Multilinking Across The T1-To-T3 Chasm

Steven Taylor

What to do when you need more bandwidth than T1 but less than T3? Try mulitlinking.

For years, enterprises have complained about the lack of incremental bandwidth services between T1 (1.5 Mbps) and T3 (45 Mbps). While the situation has improved with the availability of cable modem, xDSL and high-speed wireless services, there is still a vacuum, particularly for business customers who need service in the 3-Mbps to 20-Mbps range.

A few metro areas have transparent LAN services (TLS) available, but the TLS service footprint is limited and, given the carriers' financial situation and the general economy, it's not likely to expand significantly any time soon. In addition, ILECs and other service providers offer "Fractional T3" services, where they install a full T3 to the premises and sell the customer a portion of the bandwidth. But again, this offering has a limited footprint.

Another option is hardware-based—i.e., connect routers to multiple WAN links, which may or may not involve multiple Internet and/or intranet connections. This works well when relatively light traffic loads need to be spread across a large number of endpoints. Each session is assigned to a link, which it uses for the duration of the session. But this approach does not help enterprises with requirements for high throughput for single-session file transfers and other bandwidth-intensive applications.

Five Basic Options

There's been growing interest in "multilink" services—bundling multiple access services into a single virtual link. Not surprisingly, there are a variety of ways to skin the multilink cat, some proprietary to particular vendors. In general, however, there are five basic methods, each with its own strengths and weaknesses.

The first, known as inverse multiplexing, operates at the physical layer. Bits are simply spread out across multiple links to provide more throughput. The biggest advantage of physical-layer inverse multiplexing is that because it's a physical-layer implementation, it is absolutely protocol-transparent. ATM, frame relay, SNA, PPP? It doesn't matter; this implementation occurs too low on the OSI stack to know or care.

The disadvantage is the lack of standardization at higher speeds. At sub-T1 speeds, combining multiple DS0 (64 kbps) circuits into a single fractional T1 circuit is technically trivial, provided the DS0s all reside in a single T1. Even if they don't, "bonding" specifications, which were developed several years ago primarily for videoconferencing applications, work well.

But as soon as you get to T1/E1 or higher links, there is no industry-standard specification defining how equipment should aggregate the links. Consequently, physical-layer inverse multiplexing requires matching equipment from a single vendor on both sides of the link—at the serving wire center and on the customer's premises.

A second approach to this problem is called Multilink PPP (ML-PPP), and it has been available for more than a decade. Since IP does not inherently support any link-layer capabilities, the link layer is usually provided either by point-to-point protocol (PPP) or frame relay. PPP provides a simple link layer, and ML-PPP extends this to provide transport across multiple parallel links.

One of the strongest points for ML-PPP is its support for dial links. In fact, many of the earliest applications for ML-PPP were for using multiple dial or ISDN lines for turbocharged Internet access. But that strength also is ML-PPP's weakness. While in theory PPP can be used to transport any type of traffic, in reality it's almost exclusively used for IP. It is not usually thought of for transporting multiple traffic types.

A third approach, developed specifically for the IP layer, is a multilink protocol for multiple IP streams independent of the link layer. The good news about this approach, led by ePipe, is that it works over a wide variety of access media. The bad news is that it is based on proprietary protocol.

The fourth approach, Multilink Frame Relay (MFR), is closely related to ML-PPP, because the two protocols are quite similar. Both are based on frame (as opposed to cell) transport, but while frame relay supports multiple *virtual* connections

Steven Taylor is publisher of Webtorials.Com, an online tutorial and technical information website. He can be reached at taylor@webtorials.com

at the link layer, PPP, only supports a single link-layer connection. The word "virtual" is critical, because PPP does not have Data Link Connection Identifiers (DLCIs), but apart from that, PPP and frame relay look a lot alike.

End-to-end MFR, which is based on the Frame Relay Forum's FRF-15 implementation agreement, provides high throughput using multiple frame relay circuits. The big advantage of FRF-15 MFR is that it doesn't require participation on the part of the service providers; indeed, it can span multiple service providers. The drawback, though, is that it only works between two endpoints; in short, it's a simpler, cheaper version of point-to-point, physical-layer inverse multiplexing.

