fiber where it exists, but also provides the extended reach of VDSL over the voice-grade cabling that exists in all buildings.

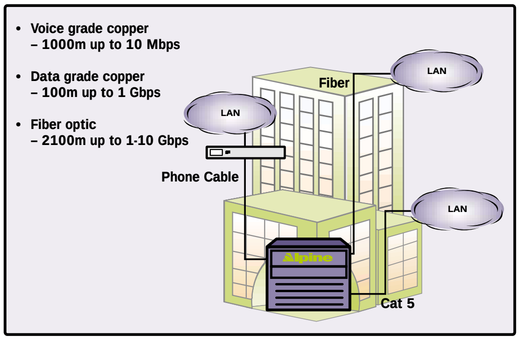

This eliminates any need to rewire the building while still offering a minimum 10Mbps symmetric bandwidth to each subscriber – over four times the subscriber bandwidth available with SDSL.

Combining VDSL and Ethernet switching provides an in-building switching platform that can scale to meet bandwidth demands and the subscriber densities to fit any building requirement.

One never knows what you'll find when you enter a building. One thing that service providers can do without, is needing a different solution depending on what media is available or how many subscriber connections are needed. That can cause havoc in terms of managing inventory and spares and places a greater burden on your engineers and support staff.

Long-range Ethernet using VDSL

VDSL provides the highest speed of any DSL technology for distances up to 3000 feet. Originally conceived by Telecommunications carriers as a "cable killer", suitable for delivering digital video broadcasting and video on demand over the standard telephony infrastructure, VSDL is also proving a strategic choice for high-bandwidth broadband access.

Capable of operating in symmetrical mode at speeds of 10 Mbps and asymmetrically with downstream speeds up to 52 Mbps, VDSL is the next step up the speed ladder beyond ADSL and SDSL. But, unlike rival DSL technologies, VDSL offers the bandwidth scalability to meet subscriber demands far into the future, making it ideal for business subscribers in multi-tenant buildings, business parks and hospitality as well as office buildings within 3000 feet of the CO.

For service providers deploying broadband access in MTUs, VDSL represents a sound strategic investment choice that will not have to be revisited for many years to come - VDSL technology lends itself well to providing long-range Ethernet connections with up to 10Mbps symmetric bandwidth over phone-grade in-building wiring as well as short-run phone cabling in the first mile. Essentially, VDSL can be used to deliver IP/Ethernet connectivity transparently over telephone grade wire at up to 10 times the distance of traditional 10/100 BASE-T on Category 5 cabling.

	Subscriber Bandwidth	System Scalability	Traffic Management	Building Wiring
Ethernet DSLAM	10/100/1000 Mbps	Multi-gigabit Switch Fabric	Advanced L2/L3/L4	Cat 1, 2, 3, 5 Fiber

Ethernet DSLAMs combine the scalability of Ethernet switches and the reach of Mini-DSLAMs.

Challenge 2: Leveraging available WAN access options

As mentioned earlier, to gain "rights of entry" to an MTU, service providers must strike agreements with property owners. Property owners may own buildings in a variety of metropolitan service areas (MSAs) in one or more geographic regions. Over the last twenty years virtually every city, even the most densely populated cities in Europe and certain South East Asian countries, has witnessed urban sprawl. In the US the percentage of office space located inside the core central city area or central business district (CBD) dropped from 66% in 1979 to 47% in 1999. This has important implications for availability of different WAN access options. Within the CBD it is estimated that 10-15% of buildings are in reasonable reach of a fiber splice point, but unless fiber is already available to the curbside, it may take 3-6 months to get fiber to the building. Outside the CBD fiber is rare.

For large buildings fiber is obviously the most attractive option, it offers almost unlimited scalability at a much lower cost per megabit than unbundled copper loops of T3/E3 or T1/E1. These days in the US one can lease dark fiber from a wholesale metropolitan fiber provider for around $1000 a month on a long term lease, while a T3 will cost as much as $3000 from the local Telco. However if the aggregate bandwidth demands are low, for example in small buildings with only a handful of subscribers, the minimum cost of a fiber link is still much higher than an unbundled T1 or E1 from the Telco, which may cost only $300 wholesale.

