Dream Networks Face Nightmare Forecasts

Sandy Borthick

Routing and switching products spawned in the exuberant late 1990s are beginning to ship—but their timing couldn't be worse.

It's a shame the market can't muster more enthusiasm for the current crop of routing and switching products. How ironic that such a strong and diverse set of options faces the most dismal purchasing forecast in years.

These routers, switches, VPN and optical Ethernet platforms were, for the most part, conceived at or just past the peak of the Internet boom, and they reflect a range of expectations about the future of carrier and service provider networks. Overall, however, the once-dominant view that the growth of the Internet would dictate every platform, protocol and product choice, is giving way to a more realistic acknowledgement: That some carriers and service providers, especially the ones who still have money to spend, aren't very interested in an all-IP world, unless it also will preserve their legacy revenue streams, processes and technology investments.

This explains the renewed vendor interest in big multiservice switches, which can be used in lieu of core routers to more gradually migrate carriers toward IP/MPLS (or not), and in repurposing MPLS tunnels to haul ATM and frame relay traffic, complete with their signaling and control protocols, across IP backbones. With these MPLS tunnels, ILECs and IXC/ISPs can take advantage of cheap IP bandwidth from wholesalers (e.g., Level 3, Global Crossing), without having to use MPLS to control their networks (Figure 1).

Some startups (including Equipe Communications, Oresis Communications, Pelago Networks and WaveSmith Networks) include in their products and even champion the once-maligned ATM. At least one vendor (Equipe) plans to attack the MPLS backbone core market by first introducing a huge ATM switch. Other startups, including Gotham Networks, Tenor Networks and Vivace Networks appear not to be including ATM, but focusing only on IP/MPLS.

Although Lucent cancelled its big core multiservice switch, the MSC 25000, in September, other traditional ATM carrier suppliers like Alcatel (Newbridge), Marconi (Fore) and Nortel continue to enhance their core multiservice platforms. Even Cisco is rumored to have a replacement for the MGX 8850 multiservice switch, dubbed the Jupiter.

Which is not to say that all-IP notes aren't still being sounded. While the ATM folk may be co-opting MPLS for their own purposes, the IP-over-everything crowd hasn't given up (Figure 2). For example, Avici, Cisco and Juniper have all recently come out with smaller versions of their backbone core routers in an attempt to "expand their addressable market," marketing-speak for selling more backbone core equipment into the metro/regional POPs and aggregation points.

