Frame Relay: Good For Another 10 Years?

Robert Bellman

Are higher speeds and inverse multiplexing what frame relay needs to beat its new access competitors and befriend the new backbone transport technologies?

Keep it simple, stupid! It's familiar advice that's too often ignored. But the folks who guided frame relay through its formative years took the slogan to heart and were rewarded with lasting success. According to Vertical Systems Group, frame-relay services—first launched in 1991—reached an estimated 1.78 million ports worldwide in 2001 and total revenue of $12.7 billion. In the U.S. alone, frame-relay services earned around $7.6 billion last year—not too shabby considering that U.S. private-line revenue was around $12 billion.

Now the frame relay industry is invoking another cliché: If you can't beat 'em, join 'em. It worked when frame relay was threatened by ATM and promises another life-extending kick as frame faces competition from IP and MPLS. Together, simplicity and cooperation comprise a formula that should earn frame relay a second decade of steady growth.

Smoothing The Wheel

Frame relay began in the late 1980s as the technical successor to X.25. Reliable digital transmission and intelligent end-stations had made X.25's belt-and-suspenders approach to error handling (at both the link and network layers) obsolete. Seeking a more efficient vehicle for packet-data traffic, the designers of ISDN combined the virtual circuits of X.25 with the lightweight error detection of a link-layer protocol. The industry soon recognized the value of streamlined packet switching, and frame relay was born as a standalone, network service.

Like most standards, however, the frame relay specifications left too much open to interpretation. Different equipment vendors could build products that adhered to the "standard" but would not interoperate. So in 1991, a handful of companies founded the Frame Relay Forum and set about developing implementation agreements (IAs)—simple, unambiguous versions of the specs. Even as the Forum swelled to more than 300 member companies, it wisely avoided the we-can-solve-everything excesses of similar groups. "We didn't

FIGURE 1 Traditional Frame Relay Services Over An ATM Backbone

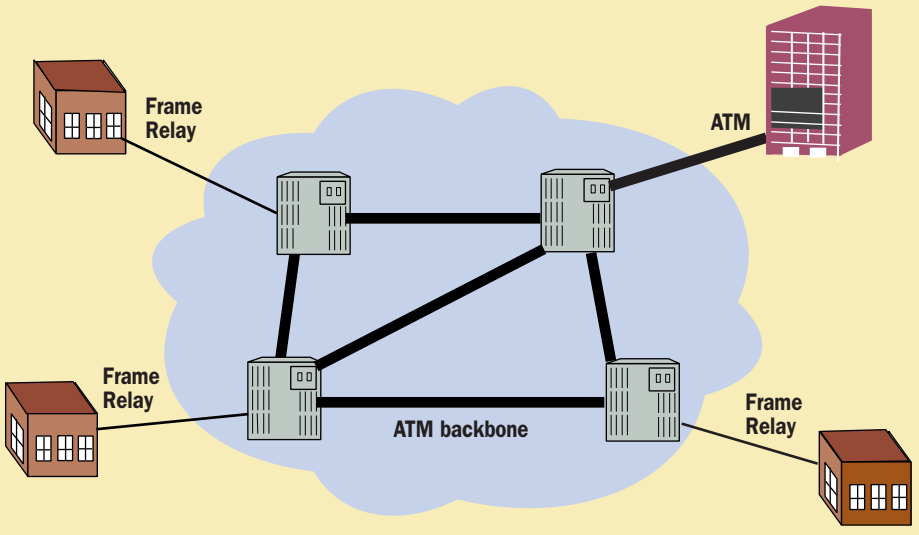

Bob Bellman is president of Brook Trail research, a consultancy specializing in network technology assessment and marketing. He can be reached at bob@brooktrail.com

Demand for frame continues to rise, while supporters work to exend its utility

want to reinvent the wheel, just make it smoother," recalled Mike Walsh, president emeritus of the Frame Relay Forum and director of product marketing at Stallion Technologies.

Based on the Forum's IAs, frame relay won acceptance throughout the industry because it was simple, familiar, reliable and affordable. Simplicity made frame relay easy to implement, and equipment manufacturers quickly rolled out frame relay access devices (FRADs) and switches. Simplicity also made it easy for service providers to support frame relay at a reasonable cost.

