Building The Optimal Edge

Larry Hettick

Equipment vendors are on the right track, but carrier edge networks are still a big challenge.

Gigabit Ethernet LANs have brought broadband speed to the end-user's desktop. All-optical core networks have delivered tremendous capacity in the wide area. Aggressive broadband access deployments are beginning to connect the two and promise incredible new service capabilities. Unfortunately, the infamous "edge network" has failed to effectively bridge the LAN-WAN capacities. Nor have edge devices—either in the service provider network or on the customer premises—been able to facilitate the new services that providers hope to offer, or the simplified and converged networks they would prefer to operate.

Part of the problem is that although the edge remains the broadband bottleneck, the clearly-defined edge between enterprise and service provider networks—the old demarc—is vanishing. For example, devices like multiservice switches that once lived only in carrier central offices are also found at the customer premises. Now called "customer located equipment," these devices are still managed and controlled by the service provider. Similarly, servers, databases and storage equipment, once located solely at the customer premises, now also reside in providers' datacenters and central offices.

This blurring between premises and network-based functions has also complicated development of the networking products that directly target the carrier edge network. For the past few years, vendors have tried combining the many traditional functions performed by end office and tandem access equipment, but have not yet hit on the perfect solution.

Instead, multifunction and multivendor edge solutions have proliferated. Each boasts a complex set of often-overlapping functions, yet no single solution has made it easy for carriers to provide broadband access and provision new service features.

In fact, carriers often find they must adopt multifunction devices with unnecessary features to solve point problems. For example, one service provider's operations division installed SONET multiplexers in buildings just to do optical loopback testing—an expensive solution to a single point problem.

Put simply, today's edge solutions have fallen short of satisfying carrier requirements. To be successful, emerging solutions will have to meet three critical objectives:

1. Bridge the bandwidth bottleneck—between user LANs and the optical core.

2. Improve the serviceability of carrier networks—make it easier to define, provision, bill, and manage services and equipment.

3. Enable converged carrier infrastructures—to simplify carrier networks and support new end-user services.

Equipment vendors tend to address these problems based on different technology solutions. They position themselves as suppliers of Ethernet, PONs, multiservice provisioning platforms, optical multiplexers, IP routing and next-generation SONET systems.

Many of them claim to meet the first two criteria, including vendors such as Accelerated Networks, Amber Networks, Appian, Atrica, Celox, Cerent, Ciena, Ellacoya, Integral Access, Luminous, Mayan, Pluris, Qeyton, Quantum Bridge, Redback, Sentient, Shasta and Tenor. And packet-based Class 5 switch replacement devices—like those made by Telica, Oresis, Tachion and Convergent—can help address the third criterion. But a closer look at these three basic requirements shows that more work is still needed.

Build A Bandwidth Bridge

The networking world was once easily understood. Time-division and statistical multiplexers gathered customer traffic for additional, circuit-based aggregation through a stable hierarchy of edge, tandem and core switching offices in the carrier networks. But the introduction of overlay data networks—starting with X.25, then frame relay, ATM and the Internet—and the need to interwork these services have eroded traditional network borders.

While overlay networks proliferated, access and transport options also multiplied. Cable, DSL and wireless joined traditional modems and brought their own, high-density access aggregation devices. In the core, SONET transport has been layered over DWDM, adding capacity and producing a variety of vendor-specific switching, routing and management options. As bandwidth demands grew, yesterday's core switches became today's edge multiplexers, and the cycle to push core switches and routers to the edge in favor of bigger core switches and routers became a fundamental network planning paradigm.

Larry Hettick is vice president, consulting with TeleChoice (www.telechoice.com), a consulting firm that catalyzes change in telecommunications business and marketing strategies. He can be reached at lhettick@telechoice.com

