Alleviating the Pain in Today's Edge Networks

TeleChoice, Inc.

1307 S. Boulder Ave., Suite 120

Tulsa, OK 74119

918.382.0007 TEL

918.382.0033 FAX

www.telechoice.com

408003

Table of Contents

The Strategic Catalyst TeleChoice for the Telecom Industry

Executive Summary

All-optical networks have delivered tremendous new capacity in the core network. Aggressive broadband deployments are providing incredible service capabilities in the access network. Unfortunately, the edge network has failed to effectively bridge the chasm between the two. Carriers have been unable to effectively access all the new bandwidth that is available. Service providers have been forced to try to fill the chasm with a complex array of point solutions resulting in many overlapping layers of functionality and capability.

This edge network wilderness is painful for service providers to navigate. As their customers demand new services, service providers struggle with implementation pain, integration pain, provisioning pain, operational pain, and cost and investment protection pain.

These pains have real financial and operational impacts. Investors are demanding capital efficiency, but service providers may be facing even more critical shortages in basic commodities such as real estate, power, and people.

Thankfully, technology advances hold promise in alleviating these pains. Vendors are beginning to integrate these technology advances into solutions that appear to effectively address the real pains of today's edge network. We believe it is technically feasible for a vendor to develop an optimal edge network platform that can span from the edge of the access network to the edge of the core network and support all the services required of carriers today and in the future. This single vendor solution could eliminate much, if not all, of the pain service providers are currently suffering.

We have developed a simplified model of all of the layers in the edge network. Based on this model, we believe that once this optimal edge network becomes available, service providers will realize real, order of magnitude benefits:

- Reductions in rack space requirements of up to 82%

- Reductions in power requirements of up to 83%

- Acceleration in service provisioning time of up to 21%

For a growing network, these issues are critical. Service provisioning represents a very high cost to service providers today—achieving a 21% reduction in those costs would be significant.

This vision is within our grasp. We are aware of a number of companies leveraging one or more of these new enabling technologies. Gotham Networks, in particular, is integrating many of the key technology advances into a new network architecture that can flexibly reach from the edge of the access network to the edge of the core. Its GN 1600 Switchless Switch™ platform supports any protocol on any interface port and its innovative architecture enables breakthrough scalability and simplicity in managing the edge network. Gotham is planning the initial release of its product in the first half of 2001 and plans to continuously build upon the architecture to move towards the optimal solution we have envisioned and described.

How We Got Here

In the beginning were dedicated data services. Dedicated data services were offered over copper lines terminating in relatively simple multiplexing equipment, aggregated into higher level circuits, carried across a relatively simple network, and easily managed from end-to-end.

Life was simple, ordered, and manageable.

As data services evolved, there came frame relay, SMDS, ATM, and VPN services. Then came the Internet, which forever changed service provider networks. No longer was there order and hierarchy. No longer was there simple manageability and ease of provisioning. No longer were there simple billing plans and single service customers. No longer was there clear, end-to-end testing and predefined restoration plans. No longer was there simplicity, clarity, and certainty.

But there were new services, new bandwidth consumption, and much new revenue. Life was good?

With new services came demand for more bandwidth and more bandwidth and more bandwidth. Then came SONET and DWDM and ATM core switches, pushing the old ATM switches to the edge of the core. Then came IP core routers, pushing the old routers to the edge of the core. Then came bigger ATM core switches and bigger IP core routers, pushing the old switches and routers out yet another level from the core.

With more bandwidth came more revenues. With new technology came lower costs per bit transported. Life was good?

Then came cable modems, DSL, ADSL, SDSL, DSL.lite, and… There were CMTS and DSLAMs and high-density access aggregation devices. Then came new service switches supporting VPNs and security services. And then came subscriber management systems and service management systems and, finally, high-density core aggregation devices. With more access capacity and services came more customers subscribing to more services with more service plans.

With more subscribers to more services came more revenue. Life is good?

No, life is not good—not for service providers. Despite the upsurge in access bandwidth and the explosion in core optical bandwidth, a deadly chasm has opened between the two. Service providers have been forced to pour more and more components into the edge network chasm to try to build a bridge and provision services for their customers. It is not working. Service provisioning intervals are painfully long (up to six months for some services). Lack of integration between platforms results in unpredictable network impacts as new services are provisioned. Relatively simple network problems become PR-nightmare extended outages. Service providers are struggling to compete and struggling to keep customers satisfied.

Instead of a bridge linking the paradise of broadband access networks with the glory of all optical core networks, the edge network has become a wilderness that is difficult, painful, and slow to navigate and filled with uncertainty.

Life on the Edge

Rapid advances in technology and unparalleled network growth are forcing service providers to continuously modify and update their network architectures. Lehman Brothers estimates that service providers will spend over $100B in capital this year. Investors and other suppliers of capital are increasingly skeptical that these investments will provide meaningful returns.