Customers readily understood frame relay; its virtual circuits were just like private lines, only they allowed traffic bursts and consumed fewer CPE ports. And once frame's reliability and attractive pricing were demonstrated, customers were eager to buy. Who wouldn't want to save 30 to 50 percent on their monthly bandwidth bill if the risk was minimal?

Another, more serendipitous factor also contributed to frame relay's success. At about the same time that frame relay services were emerging in the early to mid-1990s, network managers were looking to link their newly-built LANs across the wide area. Private lines were expensive, but frame relay was a perfect fit. So, new traffic from LAN interconnection provided burgeoning demand—supplemented by legacy traffic from SDLC, bi-sync and other heirloom protocols—and frame relay services hit the ground running, earning $1.7 billion in 1991 and growing steadily ever since.

Frame relay easily beat out its immediate high-speed data service competition—SMDS—and went on to form a symbiotic relationship with the other threat, ATM. Although ATM is more versatile than frame, runs at faster speeds and guarantees quality of service (QOS) for diverse traffic types, it is also more complicated and requires more expensive CPE. So while carriers have come to prefer ATM in their transport networks, most customers favor frame as their service interface.

Rather than campaign against ATM, even when some ATM proponents were calling frame an "interim technology," the Frame Relay Forum worked with the ATM Forum to make frame relay and ATM interoperable. Now, all the major frame relay services in the world run over ATM backbones, and low-speed frame relay traffic gathered from a company's remote offices can be delivered to headquarters over a high-speed ATM interface (Figure 1). Frame relay interfaces are considered standard equipment on virtually every vendor's CPE access device, and on every network edge, multiservice and backbone switch. Indeed, frame relay's ability to evolve—from an interface and transport technology to a general-purpose service interface—has been key to its survival, and promises to keep the frame relay market healthy for years to come.

You Can't Win Them All

Of course, not every frame relay initiative has been a complete success. Switched virtual circuits (SVCs), for example, have never really taken off. And voice-over-frame relay (VOFR) has seen only limited use.

Permanent virtual circuits (PVCs) were standardized before SVCs and gained immediate acceptance. They could be centrally configured by the service provider, and offered a straightforward alternative to private lines. SVCs, on the other hand, are more like dial-up connections. An attached device signals the frame relay network for a new virtual circuit and releases it when the session is finished.

SVCs were promoted as a basis for dynamic bandwidth allocation. Routers, for example, could automatically request SVCs to handle traffic bursts or to back up failed leased lines. SVCs could also be useful for VOFR, just as dialup is used for ordinary telephony.

But the complexity and potential instability of dynamic bandwidth allocation made SVCs too scary for service providers. There also were technical hurdles in getting switches to handle SVC signaling in volume and gather enough billing statistics to be useful, and difficulty in figuring out how to price SVCs was no doubt part of the problem.

The lack of enthusiasm for VOFR is more surprising, especially given today's excitement about voice over IP (VOIP). As a connection-oriented protocol, frame relay makes a better foundation for voice calls than IP. But, as Steve Taylor, president of Distributed Networking Associates, Inc. and publisher/editor of Webtorials.com, lamented, "The best technology doesn't always win."

Voice-over-frame was marketed to the telecom community, who wanted five-nines reliability. VOFR can't guarantee five-nines, and neither can VOIP, but VOIP is being sold to the IP community instead of the telecom and carrier folks. "IP people don't have the same expectations as carriers," said Taylor. "Plus VOIP has the IP halo." As a result, VOIP is gathering momentum while VOFR has stalled.

Faster, Denser, Cheaper

Still, these disappointments have not deterred frame relay's growth. Demand for frame relay services continues to rise, while service providers, equipment vendors and the Frame Relay Forum work to extend the technology's utility. For starters, there's a gradual move to higher speeds and more ports.