FIGURE 1 Current Edge Environment

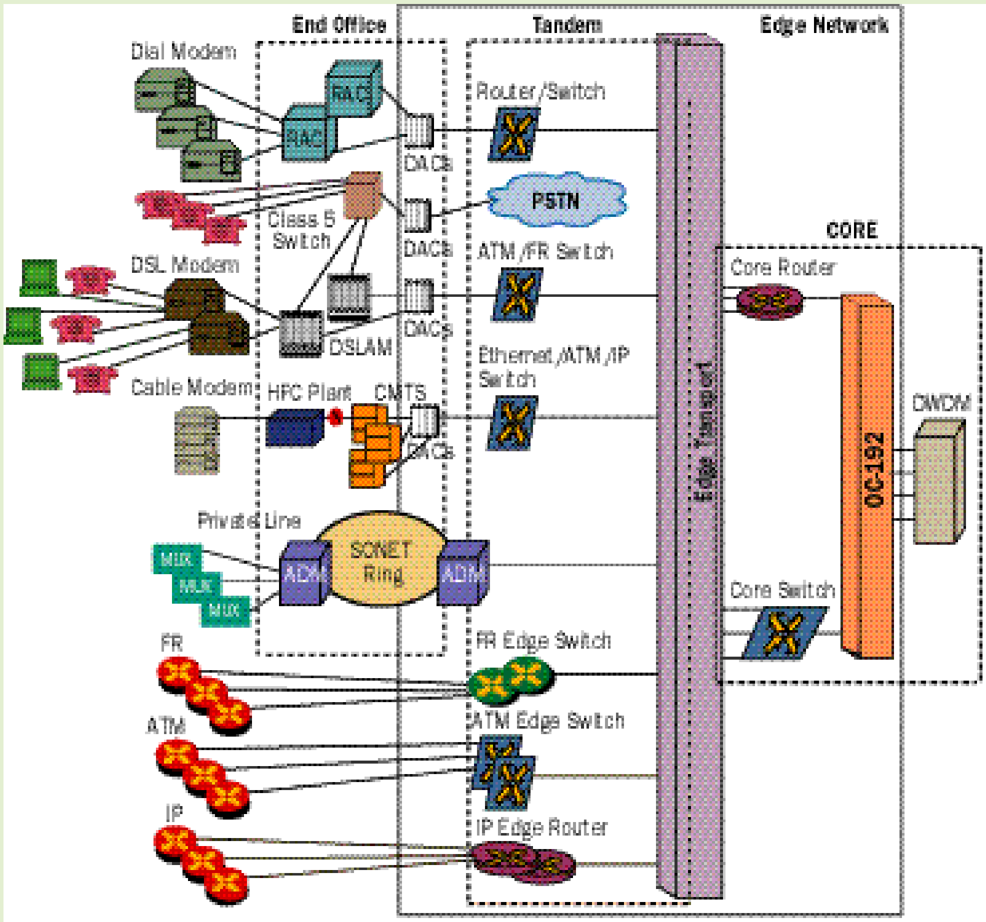

Linking the new edge and core is like rebuilding the Golden Gate Bridge without stopping traffic

Because so many different network elements feed traffic into today's network edge (Figure 1), building a bandwidth bridge from that edge into the core of the network is like rebuilding the Golden Gate Bridge without stopping traffic. Besides doubling the current carrying capacity right away, keeping the existing steel and adding more modern composites, the bridge must also add on-ramps and traffic lanes to accommodate annual growth rates of 40 percent that are expected to continue for the foreseeable future.

Carriers manage to add capacity and new services without disrupting existing services by building overlay networks and by separating access (edge) and transport (core) functions. Consequently, most equipment vendors have tried to produce one of two types of devices:

- **Access-oriented—**Like the multiservice provisioning platforms (MSPPs) offered by ATM/IP switch vendors like Cisco, Nortel and Lucent along with others which are not part of the ATM/IP crowd like Ciena, Appian, Geyser and Mayan.
- **Transport-oriented—**Like the optical aggregation equipment from Anda, Chromatis, Coriolis, Celox and Sirocco.

Most of the MSPPs can handle all the popular data protocols and interfaces, but they aren't designed as optical aggregators. In contrast, the new aggregators don't offer all the data interfaces, but they do support a full range of hierarchical aggregation from DS-3 to OC-48, and offer electrical-to-optical conversion.

Successful edge devices will handle multiprotocol data services *and* multispeed aggregation. Both ATM-over-optical and IP-over-optical uplinks to the core also will be required, since many core networks will continue to use both. Scalability on both the "line side" and the "trunk side" will also be paramount—just as it is in our Golden Gate Bridge analogy. As more end users get broadband access (see Figure 2, p. 64), more ports will be required and more traffic will be aggregated on to the core.