Sandy Borthick is BCR*'s technical editor.*

FIGURE 1 MPLS And The Multiservice WAN

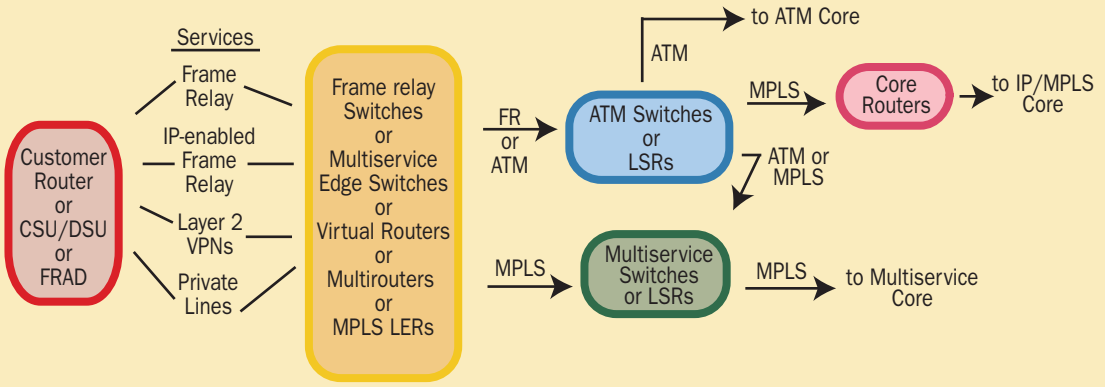

FIGURE 2 Converged IP Routing

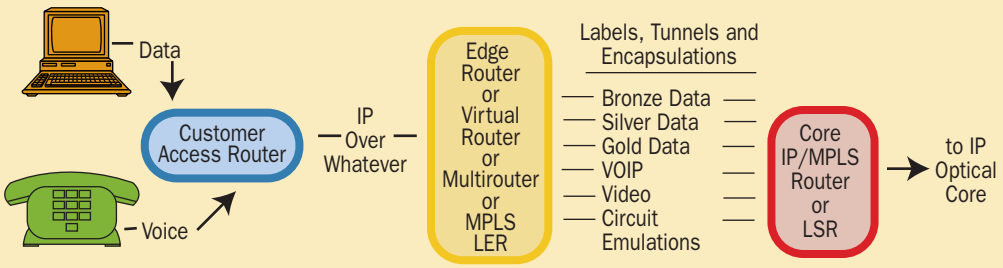

MPLS is being repurposed to haul legacy traffic across IP backbones

Other routing and switching choices abound in the access and metro markets, although the re-monopolization of the ILECs and their slow-rolling on broadband access bode ill for any near-term dramatic adoptions. Despite this prognosis, and the slow market acceptance for most first-generation network-based VPN platforms, new versions of IP-VPNs continue to pop up and figure prominently in some edge switch/router vendor visions. Other vendors continue exploring ways to streamline the Layer 2 functions between IP and the underlying SONET, DWDM—or even dark fiber at Layer 1.

Meanwhile, the mighty Ethernet, having escaped its enterprise bounds, threatens to loose its commoditizing magic on the local loop and beyond. Access options of n × 100 Mbps, Gigabit and 10-Gigabit Ethernet, plus several new flavors of optical ring and mesh architectures, have emerged to complete the access/metro technology menu (Figures 3 and 4).

In short, a broad range of routing and switching options are competing in a very tight market. It would be a buyer's market, if the buyers weren't just as financially strapped as the equipment vendors, and if anyone could see how to break the broadband access bottleneck to profitably market the capacities and services these new technologies could unleash.

Instead, the carriers have announced 5- to 20-percent capital spending cuts for next year, and an even slower pace of DSL deployment. In reality, the capex cuts began this year. The Dell'Oro Group found WAN equipment sales down 19 percent in 2Q01, following an 18 percent drop in the first quarter. On the enterprise front, the news is also bad: Even Ethernet switch sales have fallen, in what Dell'Oro predicts will be the technology's first down year ever. Worse, Dell'Oro doesn't expect the next enterprise upgrade cycle—bringing Gigabit Ethernet to the desktop and 10-Gigabit Ethernet into the corporate backbone—to hit until around 2003.

Perhaps the best that can be said of this market is that it is a level playing field of sorts—practically all the players are poor. That's not all bad, however, since at least it represents a common departure point. But how soon can we get out of here, and where do we go?

Can We Go Now?

Of course, no one knows when the network equipment market will revive, which players will survive the downturn and which technologies will be favored in the next upcycle. But everyone agrees the next upcycle is inevitable.

"Gloom is not a direction," said Atiq Raza, founder, chairman and CEO of the early-stage investment and partnering company Raza Foundries, speaking at the annual IEEE/Stanford

FIGURE 3 Ethernet Everywhere

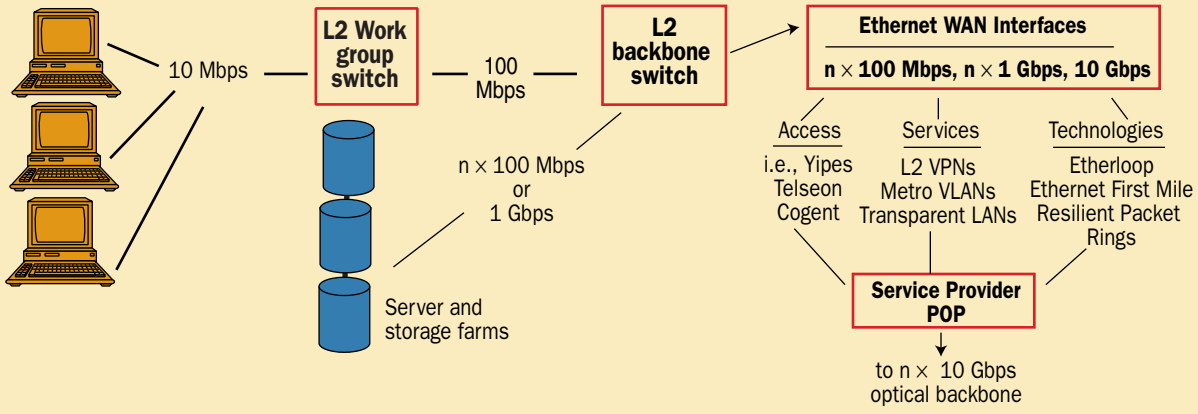

Some carriers are attending Ethernet in the First Mile (EFM) IEEE meetings

University Hot Interconnects conference in August. While the carrier spending cuts and oversupply of start-ups may look bleak today, Raza characterized the current situation as a case of indigestion in what he calls the "terabit food chain." Some of the new companies just won't survive, but for the rest, debt restructuring and inventory reductions will clear out the blockages so that new products, services and revenues can again flow freely.