The dilemma is this: You're going to be serving small buildings and large buildings. You're not always going to have access to fiber or even T3/E3 when you want it, and even when it is available, you may not have the bandwidth demands to justify the bigger pipes and will need to start with T1/E1. If you limit your addressable market only to large buildings with fiber access, you'll have limited your market to 10% of its potential size, plus you'll be butting heads with competitors all the time – of course, those are the prime sites that everyone wants - not low-hanging fruit.

What all this means is that you must be able to implement a solution that scales, using whatever WAN access technology fits, and you shouldn't have to architect a different solution for each scenario. You'll be amazed how many vendors offer one platform for when T1/E1 WAN links are used and a completely different platform for when dark fiber is available - even changing architecture from ATM to Ethernet. Managing your business is a big enough challenge so dealing with the complexity of two different architectures is something you don't need.

Most metropolitan fiber build-out is of course targeted at large buildings that currently have T3 or E3s because these are prime candidates for upgrades to higher capacity. And with the wholesale price differential between T3 and dark fiber being 3:1, it only takes bandwidth demands of 15-20Mbps for Fiber to make sense. Hence it's likely that large buildings that don't currently have fiber access will have it over the next few years. So, by making pragmatic technology choices at the beginning, you'll be able to further reduce costs by upgrading buildings that originally needed T3s to fiber when it becomes available.

The Extreme Solution: Consistent architecture with great WAN flexibility

Unlike most other competitive solutions that require an additional router or a WAN aggregation switch depending on what and how many WAN links are required, Extreme has a one-box solution. With 4-port T1 and E1 modules, 1 and 4-port T3 modules and a variety of long and short-range GBIC modules for Gigabit Ethernet, you can deploy a consistent service model no matter what WAN links are available to the building. For bandwidth scalability on low speed WAN links, the T1/E1 ports can be bonded together using multi-link point-to-point protocol (MLPPP) to scale from 1.5Mbps to 6Mbps (T1) or from 2Mbps to 8Mbps (E1). And because you can use the same platform in every building, if you need to jump from n*T1/E1 to T3 or Gigabit Ethernet, you can simply re-use the module in the next building requiring T1 or E1 links.

Challenge 3: Packet versus circuit switched WAN backhaul

Most solutions being deployed today in MTUs utilize SDSL or HomePNA (Home Phoneline Networking Alliance) technologies to take advantage of existing voice-grade wiring inside buildings. In other words service providers are most commonly deploying mini-DSLAMs in the building basement. Most DSLAM architectures are really circuit-switched access concentrators rather than packet switches and this typically means they use ATM or Frame Relay for the backhaul to the POP. Usually, each subscriber line is mapped to one or more permanent virtual circuits (PVCs) that originate at the CPE and terminate on an ATM switch at the metro POP. This is a layer 2 service model end-to-end. To give applications preferential service, it becomes necessary to provision additional PVCs – usually one for each application that should be treated differently from others.

While the use of PVCs is excellent for subscriber isolation, they do complicate solution deployment. Not only is circuit switching expensive to implement because it takes time and experience to configure PVCs, and requires meticulous tracking of virtual circuit label assignments end-to-end, it is also very costly back in the POP. ATM switches that scale to support thousands of PVCs are required, in addition to IP routers for Internet traffic handoff. Furthermore, if a service provider has multiple metro POPs and wants to

interconnect them, it usually becomes necessary to use an ATM or synchronus optical network (SONET) backbone, rather than exploiting the simplicity and economics of Gigabit Ethernet metropolitan rings. Few vendors provide all the pieces for both the POP and the MTU as an integrated solution, with important implications for end-to-end service provisioning. Different element management and provisioning tools are needed to
configure the various boxes and assign subscriber QoS and bandwidth profiles end-to-end, adding greatly to the overall complexity and administrative burden of turning-up subscriber services.