Clearly, scalability and ease of integration with existing equipment will be the keys to successful edge equipment designs, since none of today's network components can be removed without disrupting service. Replacing them would be like trying to replace the suspension cables on the Golden Gate Bridge. Broadband access switches that incorporate multiprotocol and multispeed

Overlay networks have caused management systems to proliferate in carrier networks

FIGURE 2 U.S. Broadband Subscribers (Millions)

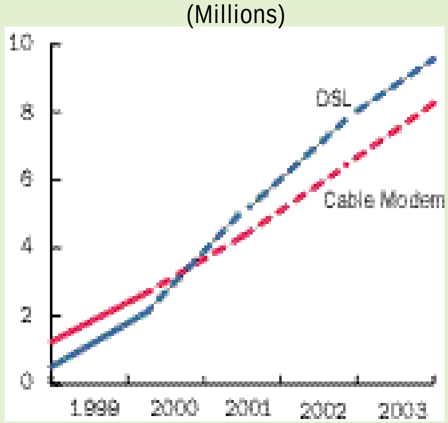

Sources; TeleChoice and Strategis Group, 2000

aggregation will help, especially if they are based on a consistent architecture that scales up and down in capacity.

Among promising candidates in this area are products from Cyras, Gotham, Pluris and Tenor. Such devices will help carriers build new capacity into the edge network without interfering with the figurative "suspension cables," and help address the carriers' second major requirement—the need for operational simplicity, or serviceability.

Add In Serviceability

Overlay networks allow service providers to keep existing services running and to experiment with new ones, but when the new services take hold, the carriers face an unpleasant side effect: The new services rarely use the same provisioning, management and troubleshooting systems as the old network. The resulting proliferation of network and service management systems adds complexity and escalating costs to network operations. These operations and management costs can amount to as much as 50 percent of the carrier's total cost to provide a service.

Successful edge devices must reduce providers' total cost of ownership by being "operator friendly." For example, many of today's MSPP edge devices use architecture with an "IO" card connected to a "protocol processor" card over a switching midplane. While this architecture is optimal for mapping speeds and feeds to the service protocol, the approach increases operational costs. Operators must manually track and anticipate card-level configurations and maintain a detailed per-card inventory.

Furthermore, the equipment's ability to scale up is limited by the midplane's speed, and the ability to scale down is limited by the complex midplane architecture. The result is that the service provider ends up adding more "pizza boxes" to the overall network.

Neither the MSPPs nor the new edge aggregators have paid as much attention to serviceability as they will need to if they want to earn carrier business. Most of the vendors claim they can integrate with carrier management platforms that are built upon the telecommunications management network (TMN) model, with its element, network and service management layers. But the truth remains that for every new vendor and for many new edge devices, the provider is burdened with yet another management system. And most systems require customization to provide mission-critical data feeds to billing and accounting, inventory support, subscriber management and other systems.

Good News Is On the Way

There are four good reasons to expect that equipment vendors will soon be able to provide broadband access switches that more fully address the service providers' bandwidth bridge and serviceability problems. These are:

1. Advanced network processors: Unlike hardwired ASICs, programmable network processors let carriers purchase "generic" cards and enable the desired service function in software. Some processors are even capable of automatically detecting the service required.

With auto-detection at the service level, generic cards can, for example, take in either native ATM cells, IP frames or TDM connections, and convert the traffic to optically multiplexed flows while maintaining the service level requested by each user flow. Ocular Networks recently announced this capability.

2. Advanced computing memory: Cheap, reliable and fast memory allows equipment vendors to load large amounts of data and software onto individual cards. This, in turn, provides the data collection and computations that are required to support increasingly demanding and detailed service level agreements (SLAs).

Managed services require that large amounts of user-specific data be stored to maintain the service level agreement. Keeping this information at the edge makes it easier for the carrier to manage each customer's individual service level, because locally-stored, call admission control data can more easily be integrated with edge-based policy servers.

Abatis Systems (recently acquired by Redback Networks) and Ellacoya Networks are among those vendors who are incorporating optimized edge intelligence into their platforms, so that service providers who use their equipment will be able to offer users "consumable IP," enabling users to buy a "data flow" service with more attractive pricing options.

3. High-capacity switching: Switching fabrics, such as those used by Brightlink, can now move traffic at optical speeds without any blocking or buffering. This allows full protocol processing without affecting network performance. Because huge capacity exists today both in the core and on