But how quickly will demand for these products and services pick up? Raza cited a survey performed earlier this year by Internet pioneer Dr. Lawrence Roberts, which found that Internet backbone traffic had quadrupled in the year ending 1Q01. According to Roberts, carriers have gotten just about all the mileage they can out of their existing infrastructures and will soon have to open their wallets to increase traffic carrying capacity just to accommodate this business-as-usual growth. (Grains of salt may be in order, however, since Roberts has a horse in the core switch/router race, the still-secretive Caspian Networks. For more on Roberts's projections, see Peter Sevcik's column in *BCR*, November 2001, pp. 10–12.)

Unfortunately, the main driver for carrier network expansion—high-speed access and services deployed by carrier customers at the edge—is still hung up in the chicken-and-egg conundrum: Carriers will only expand their networks if and when customers begin to adopt higher-bandwidth carrier services—e.g., 100-Mbps and Gigabit Ethernet access, hosted data storage and applications, videoconferencing and content delivery networks. But if the carriers don't expand their networks so that more customers can get high-speed access, they can't deliver these new services.

On a brighter note, assuming Ethernet everywhere is the favored model going forward, and high-speed access can be made more widely available, it wouldn't take many customers to drive carrier expansion. Even a few active Gigabit Ethernet customers can quickly saturate a 10-Gbps network—or, as Dr. Lynn DeNoia, a *BCR* contributor, consultant and professor likes to say, "There are only 10 tens in 100."

Another way to quantify the bandwidth requirements of an Ethernet-everywhere model, according to Cisco's Andreas Bechtolsheim, is that 500 terabits per second (Tbps) of backbone capacity are needed for either 5 million users with 100-Mbps access or 100 million with 5 Mbps. And Telechoice president Christine Heckart, speaking on an October teleconference sponsored by Net.com, maintains that if only 5 to 10 percent of the addressable market adopts Gigabit Ethernet access, all the current long-haul fiber capacity will be exhausted.

By way of comparison, at the start of this year, analyst and *BCR* columnist Peter Sevcik estimated that total trunk capacity of the Internet at the end of 2000 was 6,000 Gbps or 6 Tbps.

How likely are such demand scenarios? Ethernet clearly offers the fastest and cheapest path to high-bandwidth access services (see *BCR,* September, 2001, pp. 44–48). But for all its enterprise success, it may not have as clear a shot at dominating carrier access networks.

Access To Everything

ILECs undoubtedly have been doing the math on carrier-class Ethernet switching products from established vendors like Extreme Networks and Foundry Networks, while looking over their shoulders at high-speed Ethernet service providers like Cogent, Telseon and Yipes. Some carriers (including Qwest, WorldCom and Verizon) have been participating in the IEEE's 802.3ah Ethernet in the First Mile (EFM) task force, recently launched to extend Ethernet into the local loop for both business and residence access services.

FIGURE 4 End-To-End Fiber Fantasy

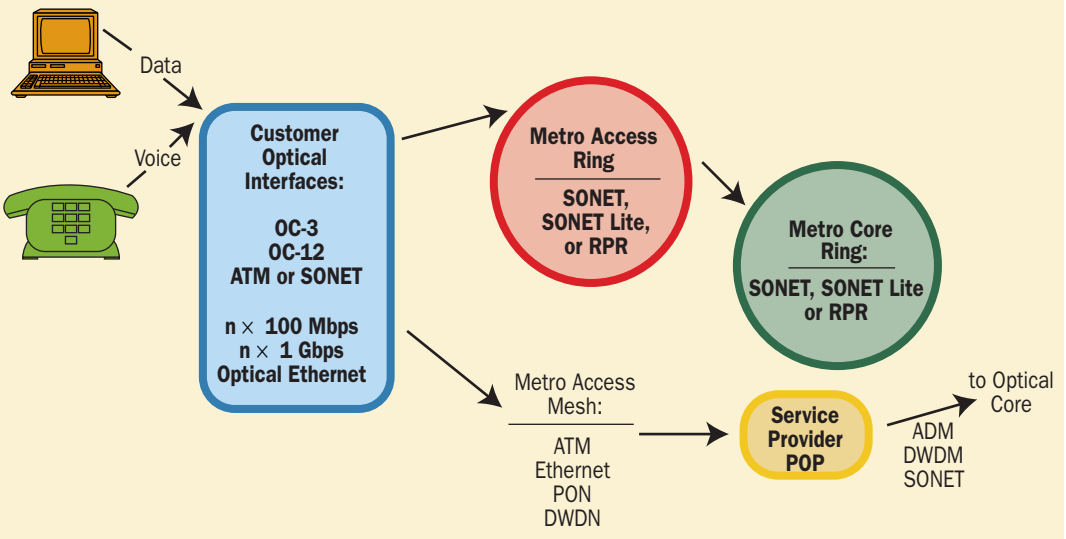

This work includes standardizing a longer reach for singlemode fiber, figuring out how to get Ethernet running over copper POTS wires and developing an Ethernet optical ring (see *BCR*, September, 2001, pp. 40–43). More than 50 equipment vendors have formed the Metro Ethernet Forum to promote their products into this market, including Appian Communications, Atrica, Inc. and many more.