According to Tim Halpin, product director for Frame Relay and ATM Services at AT&T, the most popular speeds for frame relay are DS0 (56/64 kbps) at remote sites and T1 (1.5 Mbps) at hub sites. Vertical Systems Group analysts estimate the current U.S. mix of frame speeds breaks down as follows:

FIGURE 2 Multilink Frame Relay Concept

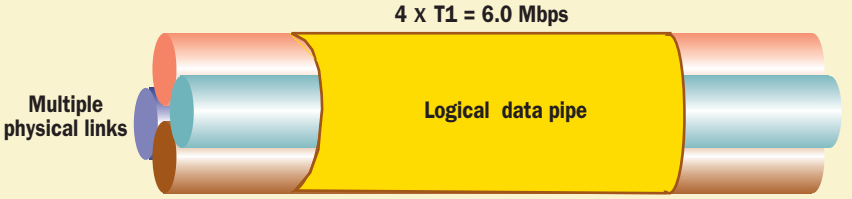

Vendors think the time is right for higher-speed and multi-link frame relay

- DS0—27 percent.
- Fractional T1—26 percent.
- T1— 44 percent.
- T3—3 percent.

Lately, however, Halpin says that customers are stepping up to fractional-T1 at remote sites and T3 or inverse-multiplexed ATM (IMA) at their hub sites. Start-up switch vendor WaveSmith Networks is among the equipment providers who view this as an opportunity. The installed base of frame relay and ATM switches has been around for years, based on products from Cascade (now Lucent Technologies), Stratacom (now Cisco Systems) and Nortel Networks. Increasing traffic is straining the scalability of these relics, WaveSmith maintains, which opens the door to 21st-century products that cost less and run faster.

In February, WaveSmith announced a new Packet Forwarding Module for its Distributed Node Multiservice Switch family, promising frame relay, multilink Frame Relay (MFR) and IP support at one-third the cost-per-port of legacy switches and 2–4 times the DS1 and DS3 port density. Moreover, WaveSmith's new module also supports frame over SONET at speeds up to OC-12 (622 Mbps).

Although the Frame Relay Forum has had specifications in place for SONET-speed frame relay interfaces since 1998, and for the inverse-multiplexed MFR since 1999, switch vendors—and presumably their customers—haven't yet shown much interest. This could be because the specs for SONET-speed ATM, IMA and the frame-to-ATM interworking were already in place and in products, and because ATM already had the carriers' blessing.

Still, the major carrier switch vendors seem to think the time is right for higher-speed, higher-density frame relay and for MFR. For example, Lucent's popular 5-Gbps ATM switch CBX 500 supports MFR and, in April, the company announced new, channelized T1/E1 port cards, providing an eightfold increase in the number of frame relay service ports supported.

Nortel's Passport family of multiservice switches also supports MFR. The Passport 15000, for example, offers a 4-port channelized DS3 card, allowing multiple DS1s within the same DS3 to be bundled as one logical MFR link. Cisco, which dominates the frame relay CPE market, includes MFR in its routers as a standard software feature of IOS. Strangely, however, Cisco does not currently support MFR on its BPX/MGX line of multiservice switches. Maybe that's why Cisco didn't show up for the Frame Relay Forum's MFR interoperability tests last summer.

Back To Basics

The switch vendors' support for MFR represents part of a larger movement, according to Paul Smith, president and CEO of access and edge equipment maker Tasman Networks. "The industry is getting back to basics while waiting for the next big thing," he said. "Instead of investing in new access technologies like PONs [passive optical networks] and wireless, service providers are looking for ways to squeeze more out of the existing copper plant."

AT&T and WorldCom, for example, have added frame over DSL to their service portfolios. Frame relay over DSL is targeted at business customers who have access to DSL and couldn't justify frame before, according to AT&T's Halpin. Compared with traditional frame-over-leased-line implementations, frame over DSL costs less, but it is also less widely available. Unfortunately, it is less reliable, too, due primarily to the longer repair times associated with DSL circuits, and because its top speed depends on the length of the copper loop.