Some vendors have attempted to improve the flexibility and intelligence of their mini-DSLAMs by aggregating "same-class" traffic from multiple subscribers onto a single PVC over the WAN. This is designed to reduce WAN complexity and PVC scaling issues. To achieve this requires a packet classification engine in the mini-DSLAM. Since mini-DSLAMs are really concentrators and not packet switches, the classification is typically done in software instead of hardware, as is the case with Ethernet Layer 3 switches. The performance impact is disastrous, with latency being added every time packets are converted to cells or vice-versa.

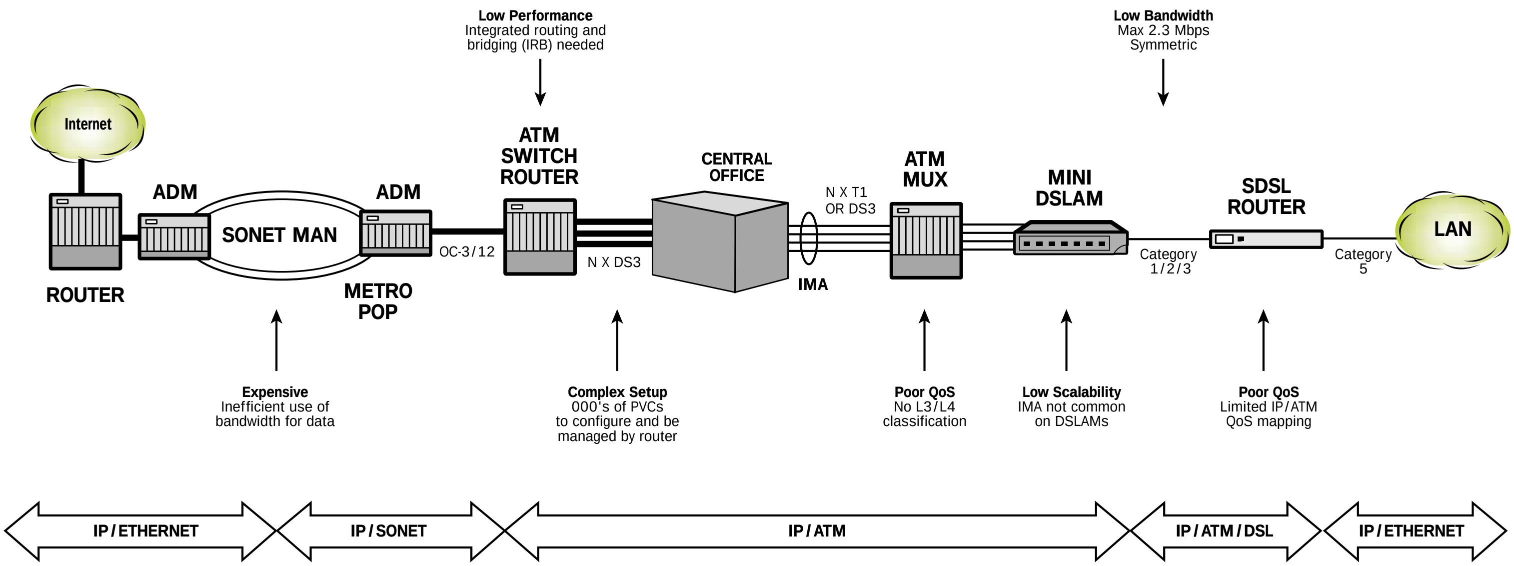

The hard way to provision services – ATM and PVCs means lots of configuration complexity.

This approach also makes it very difficult to control how much of the shared bandwidth each subscriber gets. A typical data flow looks like this: Data is converted from packets to cells at the IAD, back to packets at the mini-DSLAM for classification and back to cells over the WAN and eventually back to packets at the POP router for Internet handoff.