Much of this capital is being spent to migrate to the "next generation network." This packet-based infrastructure holds out the "converged" network carrot, or one single network from which all services can be delivered thereby generating new revenue opportunities and reducing total capital and operational expenses. A recent JPMorgan report ("Backbone!" published September 8, 2000) projected potential capital savings of over 50% and operating cost savings of nearly 25% from deploying a next generation network.

Figure 1: Indexed Growth in Access and Core Networks

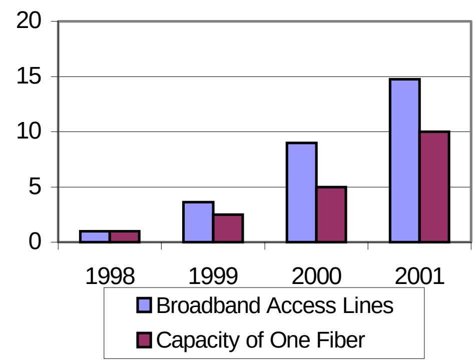

Sources: TeleChoice, KMI, Insight Research

However, for many reasons, the full benefits of the next generation network have not been realized. Although there have been meaningful technology advances in both the electrical access network and the optical core network, a chasm has formed between the two. Figure 1 demonstrates the growth in capacity in optical networks (a 10x increase) and in deployment of broadband access lines (nearly a 15x increase) between 1998 and 2001. The edge network has not kept pace with this dramatic growth. Today's provider has had to utilize a layered approach to network design and service delivery to offer next generation services. Each network layer adds complexity to network implementation, integration, provisioning, and operations…not to mention escalating capital and ongoing expense costs.

What has been created is an edge infrastructure "wilderness" of layers between the electrical access network and the optical core. This wilderness is a bottleneck to unleashing the true power of the advances in optics. Ongoing developments at the edge—the key link between the access and core networks—are fundamental to the ongoing realization of the next generation network.

Figure 2 represents a simplified model of the typical current edge network. Layers of equipment have been implemented in two dimensions. First, specific types of equipment must be implemented to support the different protocols and services supported in the network. Second, different layers of equipment must support several levels of hierarchical aggregation from low-speed connections up to the optical-speed connections at the edge of the core network. This diagram is overly simplified, for example, by representing only one layer of routing or switching at the edge of the access network. Often, multiple layers are required to accommodate low-speed to high-speed aggregation and as a result of constant introduction of the latest and greatest technologies into the network.

Figure 2: Today's Edge Network (Simplified)