MFR seems to have more going for it, offering higher bandwidth without reliability and scarcity trade-offs. MFR can inverse multiplex from N×DS0 to N×T3, but will probably see its greatest use at N×T1, due to the low cost and widespread availability of T1 circuits. With MFR, a customer who needs 6 Mbps, for instance, pays for just four T1s instead of ordering (and waiting for) expensive T3 service, and then wasting 39 Mbps. And if one of those T1s fails, the multilink connection automatically throttles back to 4.5 Mbps until the broken circuit comes back up (Figure 2).

Besides new software on network switches, MFR also must be supported on premises devices, and that's where equipment companies like Tasman Networks come in. Tasman was formerly known as Tiara Networks, but renamed itself in May and is now expanding into the low-end access and mid-range edge router markets. Cisco currently dominates those markets, so Tasman's goal is to offer both enterprises and service providers a cost-effective alternative. With

Will multi-link frame relay be a midlife kicker, or the last gasp of a mature market?

hardware architecture based on off-the-shelf CPUs, Tasman boasts that it can roll out features more quickly, support higher performance and still keep costs 25 to 35 percent below the competition. Other frame relay CPE providers with MFR capability include Adtran, Larscom and Quick Eagle Networks.

According to Smith, Tasman's routers handle multilink bundles of up to 16 T1s at wire speed, and their MFR support has already been certified for CPE use by both WorldCom and Sprint. "Now," said Smith, "they're just waiting to deploy MFR support on their POP switches." But an overriding concern with reliability makes service providers reluctant to add new equipment to their POPs. "We don't want a separate MFR box in our POPs due to reliability concerns," AT&T's Halpin said. "We want a protected line card on our current switches."

Meg Moschetto, WorldCom's senior manager for data services, also worries about rocking the boat. Today WorldCom offers an MFR-like service by bonding multiple T1s at Layer 1. "MFR should be cheaper and simpler," Moschetto admitted, "but it won't enable new services for us." So while WorldCom is playing with MFR in its labs, it's not in a rush to offer it commercially.

The I/O cards that support MFR on Nortel's Passport family and Lucent's CBX 500 can be backed up with optional redundancy, although that isn't really "protected" in the SONET sense. In any case, none of the major service providers appear to be actively marketing MFR services. Yet it's too soon to predict that MFR will suffer the neglect that befell frame SVCs.

Another potential "next big thing" in access technologies is the emerging Ethernet access specification—numbered 802.3ah by the IEEE, and dubbed Ethernet in the First Mile (EFM) by its vendor proponents in the EFM Alliance. EFM will use existing copper loops and thus could offer a technology challenge to high-speed and multilink frame relay—depending, of course, on when the EFM products and services become available and how they are priced.

Modern Maturity

Clearly, MFR strengthens and expands frame relay service offerings, but is it a mid-life kicker or a mature technology's last gasp? Steve Taylor votes for the former. "Frame relay is going to be around for a long time," said Taylor. "It has no sex appeal, but it works." Nonetheless, frame relay providers are looking beyond their traditional markets for continued revenue growth.

AT&T, for example, is looking overseas. Private lines are still expensive in many countries, which creates a natural market for frame relay. "In the U.S., our frame relay order entry, provisioning, billing and so on, are all automated," reported Tim Halpin. "AT&T is moving to make global service look like domestic with the same operations systems." The goal is consistent frame relay service "from Paris, Texas, to Paris, France."

Verizon also announced international frame relay service, and Lucent is targeting international carriers with its beefed-up E1 (2.0 Mbps) and N×E1 MFR support. "We expect quite a bit of demand for MFR overseas, where E3 (34 Mbps) service is rare," said Lucent's Jim Jordan, director of product management.

WorldCom, on the other hand, hopes to reach more domestic customers with two new service offers announced in March: Bundled Frame Relay and Economy Frame Relay. With a traditional frame relay service, the customer specifies everything separately for each site—access, ports, PVCs, CPE—and receives a discount based on the level and the length of the commitment. WorldCom's new Bundled Frame Relay provides a "network in a box," including CPE, access at speeds up to T1 and "all-you-can-use" PVCs for a single price. The *prix fixe* Economy service includes access up to 56 kbps, basic service level guarantees and end-to-end management. Both packages require shorter contracts than usual. WorldCom hopes that simplicity and risk reduction will attract small and medium-size businesses, late adopters, dial-up users and stragglers who may still be stuck on X.25.