The Extreme Solution: IP/Ethernet packet switching everywhere

Extreme's solution is based on the Alpine™ service provisioning switch – a. platform that combines powerful hardware-based packet classification with Layer 3 switching, advanced QoS and bandwidth management features, all made possible by a multi-gigabit switching fabric. It provides a packet provides a packet switched WAN deployment model that takes minutes, not days, to configure. Extreme's approach preserves native Ethernet framing over VDSL within the building and over the WAN. This eliminates the arduous, complex task of configuring PVCs and can be implemented by less experienced, less expensive engineers.

There is no format translation required as packets are transmitted over telco WAN links such as T1 / E1 and T3 or Gigabit Ethernet. This ensures very low latency and enables consistent QoS and bandwidth management end-to-end. Even 802.1 Q/p tags are preserved over the WAN allowing subscriber isolation using Virtual Local Area Network (VLANs.) With 8 queues per port, hardware-based packet classification, and jitter-free rate shaping, Extreme's solution is more than a match for any ATM based mini-DSLAM solution. The Alpine platform offers the scalability and raw packet forwarding performance to sustain per-subscriber bandwidth provisioning from 200K to 100 Mbps and differentiated service for any number of subscribers and applications.

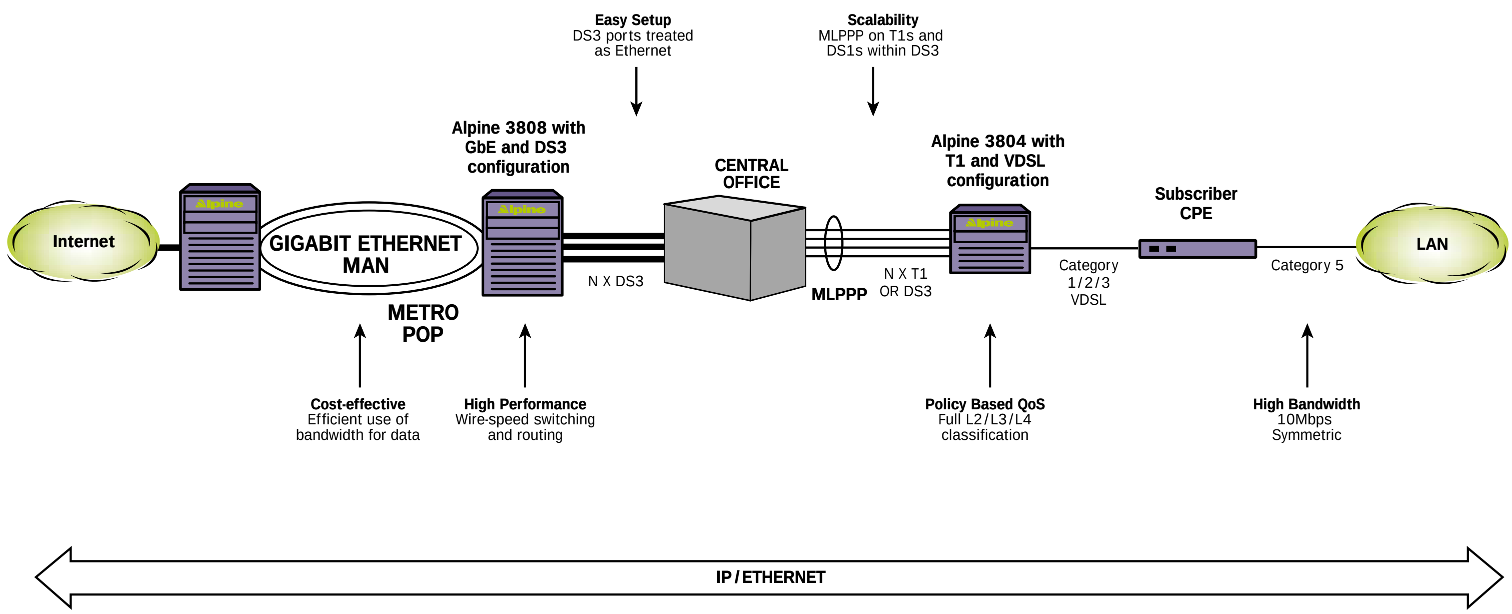

he easy way to provision services – Ethernet or IP end-to-end means simple configuration.