At the same time, armies of developers are working to supply the equipment vendors with network processors, custom chipsets and optical transceivers that can run at 10 Gbps (see *BCR*, October 2001, pp. 45–49). Right now, only a few backbone core products (e.g., Cisco 12416 and Juniper M160) have true 10-Gbps/OC-192 throughput, although practically every vendor of core and edge switch/router products is promising it in forthcoming products.

The problem has not been 10-Gbps interfaces and simple switching, which are adequate for steering big pipes through the backbone core of the network. Instead, the challenge has been getting every packet classified, marked, queued and billed fast enough to support 10-Gbps (and higher) throughputs at the network edge. It's an exciting technology race, with many companies working very hard on different hardware and software combinations to address the issue, but the question remains: Are these people working on products that carriers will use to build services that they can sell to customers—at a profit?

Figuring out which features and functions in edge switch/routers will help "pull thorough" customer demand for new services has been a losing game for several years now. Customers just haven't taken to "value-added" services like integrated voice/data access, unified messaging, follow-me forwarding, hosted applications and streaming content. At the same time, everyone knows the ILECs and IXC/ISPs aren't about to cannibalize their installed base of circuit-based voice, frame relay, ATM and private line revenues by promoting an all-IP or all-Ethernet world.

One of the most visible busts so far has been the concept of IP-VPNs as a carrier service. Remote-access VPNs have largely supplanted expensive dial-in services, but enterprise customers are still leery of using IP-VPNs instead of site-to-site frame relay and private line networks. IPSec VPNs have proven complicated and MPLS VPNs are still unproven. It hasn't helped the VPN cause that arguments swirl around the scalability of different VPN, MPLS and BGP signaling and tunneling mechanisms.

MPLS—Switch, Route Or Shim?

The BGP-signaled, MPLS-based variety of IP-VPNs, described in RFC 2547, came under criticism this past summer by several prominent IETF insiders. Others chimed in to attack or defend the concept, but such theoretical arguments are unwinnable. It's impossible to prove that something which hasn't happened yet either will or won't happen in the future. The only way to know for sure is to build the network, operate it and find out.

That's what Dave Garbin, Cable & Wireless chief network engineer, plans to do, and he doesn't anticipate any problem. "It's a tempest in a teapot," he said. "We expect to use RFC 2547 on the edge, with a routing table per customer, and virtual routing in the core."

Garbin expects basic services like Internet access and IP-VPNs to sell better initially than hosted content, voice and video services. He is in the process of migrating the core of the Cable & Wireless data network away from a leased ATM infrastructure onto a routed IP core, which will be traffic-engineered using MPLS. Although the new core will be all-IP, he is not integrating the Layer 3 routing and Layer 2 link control functions. "We believe in separating the traffic engineering from the core routing functions. Otherwise problems with BGP could affect my ability to carry traffic."

This separation of routing from control is not what MPLS proponents had planned for the protocol when they started working on it in 1998. They wanted to replace ATM and create an integrated IP/MPLS core. They began by rationalizing the approaches of several competing vendor proposals for a cut-through form of IP switching. Then it became apparent that MPLS also could be used to create VPN tunnels, and to control optical circuits and lightwaves. Most recently, MPLS is also being seen as a backbone Layer 2 mechanism for carrying everything from Ethernet, PPP and HDLC packets to legacy packetized services like frame relay and ATM.

After more than two years of being inundated by more than 100 MPLS, VPN, IP over optical, and other Layer-2-oriented Internet Drafts, the IETF finally decided to create a new "Sub-IP" area devoted to sorting through these issues and technologies.

In particular, the Provider Provisioned Virtual Private Networks (ppvpn) and the Private Wire Emulation End-to-End (pwe) working groups were launched to hash out ways in which ATM and frame relay signaling and control protocols might be interworked with, or encapsulated in IP/MPLS. The relevant pending Internet Drafts include Martini, Kompella, Koleyni, Brayley and Fischer.