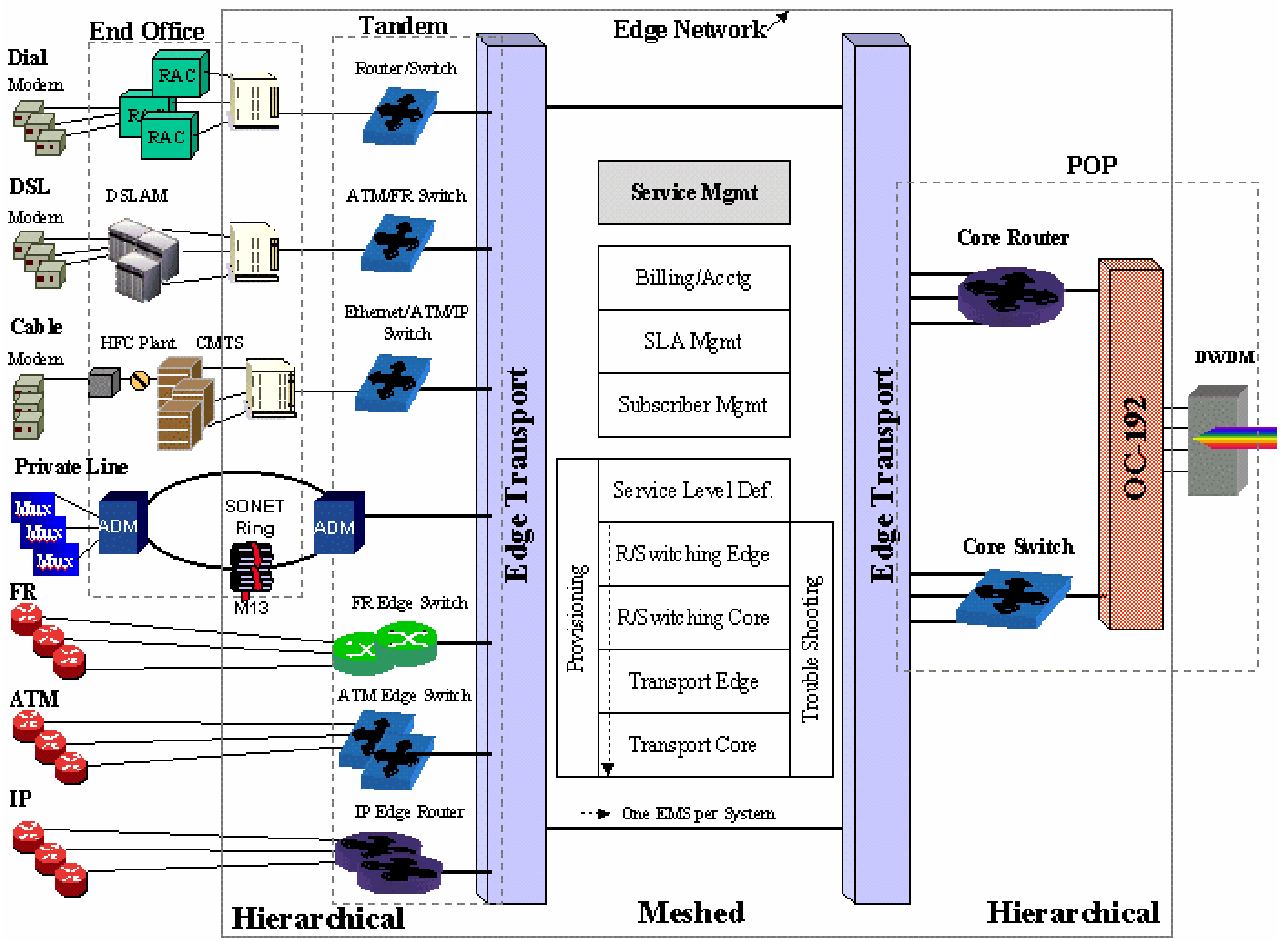

The Painful Reality

Service providers have hit the wall in trying to keep up with this growth and complexity. There is not enough space, power, or people to continue down the current path. We have identified five specific areas of pain caused by the current heavily layered edge network architecture:

Implementation Pain. It is hard to implement new services and new technologies. The cost of sending technicians out to a network location is very high ($150 to $500 for a simple truck roll) and the scarcity of qualified staff results in delayed implementations.

Integration Pain. The pain of integrating a new service or technology into the network is even greater than that of just implementing it. The complexities of the interworking issues are significant to the point that the implications are unpredictable. Increasingly, service providers are in "plug-and-pray" mode, hoping that everything will work together without any network catastrophes.

Provisioning Pain. Once a new service is implemented and integrated, service providers must be able to provision customers onto the new service. In the "good old days" of a single vendor, single platform solution, the vendor's network management system managed the mapping of service parameters to equipment settings and capacity allocation. That is no longer possible in today's complex edge network. The provisioning and planning groups must carefully manage the implications of service level provisioning on all the other layers in the network. Unfortunately, this is often more feasible in a reactive rather than proactive manner.

Furthermore, the service provider must keep on top of the equipment inventory in the field for each service. In the best case, edge devices have different cards supporting different services and different functions, and the service provider must carefully manage which cards are currently deployed. Often, entire equipment shelves and racks must be dedicated to a specific function. In many cases, another expensive truck roll is required to install the right card to support the specific service being provisioned. This, of course, can add days or weeks to the provisioning cycle.

Operational Pain. Operating today's complex edge network is a nightmare. What should be simple tasks, such as network optimization and grooming, become almost unsolvable challenges due to the often-conflicting information in the diversity of vendor systems. Troubleshooting a simple network problem can easily translate into an extended network outage due to the lack of a comprehensive view across all network elements.

Service providers have attempted to overcome these challenges by throwing more people, more training, and more organizational layers at the problem, but there are no short cuts. Much of the information about the interworking of the different layers typically resides in the heads of a few critical people. There's no way for service providers to effectively multiply this knowledge or divide it among function, layer, or service specific departments. In a few cases, the results have been catastrophic, with poor communications and poorer understanding making bad network outages only worse.

Cost and Investment Protection Pain. All these pains introduce operational costs. Additionally, the need to buy multiple layers results in inefficient capital spending with multiple vendors and often results in redundant and overlapping functionality. The introduction of new technology often results in implementing a new layer of cost (capital and operations) into the network well before the existing layers are fully depreciated or fully utilized. Even if wise decisions are made in efficiently populating the right equipment, as carriers grow, they outgrow small-scale devices and must perform a forklift upgrade to a larger-scale device.

As a simple example of the scale of this problem, Tim Horan of CIBC evaluated WorldCom's 3Q00 financial results and identified that the company had capital expenditures for the quarter of $2.5B while only depreciating $850M. The company appears to be deploying new equipment three times more quickly than it is depreciating old equipment. Something is wrong with this situation.

Treating the Symptoms, Suffering the Side Effects

Service providers have been living with this pain for nearly a decade. To date, there has not been a comprehensive solution. To hold out in the face of pain, most service providers have been implementing partial solutions that help reduce the problem in the short term.

The most immediate need is always to maintain a competitive service offer and to satisfy customer demands. Point solutions from vendors have been the only available approach to meeting this need. "Boxes" are being deployed for aggregation, for service creation, and for service management. These are single purpose elements from various vendors inserted for a specific purpose. In addition to the implementation and integration issues, each brings with it its own network management system to be incorporated.