Extreme's approach for transporting Ethernet frames over telco WAN links utilizes bridge control protocol (BCP) over point-to-point protocol (PPP). This is completely standards-based, and offers proven interoperability with other vendors. However, as this is relatively new, there will be many scenarios where a service provider prefers to leverage an existing IP router at the POP to aggregate MTU WAN connections. In this instance, Extreme's solution also supports internet protocol control protocol (IPCP) over PPP, which is widely supported in the industry, including Cisco's IOS.

Challenge 4: Choosing between Layer 2 and Layer 3 service models

One of the major challenges facing retail IP services providers is the shortage of public IP addresses. So much so, that many service providers are actively buying back the IP addresses they gave out to customers years ago, before the problem became so acute. With CO-based DSL deployments the problem is relatively manageable. In so far as one DSLAM used in conjunction with network address translation (NAT) can represent hundreds of subscribers with just a few public IP addresses. However, when deploying a “building POP” the number of subscribers in each building is an order of magnitude smaller than the number aggregated on a DSLAM in the CO. So service providers need to pay serious consideration to implementing an addressing scheme that provides subscriber isolation while also conserving IP addresses. It’s all about where to apply Layer 2 and Layer 3.

With many vendors’ solutions, there isn’t much flexibility over where Layer 2 and Layer 3 service models can be applied. In a typical DSL/ATM solution, there is no option other than a Layer 2 model from the subscriber port, right through to the POP. Worse still, as already mentioned, most vendors’ solutions vary depending on the WAN technologies available, the in-building wiring available, or the subscriber density required. This makes it virtually impossible to deploy a consistent service model for every building. That’s another deployment and configuration nightmare that will impact both time-to-service and after sales support - yet another strain on your resources and bottom-line.

The Extreme Solution: Complete flexibility with IP/Ethernet service models

Extreme’s approach offers total flexibility over where to deploy Layer 2 and Layer 3, and since Extreme offers a uniform solution, no matter what the MTU conditions, you can implement a consistent service model across your entire customer base. Because the Alpine platform is a Layer 3 switch, Extreme can support either Layer 2 or Layer 3 end-to-end from each subscriber port through to the metro POP. We even support a hybrid model with Layer 3 switching implemented at the building edge.

	Extreme Deployment Models			Other Vendors
	Switched LAN Switched WAN	**Switched LAN Routed WAN**	**Routed LAN Routed WAN**	**ATM or F/R PVC based**
	L3	L3	L3	L3
WAN	L2	L3	L3	L2
Building	L2	L2	L3	L2
Deployment Simplicity	√√√√	√√√	√√	√
Traditional Architecture	√√	√√√	√√√√	√√
IP Address Conservation	√√√√	√√√	√√	√√√√
WAN Payload Efficiency	√√√	√√√	√√√√	√

Legend

Best	√√√√
Better	√√√
Good	√√
Bad	√

And of course, none of these models obviate the use of L3 / L4 classification for QoS and bandwidth management. In other words, you can leverage Extreme's wire-speed L3 / L4 packet classification to enable preferential services and govern bandwidth even when a Layer 2 service model is implemented. So you can choose any model you like - route at each subscriber port, route at the building POP or route at the metro POP. Each service model has its merits and trade-offs – so we've left the decision over which one is best for you, in your hands.