A variation on MFR arises from the Frame Relay Forum's FRF-16 implementation agreement, which specifies details for the MFR at the User-to Network Interface (UNI) and Network-to-Network Interface (NNI). This enables multiple physical links to behave as a single frame relay UNI/NNI, and is an excellent way to increase throughput for an access link that is limited by the speed of the physical facilities. Think of it this way: You can use two access links as if they were one. This option can make sense for enterprises that are increasing access speeds at many branch locations to up to T1, and consequently need more than T1 but less than T3 at the central site.

There is, however, a major problem with FRF-16 MFR: It requires an offering by a service provider, and so far none have stepped up to the plate, at least in the U.S. That may soon change, but it's going to be some time before this option is

Multilink frame can span multiple service providers

Table 1 Sampling Of Mulitlink Products

	Physical Layer Inverse Multiplexing	Inverse Multiplexing over ATM (IMA)	Multilink PPP	Multilink Frame Relay (MFR)
Adtran www.adtran.com	ATLAS 550, 830, 890	Total Access 3000	None	None
Cisco www.cisco.com	None	7000 Series (7200, 7300, 7401, 7500, 7600)	7000 Series (7200, 7300, 7401, 7500, 7600)	7200 and 7500 series (later this year: 7300, 7600)
Larscom www.larscom.com	* Mega-T * Mega-E * Orion 4000	Orion 2000	Larscom 6000	Larscom 6000
Kentrox www.kentrox.com	None	*CellSMART 201, 202, 203 204, 205 *AAC2,AAC3	None	None
Nortel Networks www.nortelnetworks.com	None	Passport 7000 series	*Backbone Node portfolio, including Backbone Link Node and Backbone Concentrator Node *Access Stack Node *Passport 5340 *Passport Advanced *Remote Node *Passport 2430	*Backbone Node portfolio, including Backbone Link Node and Backbone Concentrator Node *Access Stack Node *Passport 5340 *Remote Node *Passport Advanced *Remote Node *Passport 2430
QuickEagle www.quickeagle.com	*DL3800 *DL3800E	None	*4240 *5842 *5844 *5840	*4250 *5850
Tasman www.tasmannetworks.com	None	None	*1004 *1200 *1400 *1450 *6200 *6300 *6302 *7030	*1004 *1200 *1400 *1450 *6200 *6300 *6302 *7030
WaveSmith www.wavesmithnetworks.com	None	DN 2100, 4100,7100	None	DN 2100, 4100, 7100

Your decision will likely be affected by what's offered in your area—and what's not

widely available. (In fairness, it should be noted that FRF-16 MFR is used, but not promoted as such, for some high-speed Internet access services, such as those offered by WorldCom. MFR also is offered in Europe by Deutsche Telekom.)

Finally, there's a fifth approach—Inverse Multiplexing over ATM (IMA). With IMA, a single ATM UNI can be supported over multiple physical links. From a technology perspective, IMA can provide a smoother flow of information because of the fixed length of ATM cells. This removes some of the jitter that can occur, at least theoretically, when drastically different frame sizes are sent. However, because it's an ATM-based offering, IMA must contend with the overhead penalty ATM imposes, compared with frame-based technologies.

That said, IMA's biggest advantage is that it's widely available, and the service providers have considerable experience. Service providers like AT&T have used IMA for several years.

Conclusion

Which option is best? Not surprisingly, the answer is: It depends. All the options accomplish the fundamental goal; there is no overwhelming technical reason to choose one over the other. And a critical decision factor is outside your control: Which options does your service provider offer?

Some may offer only one, others may offer a set of options, and the offers change over time. For example, for years, AT&T insisted that frame relay would stop at T1; anything above T1 would be based on ATM. Today, guess what? AT&T offers T3 frame services. And while, so far, it has stopped short of offering MFR, it continues to evaluate whether MFR should become an offering within its service suite.

Ultimately, the choice of which multilink service to use comes down to personal opinions, availability and pricing. And there's no reason why an enterprise needs to choose only one type of multilink service. Regardless of which form—or forms—you choose, one thing is clear: Multilink services will play an important role as we evolve to higher speed services□

Companies Mentioned In This Article
AT&T (www.att.com)
Deutsche Telekom (www.telekom.de)
ePipe (www.ml-ip.com)
WorldCom (www.worldcom.com)