A slightly better approach is the recent trend in network devices that collapse two or more layers into a single product. Specifically, two breeds of the next generation edge switch have garnered the most attention in recent months, MSPPs (Multi-Service Provisioning Platforms) and core aggregation devices.

MSPP devices typically reside in the service provider's POP and focus on enabling a myriad of services from a single box via various service-specific network cards—voice, frame relay, ATM, DACS, Internet access, etc. The wide range of functionality and services delivered via such a platform is highly attractive to providers. There can be tremendous benefits, for example, to carriers seeking to expand geographic coverage. This one device can enable the full range of services most commonly required by the enterprise community while at the same time being managed remotely and with automated provisioning. However, this is only one step in the right direction. Typically, these devices still have service-specific cards requiring careful inventory management to efficiently manage capital while effectively meeting customer needs. These platforms also only fit into a single-scale position in the network—typically as an aggregator of access data.

A second breed of edge vendors has primarily focused on the aggregation component at the edge of the optical core network. These devices aggregate all forms of traffic from multiple, existing interfaces (e.g., TDM, frame/ATM, Ethernet, etc.) for transport across high-speed optical trunks. Again, this is a step in the right direction but does not address many of the real issues in the edge network.

Although a dramatic improvement, these solutions often result in, at least for the short term, the addition of new layers to the network. Their service-specific components still require discrete implementation and operations focus, and they generally fall far short of providing a comprehensive solution.

The Bottom Line

The net impact of the edge network wilderness on service providers has been chronic and debilitating:

- Slow competitive response in launching new services
- Delays in customer service provisioning resulting in delayed revenues and lost sales
- Increased operational costs and shrinking margins (e.g., Sprint reported operations costs in its long distance division increased from 14% of revenues for the first nine months of 1999 to 17% of revenues for the same period in 2000)
- Mushrooming capital investments with disappearing return on capital (Lehman projects that long distance carrier incremental revenue per dollar of capital spent will shrink from $0.370 in 1998 to $0.293 in 2001)
- Consumption of all remaining POP and collocate space and power, severely limiting growth
- Difficulties in troubleshooting resulting in increasing MTTR (Mean Time To Repair)
- Overworked and overstressed operations personnel
- Falling customer satisfaction and loyalty

The Building Blocks of a Solution

Good news appears imminent. Technology advances on numerous fronts over the past couple of years are enabling equipment vendors to deliver sophisticated solutions to service providers. Technology vendors are beginning to leverage these technology advances to solve the edge networking problems. As vendors begin to combine the benefits of each advance into comprehensive solutions, pain-free edge networking is coming into view.

Specific technology advances include:

Innovation	Description	Impact
Wirespeed Interfaces	Router and switch interfaces can now operate at full line speeds. Prior to this technology becoming available, input streams needed to be buffered on the input side of the switch interface.	Input buffering created a number of problems that can now be resolved. Most obviously, if the interface card could not process incoming traffic at full line speeds, then the switch had become a significant delay component in the end-to-end service. Input buffers also prevent efficient protocol-specific traffic management to handle prioritized traffic and to support quality of service levels. Wirespeed interfaces allow the switch to achieve its full potential in supporting ultrafast switching and transport while also supporting protocol and service level processing of individual packets.
Line-Speed Network Processors	Processors that are "programmable," as opposed to ASICs (Application-Specific Integrated Circuits) that are "hardwired" to perform just one function very efficiently.	Network processors allow carriers to purchase "generic" cards and simply enable the particular function they need performed out of that card (e.g., IP processing vs. frame relay processing) Benefits: Faster processing than traditional software-based solutions, more flexible than specialized ASIC processors, lower inventory costs than ASICs in complex networks.
Standard Protocol Stacks	World-class support for the most important data protocols, packaged as licensed software for integration into equipment vendors' products.	Allows vendors to focus their time and resources on developing truly innovative capabilities while providing world-class functionality for all common protocols and services.
Off-the-Shelf Components	World-class building blocks that provide standard functionality for equipment.	Provides highly reliable capabilities while allowing the technology vendor to focus time and resources on true innovation.

Innovation	Description	Impact
High-Density, High-Speed Memory	Lots of inexpensive, fast, and reliable memory in small packages.	Allows equipment vendors to load large amounts of data and software onto individual cards to provide much functionality and flexibility.
Optical Scale Switching Fabrics	Switching fabrics that operate fast enough to move traffic at optical speeds without blocking any traffic or requiring any buffering.	Elimination of buffering within the switch allows full protocol processing without impacting network performance.
Robust, Real-Time Operating Systems	Computer operating systems that support real-time applications with high reliability and failure recovery.	Provides enhanced reliability and responsiveness for all network capabilities.
Third-Party Operational Support Systems	World-class software systems for the operation of multivendor telecom networks.	Provides a stable and common OSS supporting the entire network, minimizing training and other people costs, and enabling the technology vendors to focus resources on their key areas of innovation.

