Cutting the Cost of Doing Business Using Inverse Multiplexing

With the ever-increasing importance of bandwidth-intensive communications in their day-to-day business, more and more small-to-medium-sized enterprises are outgrowing their T1 or E1 wide area connections. When this happens, they find themselves facing a "bandwidth gap". The next-higher-rate service typically available from their service provider is T3 or E3—a very large step up in bandwidth (roughly, from 1.5 to 45 Mbps for T1/T3, or from 2 to 34 Mbps for E1/E3) that comes with a correspondingly large increase in price. It should not be surprising, then, that an enterprise that has outgrown a T1 or E1 wide area connection often cannot cost-justify leaping the bandwidth gap all the way to T3 or E3.

While service offerings scaled to the intermediate bandwidth needs of these enterprises have been slow to appear, a relatively straightforward solution exists and can be easily implemented by the enterprise. Employing inverse multiplexing (called "inverse muxing" or "imuxing" for short), an enterprise can build a higher-speed wide area connection by bundling multiple T1 or E1 circuits into a single high-speed link. Thus, inverse multiplexing offers bandwidth-hungry enterprises access to scalable, economical WAN connections of up to 12 Mbps (eight T1s) or 16 Mbps (eight E1s).

Typical Savings Using Inverse Multiplexing

A recent study by RHK, an independent analyst organization that reports on the telecommunications market, found that the cost of a T3 circuit is 11 to 12 times the cost of a T1 circuit. For example, an enterprise might be paying $240 a month for T1, with its service provider offering T3 at $2880 a month—twelve times the T1 cost. (These costs are typical, as reported by RHK in 2002.)

If the enterprise has outgrown its T1 WAN connection, it can pay the price for T3, or it can implement inverse multiplexing for substantial monthly savings. Inverse-muxing four T1s produces a WAN connection of 6 Mbps at a cost of $960 a month, a 67% savings over the T3 alternative. A 9 Mbps inverse-muxing connection (six T1s) saves 50% over T3, while a 12 Mbps connection (eight T1s) saves 33%. Annual savings for four inverse-multiplexed T1s compared to a T3 connection amount to over $23,000.

An Overview of Inverse Multiplexing Technologies

The basic concept of inverse multiplexing, illustrated in Figure 1, is simple. A data stream that is too large for a single transmission path is broken into smaller pieces, and the pieces are transmitted over separate transmission paths to the receiving end, where the pieces are reassembled into the original data stream. Since the 1990's, a number of devices have employed this strategy in different ways to create otherwise-unavailable multi-megabit wide area links using readily-available T1 or E1 circuits. Today there are several inverse-multiplexing solutions that address different performance and interoperability issues.

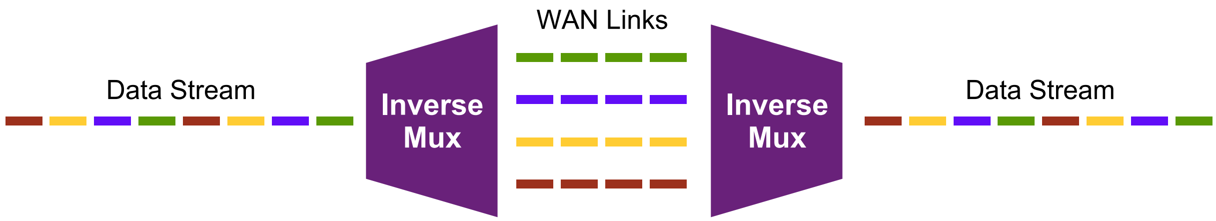

Figure 1. The Basic Concept of Inverse Multiplexing

Bit-Based Inverse Multiplexing

The simplest inverse multiplexing mode, and the first to be used in wide-area networking products, employs a bit-based approach in which individual bits are distributed in sequence to the multiple transmission circuits on a round-robin basis. Because the inverse multiplexer is operating at the Layer 1 level of the serial bit stream, this method is entirely independent of the protocols of the data being transmitted and of the network equipment. Thus, it can be used in virtually any application. Relatively little overhead is required to enable reconstruction of the received bits into the original sequence. And the inverse multiplexer, while adding another device to the network infrastructure, can provide a high level of performance monitoring and diagnostics for the individual circuits as well as the multi-link bundle.

There is no industry standard for bit-based inverse multiplexing; therefore, products implementing this approach must use a proprietary protocol for disassembling and reassembling the data stream. However, since the inverse multiplexing operation is transparent to both the local and wide-area networks, only the inverse multiplexer at each end of the link must recognize and support the proprietary protocol.

The desire for standards-based inverse multiplexing has resulted in three major higher-layer solutions. These approaches have particular benefits and drawbacks directly related to the protocols on which they are based.

Inverse Multiplexing for ATM (IMA)

For ATM-based networks, the ATM Forum has standardized a method of inverse multiplexing in which the data is distributed to multiple T1 or E1 circuits on a cell-by-cell basis. The advantages and disadvantages of Inverse Multiplexing for ATM (IMA) are essentially the advantages and disadvantages of ATM. The protocol provides some advanced Quality of Service (QoS) capabilities, but at the expense of a high percentage of bandwidth being used for overhead. IMA can be an excellent solution in a wide-area ATM network, but it is not typically the ideal choice for other environments.

Multilink Frame Relay (MFR)

Multilink Frame Relay is an inverse multiplexing standard for Frame Relay environments developed by the Frame Relay Forum. In MFR, multiple virtual T1 or E1 circuits are bundled to create a higher-rate channel which is recognized as a single physical interface at the data link layer. The data stream is distributed across the virtual circuits on a packet-by-packet basis.

MFR is implemented in a Frame Relay environment with MFR-compatible Frame Relay switches. Its advantages include standards-based Service Level Agreements (SLAs), support for variable frame sizes and fragmentation, low latency, and minimal overhead bandwidth.

Multilink PPP (MLPPP)

Like Multilink Frame Relay, Multilink PPP is a packet-based inverse multiplexing method employing multiple Point-to-Point Protocol (PPP) links as standardized by the Internet Engineering Task Force (IETF) in RFC1990. MLPPP was developed in response to some of the limitations of proprietary load sharing schemes employed on many routers. In addition to vendor-independence, its advantages include efficient frame mapping, low overhead bandwidth, connectionless IP environment, and packet fragmentation to prevent monopolization by large data packets. Unlike IMA and MF, MLPPP is not dependent on switches supporting a particular protocol; rather, it is like bit-based inverse multiplexing in requiring only that the receiving end supports MLPPP.

Choosing an Inverse Multiplexing Solution

The bandwidth gap is a very real problem: as RHK observed, "service providers tend to put customers into T1 or T3 buckets." Inverse multiplexing offers a reasonable growth path in the multi-megabit bandwidth range which is critical to so many small-to-medium-sized enterprises. The variety of inverse multiplexing options currently available allows each enterprise to tailor an inverse multiplexing solution to its particular situation and needs.

Larscom has been a leading provider of inverse-mulitplexing solutions since pioneering bit-based inverse multiplexing in the early 1990's. Still strongly committed to inverse multiplexing, Larscom currently offers bit-based inverse multiplexing in its Mega-T, Mega-E, and Orion 4000 products, IMA in its Orion 2000, and MFR and MLPPP in its Larscom 6000.

For more information about inverse multiplexing solutions, call us at 888.LARSCOM or 408.941.4000, visit *www.larscom.com*, or contact your Larscom representative.

ML03-IMX-01, November 2002. Printed in USA.