Another popular tactic—indicative of a maturing technology—is an increased emphasis on better service. WorldCom's Frame Gold service relies on an intelligent CSU, supplied by the carrier and installed at the customer premises, that enables a 30-minute or less guaranteed response time for circuit outages. Similarly, AT&T's Frame Relay Plus also uses CSUs to collect RMON-based performance information, which is then made available to customers via a Web interface.

However they do it, incumbent frame relay providers are well advised to tighten their hold on their current customers. Today, interexchange carriers (IXCs) like AT&T, WorldCom and Sprint, and international carriers like Equant, dominate the frame relay services market because their geographic reach is unlimited. Local exchange carriers (LECs) like Verizon are constrained by law to regional services, which has kept their market share small. If and when the government gives LECs the go-ahead to provide long-distance services throughout the U.S., the competitive map could shift dramatically.

With any luck, expanded competition will push frame relay prices lower and—certain vendors hope—open up the equipment market as well. "The LECs are preparing to enter the long-distance frame relay market," said Chad Dunn, director of product management at WaveSmith Networks, "which creates a virtual greenfield for switch vendors." Like the services market, today's frame relay equipment market is dominated by a few major suppliers: Lucent, Cisco, Nortel. The entry of the LECs is just the sort of pot-stirring

FIGURE 3 Frame Relay Access To IP/MPLS Backbone

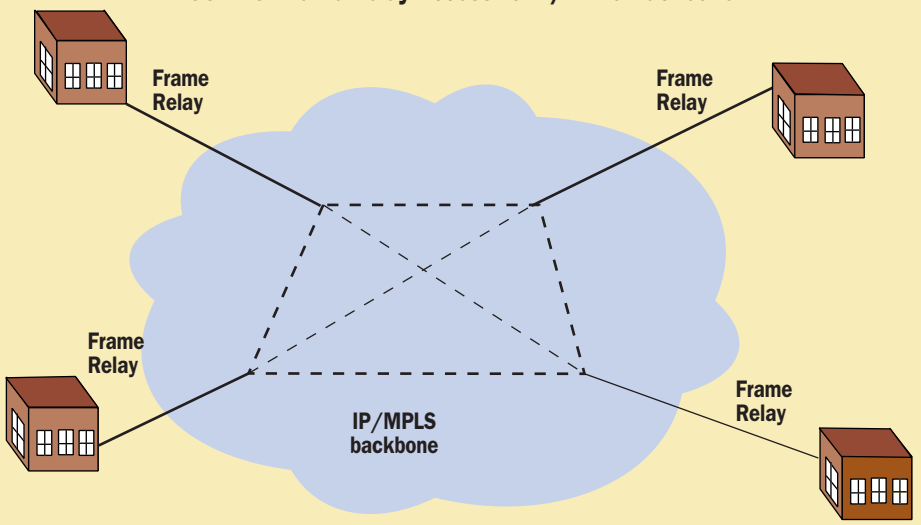

Hybrid IP/frame services can save money—depending on how your network's configured

that new companies like WaveSmith, Tasman and others need.

Of course, a cash-rich LEC could enter the long distance frame relay market by buying a struggling IXC like WorldCom or AT&T. Competition among carriers would then decrease—bad news for pricing—and equipment vendors would be left facing even fewer carrier customers.

Living With IP

In the long run, however, frame relay's continued success depends on its relationship to IP. According to Vertical Systems, IP currently accounts for 80 percent of frame relay traffic and the percentage is growing. With MPLS and IP VPNs offering an alternative to frame relay, it's in frame's best interest to cozy up to IP in as many ways as possible.

And that's just what's happening. AT&T's IP-Enabled Frame Relay Service and WorldCom's Private IP Service combine frame relay access with MPLS backbones. Sprint's IP Intelligent Frame Relay is similar, but uses virtual IP routing instead of MPLS.