The End Vision…A Civilized Edge

We believe that vendors will soon deliver solutions that directly address the critical issues in service providers' edge networks:

- Eliminate the current performance/service compromise
- Help service providers speed new technologies and services into the network
- Reduce the time and manual intervention required to perform simple network tasks, addressing the people shortage impacting all providers
- Eliminate today's complexities of traffic aggregation and service creation through individual components
- Collapse edge aggregation, service creation, switching, routing, and optical access into one flexible layer

The Optimal Solution

We envision an optimal edge solution as depicted in Figure 3. This network design eliminates all the service-specific and hierarchical aggregation layers in today's edge network. Although this architecture still must be implemented in three types of physical locations, the architecture represents a single consistent layer across all locations.

Figure 3: Optimal Edge Network

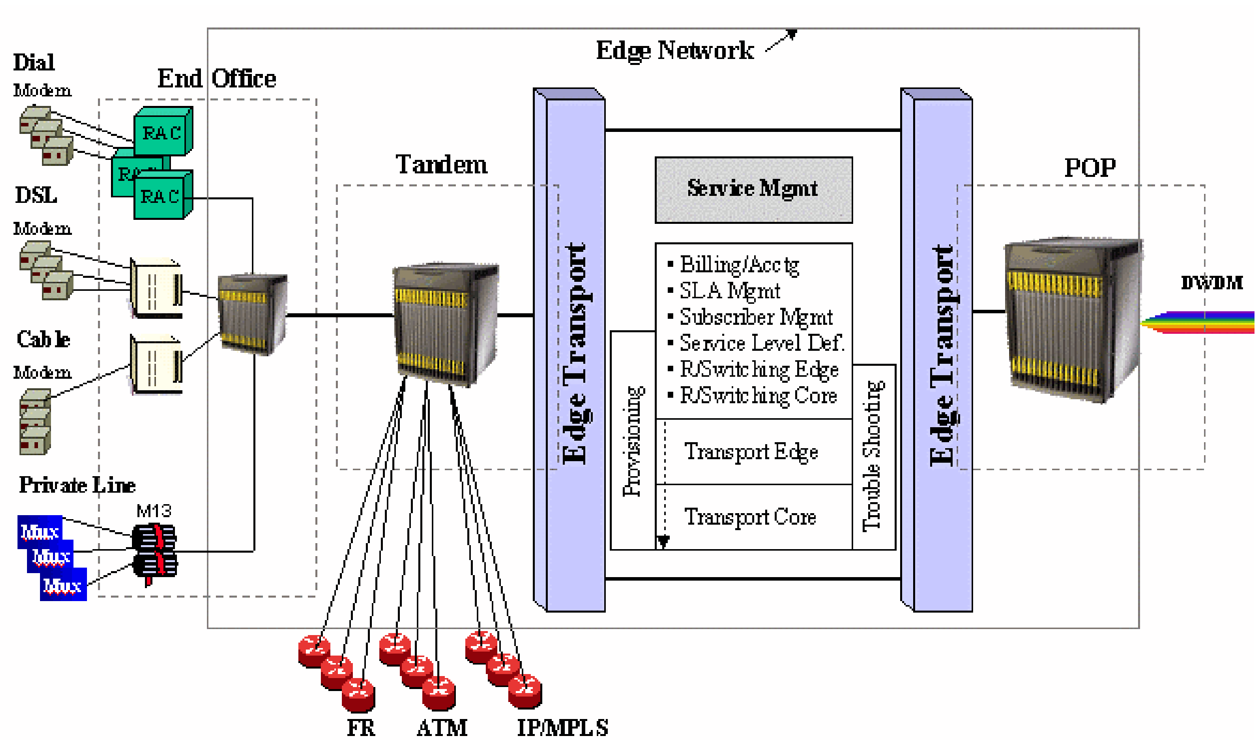

This optimal edge solution encompasses:

Single Vendor, Single Architecture. The optimal solution eliminates today's multivendor environment required across the edge of the network. This single solution will collapse layers within all three key areas—aggregation at the edge of the access network, service support, and aggregation at the edge of the core transport network.

Dynamic Programmable Protocol Support. The optimal edge network device will be capable of supporting any current protocol on any port, as well as allow that protocol to run in its native form throughout the network—from source point to end destination. As well, the network will have the flexibility to dynamically support future protocols via simple remote software upgrades. Service providers will no-longer need to carefully manage service-specific equipment inventories, but will be able to flexibly provision a service on any network card. This will result in fewer truck rolls, improved capital efficiency, better utilization of equipment and capacity, and quicker provisioning. Since any card can back up any other card, network reliability and restoration will also be significantly improved. Overall, service providers will make much more effective use of their capital investments and their POP space and power.

High Touch, High-Scale Service Level Support. New technologies will overcome the performance vs. packet-level processing constraint that vendors have traditionally faced. Breaking this compromise and enabling high-touch, processor-intensive services while maintaining the ability to scale to next generation environments allows for the removal of several layers/boxes at the edge. Rather than an aggregation device and service enablement device, the single vendor, single architecture platform accomplishes both by applying value-added touch functionality (VPN, QoS, traffic engineering, etc.) at line speeds.