Challenge 5: End-to-end Subscriber-Level Service Provisioning

This is really the most complex aspect of broadband services deployment. It's not just about having QoS and bandwidth management capabilities in the devices, one also needs the tools to define service profiles and enforce them. These are two sides of the same coin. Everyone accepts that for any real-time applications such as voice and video communications QoS is prerequisite. But let's not forget the management aspect, defining service profiles and turning-up services, takes time and money. Doing it the hard way using command line interfaces (CLI) is far too time consuming on a large scale. You'll soon need provisioning tools that allow you to define service profiles and apply them to all the necessary devices in the data path in one fell swoop.

The Extreme Solution: Advanced Traffic Engineering and Provisioning Tools

Extreme's QoS and bandwidth management capabilities are among the most advanced in the Ethernet switching industry, and are well documented in other Extreme white papers and technical briefs (www.extremenetworks.com/technology/wp.asp). In brief, the latest "i" chipset, which is employed in the Alpine platform supports eight priority queues per port. Packets can be classified and assigned to queues on the basis of a wide variety of attributes: Port, MAC address, IP Subnet, IP address, VLAN ID, 802.1Q/p tags, Diffserv code-point, and TCP/UDP port number. And it's all done at wire-speed in hardware with microsecond latency. Extreme's bidirectional rate shaping features minimum and maximum settings allowing you to control how much bandwidth applications or subscribers get. Bandwidth can be sliced to a granularity of 2% of port capacity, and with virtually jitter-free traffic shaping it can preserve the integrity of real-time traffic even over a narrow pipe.

There are many solutions that claim to deliver all of the above mentioned QoS and bandwidth management capabilities, but the operational cost and deployment complexity of implementing them is immense – because you need ATM to match Extreme's QoS capabilities. ATM is complex and costly enough, and addressing standard requirements such as IP routing and traffic priority classification at the

IP layer only compounds the problem. The cost of service provisioning escalates when disparate systems and technologies are involved. Hence, a non-uniform service delivery infrastructure requires a myriad of management applications to set policies and provision bandwidth – making it very complex to provision services end to end. Even assuming you can train and retain the experts you'll need.