With traditional frame relay, a customer needs separate PVCs between every pair of sites that communicate directly. As the number of pairs grows, the number of PVCs multiplies rapidly—the N-squared problem—and so does the company's frame relay bill. With the new hybrid services, each customer site uses just one frame relay PVC to access the network, and the network provides complete connectivity among all sites (Figure 3).

The new hybrid frame/IP services will be more valuable to customers with lots of sites and lots of traffic. For example, a low-bandwidth PVC into AT&T's IP-Enabled Frame Relay Service is priced higher than an ordinary point-to-point PVC. But at committed information rates above 1,024 kbps, there's no price difference, even though the IP-enabled PVC offers unlimited connectivity. And adding a new site entails just one more PVC, not N more. For companies with lots of traffic and many endpoints, the cost savings over traditional frame relay can be substantial (see *BCR*, March 2002, pp. 24–27).

In addition, certain applications become more feasible with hybrid frame/IP services. VOIP, for example, is difficult with pure frame, because you have to provision a separate PVC for every pair of sites that want to talk. This barrier dissolves with a routed or MPLS backbone providing complete connectivity. (The hybrid frame/IP networks will still require special packet handling for latency and jitter, not unlike the need to manage priorities in traditional frame networks.)

Yet, not everyone is convinced that the hybrids will take off. "Most private-line networks are designed as hub-and-spoke or star configurations; they aren't meshed peer-to-peer," said Rosemary Cochran, principal at Vertical Systems Group. Since most applications are written to use hub-and-spoke networks rather than meshes, Cochran doesn't see frame/IP hybrids replacing many private-line networks. Nor does she regard IP VPNs as a serious threat to frame relay for site-to-site applications. "The battle between frame relay and IP-VPNs is mainly in the press," she said.

Maybe that's why the IETF is turning the IP-to-frame relationship upside-down. Given the popularity of frame relay, IP/MPLS cannot displace ATM as a backbone technology unless it can carry frame relay traffic. So various IETF committees have produced the Martini draft, L2TPv3 and other proposals that specify how to transport Layer-2 protocols like frame relay and ATM over IP and MPLS (see *BCR*, February 2002, pp. 29–35).

These initiatives give service providers who already have IP or MPLS backbones a standard way to add frame relay service to their product portfolios. But they probably won't convince

Why would the IXCs want to abandon their ATM cores for frame over IP?

current frame relay service providers to abandon their ATM core networks. Why would they replace a working infrastructure with something new just to offer the same services? And would it really make sense to run IP over frame over L2TPv3 over IP? Incumbent frame relay service providers are more likely to prefer the faster, more capable ATM multiservice switches offered by both incumbent and new vendors. These give service providers a far less disruptive growth path, including the option to exploit MPLS for new services like IP-enabled frame.

Keep On Keeping On

By avoiding complexity and cooperating with competitors, frame relay has come a long way over the past 11 years. The same principles should help frame relay prosper as a reliable interface to new services in the future.

As Mike Walsh put it, "The real value today is not what you do *with* a technology, but what you do *between* technologies." That's why the Frame Relay Forum is working with the MPLS Forum on frame-to-MPLS interworking and soliciting contributions on frame-to-Ethernet and voice-over-packet interworking. Let's hope that along the way they don't forget the most important rule: Keep it simple, stupid!□

Companies Mentioned In This Article

Adtran (www.adtran.com)
AT&T (www.att.com)
ATM Forum (www.atmforum.com)
Cisco (www.cisco.com)
Distributed Networking Associates (www.distributed-networking.com)
Equant (www.equant.com)
Frame Relay Forum (www.frforum.com)
Larscom (www.larscom.com)
Lucent (www.lucent.com)
MPLS Forum (www.mplsforum.org)
Nortel Networks (www.nortelnetworks.com)
Quick Eagle (www.quickeagle.com)
Sprint (www.sprint.com)
Stallion Technologies (www.stallion.com)
Tasman Networks (www.tasmannetworks.com)
Verizon (www22.verizon.com)
Vertical Systems Group (www.verticalsystems.com)
WaveSmith Networks (www.wavesmithnetworks.com)
WorldCom (www.worldcom.com)