Simplified Provisioning and Management. By implementing a single-platform solution across the edge network, many of the complexities of today's multivendor solutions will vanish. Provisioning will become the simple point-and-click exercise that has been promised for years. Interdependencies between aggregation, switching, routing, and transport will automatically be managed since all components will be part of an integrated solution. Troubleshooting will be similarly simplified, and since interdependencies have been eliminated, many problems will be automatically and immediately resolved by the system.

Scalability. Today, service providers must buy cost-effective small boxes for the access edge of the network and very large-scale boxes for the edge of the core network. Recent technology advances enable vendors to architect their products to scale from very low-cost access solutions to optical-scale core solutions.

In fact, given this optimal solution, we almost need a new definition of scalability that includes scaling from small to large, scaling across different protocols and interfaces, and scaling across the various functions that are performed in the edge network.

Addressing the Pain

This optimal solution would directly address each of the edge network pains that we have identified:

	Implementation	Integration	Provisioning	Operational	Cost/Inv. Prot.
Single Vendor/ Architecture	Implementation of new services merely involves upgrades to a single platform.	Fewer platforms means simpler integration.	Services are automatically provisioned across the entire edge net.	Reduced training. No conflicting management system info.	Focused vendor management. Highly leveraged capital.
Dynamic Programmable Protocol Support	New services are downloaded as software - no truck rolls.	New services are downloaded, not integrated.	Any card can serve any function. No truck rolls based on customer orders.	Flexible redundancy. Rapid restoration and recovery.	Much greater equipment utilization. Longer lifespan.
High-Touch, High-Scale Service Level Support	A single platform supports both fast and smart functions.	Both high-touch and high-scale functions are in a single box.	No longer need to worry about provisioning across multiple platforms.	Network latency and timing issues greatly reduced.	Fewer boxes. Fewer cards. Higher utilization. Greater reuse.
Simplified Provisioning and Management	Easier testing and faster time to revenue.	Interworking issues in the edge network are handled by software.	Point and Click. Really!	Automatic network restoration becomes possible.	Inventory visibility. Increased use of deployed capital. Less redundancy.
Scalability	No need for fork lift upgrades.	A single platform scales from the edge of access to the edge of the core.	No need to worry about scale-based provisioning issues.	Scalability in all dimensions greatly simplifies operations and support.	Cost scaling from small to large as network grows. No fork lift upgrades.

Modeling the Bottom Line Impact

Although each service provider's edge network will be unique, we have created a comprehensive model representative of a service provider's edge network. It encompasses components that reside in end offices, in tandem offices in the metro network, and in the core backbone network POP. It also includes the service management components from the element management systems up through service-level management, subscriber management, and the collection of ratable information for billing. We have focused on data services.

The components in our model are represented in Figure 2 (for the current architecture) and Figure 3 (for the optimal architecture). We have based port, card, shelf, and rack counts on existing equipment widely deployed in carrier networks today (for the current architecture) and on equipment being prepared for release by innovative vendors (for the optimal architecture).

In general, the primary function of the edge equipment in the end office and in the tandem office is to aggregate services into higher-density data streams. Tandem office equipment also implements much of the service-level capabilities. The primary function of the edge equipment in the core POP is to aggregate, shape, and groom the service traffic onto the optical core backbone network.

In general, the equipment in the end office tends to be service-specific. Using some existing multiservice platforms, it may be possible to support multiple protocols and services within a single equipment shelf although typically with each card supporting a single protocol/service.

For a typical service provider edge network supporting four different services, our model has identified significant measurable improvements in each of three areas: rackspace, power, and provisioning time.

Figure 4 reveals our specific edge network analysis. In summary, by implementing our optimal network architecture in our model network:

- Provisioning effort would be reduced by 21%
- Edge equipment rackspace required in each end office would be reduced by 84%
- Edge equipment rackspace required in each tandem office would be reduced by 67%
- Edge equipment rackspace required in each carrier POP would be reduced by 78%
- Throughout the entire edge network, rackspace would be reduced by 82%
- Power required by edge equipment in each end office would be reduced by 86%
- Power required by edge equipment in each tandem office would be reduced by 50%
- Power required by edge equipment in each carrier POP would be reduced by 75%
- Power required by edge equipment throughout the edge network would be reduced by 83%

Since each service provider's network will vary in some ways from our model, actual savings will likewise vary.