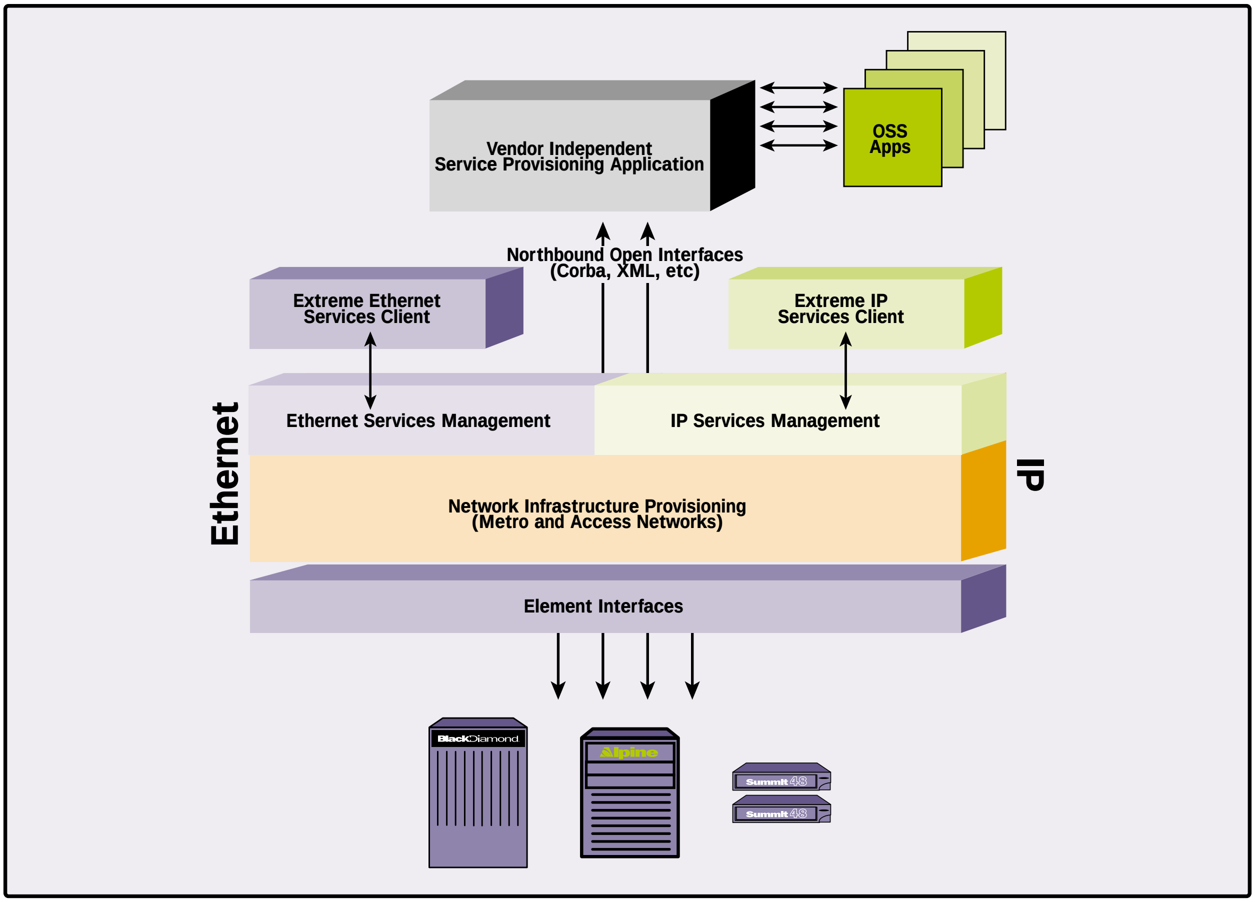

Extreme Networks service provisioning architecture.

Extreme understands the magnitude of the provisioning problem and is delivering practical service provisioning tools to facilitate infrastructure build-out and service turn-up. First of all, we recognize that service provisioning is only a small part of the entire OSS infrastructure and are working closely with partners that offer complimentary components. We're being realistic, and focusing on how to simplify the infrastructure provisioning and facilitate service creation on that infrastructure. Our service provisioning management tools for metro and access will provide northbound APIs (Corba, XML, etc.) to more comprehensive provisioning applications. Through the enablers we build, partner vendors focused on service provisioning will be able to optimally manage Extreme equipment.

Summary

To bring this to a conclusion, the main messages are: In the retail services business, marketing to SMBs, don't rely on getting lots of value added services. There is not a good track record of high take-up rates! You need to be able to make money on Internet access and only 50% additional revenue from other services revenue. That's virtually impossible when you rely heavily on incumbent providers for local loop access and co-location, and harder still if your technology choices give you a similar cost structure to that of the incumbent. To break the cycle you must focus on markets that allow you to reduce or eliminate dependence on the incumbent telco. That's what metro providers are doing by leveraging dark fiber from utilities and selling wholesale bandwidth. Likewise, competitive providers selling retail services should target multi-tenant buildings. The MTU market lets you deploy POPs in buildings and sell services to tenants in each building. Your costs are lower, and you don't deal with the incumbent telco as much. Best of all, you give less of your revenue away to your archrival the incumbent telco. The right technology choices can give you a cost model that is disruptive, enabling profitability in two years. Ethernet is that disruptive technology.

---

[1]Access to Telecommunications Services – Commercial Tenant Findings

Knowledge Systems and Research, Inc & Strategic Policy Research Feb 2001

[2]Center on Urban & Metropolitan Policy - Office Sprawl: The Evolving Geography of Business

3585 Monroe Street Santa Clara, CA 95051-1450 Phone 408.579.2800 Fax 408.579.3000
Email info@extremenetworks.com Web www.extremenetworks.com

L-WP-MTU-106