FIGURE 3 Next Generation Edge Environment

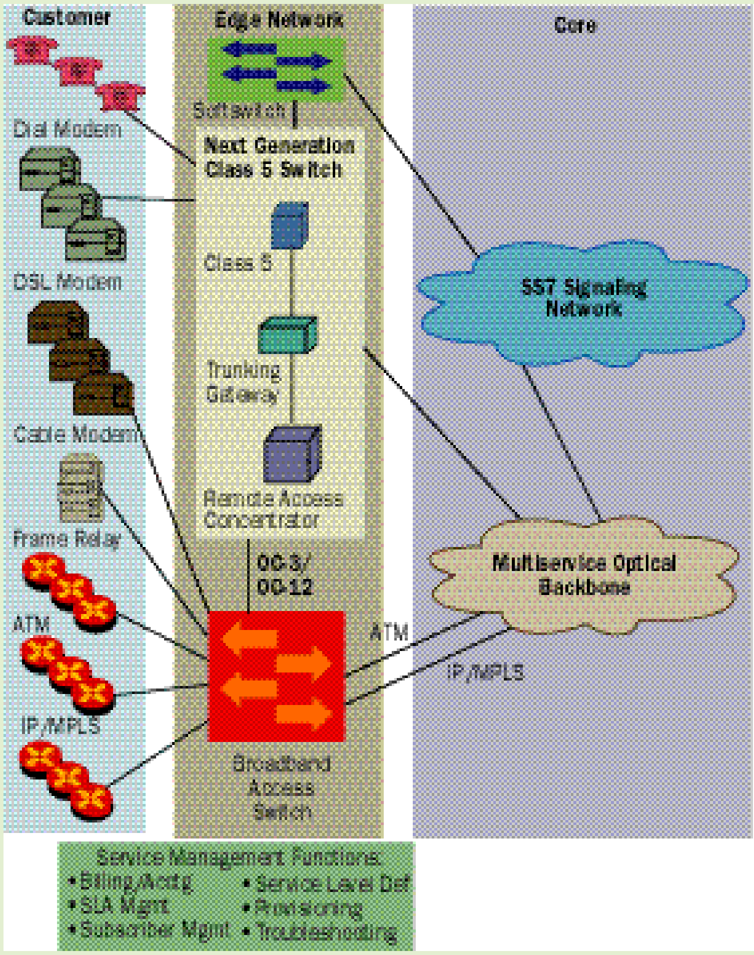

the customers' premises, using these switch fabrics in the edge will make it easier to allocate and manage capacity.

4. Standardized software support: While edge equipment vendors have not devoted enough attention to legacy-integrated management systems, software-focused vendors like Vertel and CrossKeys have been developing standards-based solutions that feature both TMN and common object request broker architecture (CORBA) software. Telcordia (formerly Bellcore) has also developed the Operations Systems Modifications for the Integration of Network Elements (OSMINE) solutions, which provide a standard means to integrate new software and hardware systems into the existing Regional Bell Operating Company (RBOC) operations support systems (OSS).

Edge equipment vendors who adopt such off-the-shelf packages will save development time and offer carriers several real benefits: A consistent operations interface, faster time-to-market for new services and reductions in operations integration and training time.

To date, only a few edge equipment vendors, including Gotham Networks, have taken full advantage of all four of these technology advances. Furthermore, although these developments can help edge device vendors more easily build broadband access switches that bridge the bandwidth gap between LANs and WANs and improve serviceability for the provider, one critical problem remains unsolved.

That problem starts with the effects of voice and data network convergence, and with the necessity to view user needs not just as technology requirements, but as service requirements.

Convergence And The Edge

Convergence may be the most misused term in today's communications jargon. It's not just about putting voice over data networks—although that is happening in some customer premises devices (personal computers, routers and PBXs), and in some network access services. Convergence is also happening in some carrier networks, where multimedia gateways segregate and remix a variety of traffic types.

New services are also beginning to feature converged functionality, such as "click-to-help" buttons on websites that combine user profiles with additional information and initiate a callback to the user from the right customer service rep at the right call center.

Not all converged services require converged user interfaces. For example, Grandma could use her POTS phone to leave a voice mail message at her local pharmacy requesting a prescription refill. That message could trigger a process in the local pharmacy's converged system or in a remote call center, retrieving Grandma's pertinent information (prescriptions, refills, insurance provider and doctor's contact information) and initiating either a live or recorded ("your prescription is ready") callback to Grandma.

As converged devices, networks and services grow, the edge will become even more ill-defined, and edge equipment devices will face additional challenges. Consider a personal computer or IP-enabled PBX that creates a voice-over-IP packet and sends it over a DSL or cable modem line to a broadband access node. Then the network operator terminates the IP call, and reinitiates it (using a softswitch or media gateway) as a POTS call so it can be completed on the PSTN.

An edge device that could perform these functions would include a media gateway or use a core-based softswitch (Figure 3). It might include some or all of the 2,000+ features enabled by a Class 5 voice switch, or rely on an adjunct device to enable features.