Figure 4: Edge Network Cost Model

Edge Network Cost Analysis

	Traditional Architecture			Optimal Architecture		
	End Office	Tandem	POP	End Office	Tandem	POP
Infrastructure Costs						
Service Platforms	4	2	2	1	1	1
DS-3 Ports/Card	2	6	48	12	12	48
Cards/ Shelf	16	16	10	16	16	16
Shelves/ Rack	2	2	1	3	3	3
Theoretical Racks/ Office/Platform	0.75	1.5	9.5	0.44	1.00	3.96
Actual Racks/Office/Platform	1	2	10	0.66	1.33	4.33
Fill Rate/ Last Shelf/Platform	50%	50%	50%	50%	50%	50%
Total Ports/Platform	64	288	4560	64	288	4560
Total DS3 Ports per Rack	64	192	480	576	576	2304
Total Racks/Office	**4**	**4**	**20**	**0.66**	**1.33**	**4.33**
Total Power per Office (Watts)	**16,000**	**10,000**	**80,000**	**2,222**	**5,000**	**19,792**
Total Provisioning Events per 100 Service Orders	**177**	**177**	**177**	**140**	**140**	**140**
Savings in Rackspace				**84%**	**67%**	**78%**
Savings in Power Consumption				**86%**	**50%**	**75%**
Savings in Provisioning Events				**21%**	**21%**	**21%**
Offices per Edge Network	50	5	1	50	5	1
Total Racks Required	**200.0**	**20.0**	**20.0**	**33.0**	**6.7**	**4.3**
Total Power Required	**800,000**	**50,000**	**80,000**	**111,111**	**25,000**	**19,792**
Total Edge Network Racks			**240**			**44.0**
Total Edge Network Power			**930,000**			**155,903**
Savings in Rackspace			**82%**			**82%**
Savings in Power Consumption			**83%**			**83%**

Specific Application Example: Out-of-Region Expansion

Most service providers are pursuing growth into new regions. As a provider enters a new geography, it faces a number of planning challenges. The objective, in implementing its network architecture in the new location, is to invest the minimal amount of capital while enabling a rapid ramp in revenues from customers subscribing to high-margin services. This is a real challenge with today's technology and network architectures.

Expansion Challenges

Service providers typically start building their new edge network with one shelf or chassis for a specific type of equipment and populate the cards in that shelf/chassis based on actual and projected demand. As the shelf or chassis approaches being fully populated, the next empty shelf or chassis is added. Each shelf consumes precious space in collocate space and POPs. Each shelf and each card also consumes precious power and generates critical heat, both of which are becoming key constraints in service provider network growth.

Since cards and equipment tend to be function-specific, today's service providers typically have very inefficient use of this precious space, since they have had to overpopulate with equipment to meet the relatively unpredictable demands of their customers. Especially as providers move into new geographies, the cost of over-populating shelves must be balanced against both the high cost of a truck roll to install additional equipment and the impact to cash flow of the revenue delays resulting from unnecessary provisioning delays.

Thanks to Murphy's Law, no matter how smart the service provider is, it will over-deploy equipment for under-demanded services and under-deploy equipment for services in high demand. The bottom line impact is unnecessarily high capital costs, rack space consumption, power consumption, heat generation, truck rolls and provisioning costs, as well as delays in provisioning and resulting revenue flow. Furthermore, since the provider will likely be training new operations personnel in the new region, the need to be functional with a wide range of equipment can introduce customer satisfaction challenges. These come at a time when it is critical for the provider to establish trust in the brand it is introducing into the new market.

Overcoming the Challenges

In our optimal scenario, today's complex architecture would be replaced by a single-vendor, single-platform solution which is implemented with a single solution in the edge office, a single solution in the tandem office, and a single solution in the core POP. Implementing our optimal system would provide significant improvements alleviating many of the costs and risks associated with expanding out-of-region with current architecture solutions.

Specifically:

- A single-vendor solution across the edge network will eliminate much of the operational complexity, speeding the provisioning process and simplifying troubleshooting. Technicians will only need to be trained on a single type of equipment, reducing risks that could damage the brand.

- Dynamic, programmable protocol support will allow any card in the network to serve any function, eliminating the need to over-populate the network with function-specific and protocol-specific cards in anticipation of customer demands. This will provide significant improvements in equipment utilization and deployment velocity, directly improving capital investments and revenue flows.

- Since a shelf of equipment will support any and all functions in the edge network, there will be consolidation of cards that today are spread across different types of equipment. This will create greater densities of card population in each level of the network. This will also have a direct impact on power consumption and heat generated.

The optimal architecture would provide significant improvements in all areas of pain described in this paper. We believe the bottom line impacts could be even more significant for an out-of-region expansion application than we have modeled, since our model assumes a relatively mature existing network with relatively high densities of provisioned services.

The Strategic Catalyst TeleChoice for the Telecom Industry

One Vendor's Solution: Gotham Networks

Is this a dream, or is it an achievable reality? A number of technology vendors are pursuing components of our optimal solution, and Gotham Networks may be closest to it. Gotham is introducing a product architecture directed towards the entire edge network vision and an initial product that makes strong progress towards that goal.