Although all the drafts have a common purpose—to encapsulate and transport legacy protocols in IP/MPLS tunnels—they vary in the specific ways they handle the legacy signaling and control protocols. Draft Martini is arguably the most popular, and most vendors have announced plans to include Draft Martini functions in their products. The working groups still have to iron out differences among the drafts, but Cable & Wireless' Garbin isn't worried about that. "They'll come

C&W plans to "martini-ize" legacy frame and ATM traffic over a new, IP-routed core

Will the ILECs just starve out the Ethernet access players—like they did with DSL?

together," he said. "Kompella has the signaling—it's really a mapping of frame relay DLCIs to [MPLS Label-Switched Paths (LSPs)]—and it will carry Martini encapsualtion. We like Draft Martini so much we made a verb of it—we plan to 'Martini-ize our traffic."

Whose Problem Is This?

Arguably, the backlog of "Sub-IP" Internet Drafts, the advent of IP switching and MPLS all demonstrate just how far the IETF has moved, or been pushed, away from its roots. Once tied closely to the network designers and operators it served, its open, collegial format has been overrun by vendors seeking its Internet Draft and RFC imprimaturs as high-tech marketing tools.

The IETF leadership has been largely powerless to stop this, and, in fairness it is hardly the IETF's fault that so many vendors are offering so many ways to do so many things with, to and for IP! Venture capitalists kept funding all sorts of new Layer 2/3 components and boxes, while the IETF keeps, in a sense, sanctioning the technology, and the carriers keep locking it up in their labs for a few years to "evaluate" it.

This has produced a lot of creative switch/routing technology, but it hasn't propelled the market for services based on these technologies. If anything it has slowed the carriers from offering new services, while they try to pick the best platforms, products and protocols and figure out how to introduce them as services without killing their legacy revenue streams.

Conclusion

In the current market, vendors are soft-pedaling technical differentiation and instead emphasizing their products' cost-effective scalability, reliability and the opportunity to gain share in "market transitions." This makes for a sensible pitch in these uncertain times, but carriers would have to be crazy to trumpet their adoption of any cost-cutting technology without announcing some high-margin services to layer on top of it. Otherwise, customers would expect the economies to be passed through, and the carriers would end up cannibalizing their existing services.

Unless someone figures out how to make IP-based network services attractive, and how to break the local access bottleneck, the future belongs to the past. The consolidated IXC/ISPs and the re-monopolized ILECs have already withdrawn or slowed all their integrated and broadband offerings, and will likely try to just starve out their latest Ethernet-based competitors—like Cogent, Telseon and Yipes—much as they did with DSL providers Covad, NorthPoint and Rhythms.

The only good news in this—and it's pretty weak—is that the product "playing field" is more level today than it has been in years. No longer do incumbent vendors like Cisco, Lucent and Nortel enjoy unquestioned, cozy relationships as suppliers of record to their carrier customers. Most of them spent the last few years, and a lot of money, buying up new technology start-ups and trying to incorporate often-immature gear into their product lines. They are as highly leveraged, vulnerable and hungry as some of the startups.

There's really nothing the equipment and component suppliers can do about this gloomy prognosis, except perhaps to arm and encourage those service providers that might yet take on the titans, break the access bottleneck, and roll out some real competition□

Companies Mentioned In This Article

- Alcatel (www.alcatel.com)
- Appian Communications (www.appiancomm.com)
- Atrica, Inc. (www.atrica.com)
- Avici (www.avici.com)
- Cable & Wireless (www.cw.com)
- Caspian Networks (www.caspiannetworks.com)
- Cisco (www.cisco.com)
- Cogent Communications (www.cogentco.com)
- Equipe Communications (www.equipecom.com)
- Extreme Networks (www.extremenetworks.com)
- Foundry Networks (www.foundrynetworks.com)
- Global Crossing (www.globalcrossing.com)
- Gotham Networks (www.gothamnetworks.com)
- Juniper (www.junipernetworks.com)
- Level 3 Communications (www.level3.com)
- Lucent (www.lucent.com)
- Marconi (www.marconi.com)
- Nortel Networks (www.nortelnetworks.com)
- Oresis Communications (www.oresis.com)
- Pelago Networks (www.pelagonet.com)
- Qwest (www.qwest.com)
- Raza Foundries (www.razafoundries.com)
- Telseon (www.telseon.com)
- Tenor Networks (www.tenornetworks.com)
- Verizon (www.verizon.com)
- Vivace Networks (www.vivacenetworks.com)
- WaveSmith Networks (www.wavesmithnetworks.com)
- WorldCom (www.wcom.com)
- Yipes (www.yipes.com)