Oresis Communications and Telica are among the vendors who provide Class 5 replacement switches that feature small footprints, high port densities and many voice and data features. New entrants like Zhone Technologies, and incumbents like Cisco, Nortel, and Lucent, also offer a consolidated architectural view of the ideal convergence-friendly edge.

But the reality is that all these vendors must still use multiple edge devices in order to deliver on their simplified architectural ideal. So far, the promise of convergence-based services still comes only at the cost of added complexity for network operators.

Successful edge networks—if not individual edge devices—must support each converged service, recognizing and handling all types of voice and data traffic. Voice traffic from PCs or IP-enabled PBXs must be recognized, classified to enable voice-grade quality of service, and differentiated from other voice flows (such as voice mail stored on the email server). Moreover, the edge must account specifically for calls based on their origination over DSL, frame relay, dial-up modem or wireless access. Hence, the edge must provide media-agnostic service interworking between multiple access technologies.

In addition, the edge should also provide an H.323 or SIP signaling gateway function between the enterprise network and the POTS network. By doing so, the edge will provide convergence between the IP signaling network and the POTS SS7 signaling network.

Finally, new converged services—like Grandma's prescription refill—will require additional edge network functions. Wherever Grandma's profile and history are stored, the edge node nearest to her must link her phone number with that information.

Conclusion

Over the coming months and years, the optimal edge will evolve to eliminate the current performance compromise created by today's "dirt paths" to the information superhighway. Successful vendors will find ways to collapse edge aggregation, service creation, switching, routing and optical access into flexible broadband access and next- generation Class 5 switches for the new edge network.

These new edge systems will also reduce the time and manual intervention required to perform simple network tasks, addressing the people shortage affecting all providers. And the edge of tomorrow will eliminate today's complexities of overlay traffic aggregation and service-by-service management platforms. New network designs will eliminate all the service-specific and hierarchical aggregation layers in today's edge network.

Although each service provider's edge network will be unique, Figure 3 represents a comprehensive model of tomorrow's optimal edge network. The optimal edge network will support any current protocol on any port, and allow that protocol to run in its native form, end-to-end throughout the network. Aggregation, switching, routing and transport will automatically be managed, since all components will be part of a scalable, integrated solution. These solutions will collapse today's many layers and multiple devices to bring about true network convergence—aggregation on the premises and at the edge of the access network, as well as converged services and streamlined aggregation at the edge of the core transport network□

Companies Mentioned In This Article

Abatis Systems (now Redback Networks) (www.abatis-sys.com)
Accelerated Networks (www.acceleratednetworks.com)
Amber Networks (www.ambernetworks.com)
Anda Networks (www.andanetworks.com)
Appian Communications (www.appiancom.com)
Atrica, Inc. (www.atrica.com)
Brightlink Networks (www.brightlink.com)
Celox Networks (www.celoxnetworks.com)
Cerent (now Cisco, www.cisco.com)
Chromatis (now Lucent, www.lucent.com)
Ciena Corp. (www.ciena.com)
Cisco Systems (www.cisco.com)
Convergent Networks (www.converntenetworks.com)
Coriolis Networks (www.coriolisnet.com)
CrossKeys Systems (www.crosskeys.com)
Cyras Systems (www.cyras.com) (Ciena to acquire Cyras)
Ellacoya Networks (www.ellacoya.com)
Geyser Networks (www.geysernetworks.com)
Gotham Networks (www.gothamnetworks.com)
Integral Access (www.integralaccess.com)
Lucent Technologies (www.lucent.com)
Luminous Networks (www.luminousnetworks.com)
Mayan Networks (www.mayannetworks.com)
Nortel Networks (www.nt.com)
Ocular Networks (www.ocularnetworks.com)
Oresis Communications (www.oresis.com)
Pluris Inc. (www.pluris.com)
Qeyton Systems (now Cisco, www.cisco.com)
Quantum Bridge Communications (www.quantumbridge.com)
Redback Networks (www.redback.com)
Sirocco (now Sycamore Networks, www.sycamorenetworks.com)
Sentient (now Cisco, www.cisco.com)
Shasta (now Nortel, www.nt.com)
Tachion Networks (www.tachion.com)
Telica (www.telica.com)
Tenor Networks (www.tenornetworks.com)
Vertel Corporation (www.vertel.com)
Zhone Technologies (www.zhone.com)

The edge must provide media-agnostic service interworking between multiple access technologies