Gotham fully understands the pain service providers feel in their edge networks. In fact, Gotham takes it a step further, referring to the problem as the "friction" caused by too many network layers. In introducing its first product, the GN 1600 Switchless Switch™, Gotham calls it a Friction-Free Networking™ solution. Gotham's concept of network friction encompasses all the aspects of pain we have described, as well as the very clear reality that all the extraneous network layers slow down every operational process for a service provider.

Gotham's Switchless Switch™

Gotham's Switchless Switch™ leverages most of the breakthrough-enabling technologies we have identified. Each card in the switch is virtually a switch unto itself, using programmable line-speed network processors and world-class protocol stacks to support any service, any protocol on any port. Each card includes the processor and network fabric functions usually built as shelf-level or switch-level components. This allows the Switchless Switch™ to linearly scale from a very small switch to a very large switch. Not only does it scale in size, but it also scales in cost, with the service provider's investment directly related to the number of subscribers (see Figure 5).

Figure 5: Investment per User

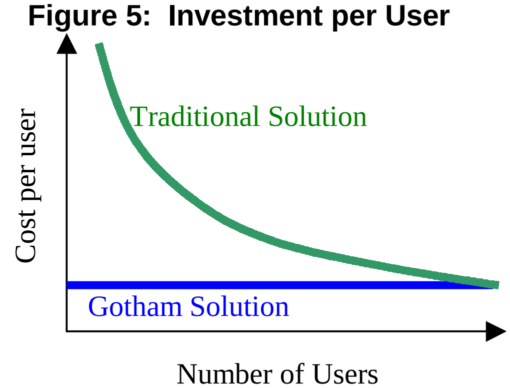

Source: Gotham Networks

This architecture has a number of obvious and subtle benefits. One of the interesting aspects of the product is that, not only are the cards and the physical ports protocol-independent, but in fact they can support multiple services simultaneously. For example, a channelized DS-3 physical interface could be supporting hundreds of different channels, running a wide variety of services/protocols, all supported by a single switch card. Another interesting component of the switch's scalability is the range of physical interfaces supported. Where a traditional edge network architecture may include two or more layers of routers, for example an edge-edge router supporting DS-1 to DS-3 level ports, and an edge router supporting DS-3 to OC-48 level ports, the Gotham Switchless Switch™ can incorporate both levels into a single device. Other clear benefits include the simplified management and provisioning, optimized capital leverage, and the resulting reductions in demands for rackspace, power, and people.

The Universal Service Card™

At the heart of Gotham's architecture is its Universal Service Card™. Since each card contains switching, control, and protocol processing and is dynamically programmable, the functions performed by each card can change on demand. This eliminates the need for service providers to ensure they have the right service-specific and function-specific cards deployed in the right locations. It also allows dynamic reconfigurations to support changing customer needs. This feature also allows any card to serve, on demand, as the backup for any other card in the switch. This results not only in 1:N redundancy for interface functions, but also 1:N redundancy for all software processes in the switch. This represents a potential exponential improvement in overall system reliability.

Each card serves not only as a service or protocol interface into the network, but can also perform higher-level protocol processing. Since each card has line-speed interfaces and fully programmable network processors, this feature breaks the traditional compromise between network speed and high-level protocol processing.

Gotham's solution turns many of the realities of today's edge network inside out. Think of it as creating an edgeless edge network. Because of its flexibility, it can easily integrate into today's network and support existing legacy applications. As service providers understand Gotham's capabilities, and as the company introduces additional features, it should be able to displace an increasing number of existing layers in the edge network, perhaps eventually becoming what we have described as the optimal solution.

Alleviating the Pain?

The pain that service providers feel in their edge networks is being caused by the over abundance of layers they have thrown into the edge trying to bridge the chasm from the access network to the optical core. This pain truly has become debilitating—keeping service providers from growing their networks and offering new services—potentially even threatening their fundamental ability to compete.

The stop gap approaches that have been available to date have only made matters worse, adding even more layers to the network. Service providers have been pouring capital into their networks and seeing questionable returns. Both investors and end users are losing patience although for very different reasons. Carriers have encountered critical shortages in the capital, floorspace, power, and people required to continue to fight this battle. It is starting to look more and more like a life-threatening situation, and it is time to call 9-1-1.

Gotham appears to be ready to answer that call, as it is really focused on alleviating the pain that service providers feel in their edge networks. If it can deliver fast enough, there should be a line of carriers waiting to see this doctor.

About TeleChoice

TeleChoice assists companies in creating new markets around innovative business models, technologies, products, services, and applications. As the strategic catalyst for the telecom industry, TeleChoice helps start or greatly accelerate the process of crystallizing a business or market strategy, value proposition, or differentiated position. Playing a strategic role, TeleChoice enables clients to launch new businesses, new markets, and new products and services rapidly and successfully.

Supporting service providers and the technology vendors that serve them, TeleChoice focuses on leading-edge public network technologies. Since being founded in 1985, we have been differentiated by our proven ability to transform new technologies into successful products and services. Our portfolio of offerings helps clients conceptualize, launch, market, and capitalize on innovations in networking—faster, more efficiently, and more profitably.