



**SPEAKEASY
BONDED T1s**

July 2009

Corporate Office
800.556.5829

East Coast Regional Office
866.575.4941

Partner Program
888.886.6116

1201 Western Ave
Seattle, WA 98101

www.speakeasy.net

TABLE OF CONTENTS

Table of Contents1
Overview2
What is Bonding?3
Inverse Multiplexing4
 Bit-Based Inverse Multiplexing5
 Inverse Multiplexing for ATM (IMA)5
 Multilink Frame Relay (MLFR)5
 Multilink PPP (MLPPP)5
Multilink Point-To-Point Protocol (MLPPP)6
 MLPPP Packet Transmission6
 MLPPPoA and MLPPPoFR7
The Speakeasy Advantage7
For Further Reading8

- Deleted: 3
- Deleted: 4
- Deleted: 5
- Deleted: 5
- Deleted: 5
- Deleted: 5
- Deleted: 6
- Deleted: 6
- Deleted: 7
- Deleted: 7
- Deleted: 8

OVERVIEW

Small- and medium-sized businesses have been using T1 circuits for their data transmission needs since the '60s. In the past, the 1.5Mbps of bandwidth available with a single T1 circuit was sufficient for most businesses. Today's voice, data, audio and video applications demand increased bandwidth and the multi-megabit access products that can provide such capacity.

Businesses attempting to meet their need for increased bandwidth frequently believe that the next step in bandwidth access is a fractional DS3, starting at 9-15Mbps of bandwidth, or a full DS3, which provides 45Mbps of bandwidth. The cost of this connectivity option jumps from an average of \$500 per month for a T1 line to \$4,000 or more per month for a fractional DS3 line. Not only is this option cost-prohibitive for many businesses, but it often far exceeds the actual amount of bandwidth needed. Additionally, DS3 infrastructure is not available to most businesses in the U.S., and where available, can take months to provision.

Bonding is a cost-effective alternative for most businesses. Existing copper T1 infrastructure is readily available and relatively inexpensive. Bonding T1 lines together provides a virtual multi-megabit access path with high throughput and low latency. Speakeasy Bonded T1 comes in 3.0 Mbps, 4.5 Mbps or 6.0 Mbps, providing two to four times the speed of single T1 – a good intermediary step between a single T1 circuit and a more costly DS3.

There are many benefits to bonding T1 lines. Bonding simplifies routing tables because packets are routed on a per bundle basis instead of a per circuit basis. Aggregated links also provide some redundancy in that if one circuit goes down, the other(s) will still be able to pass traffic up to their aggregate capacity. Per packet load distribution between the member circuits is handled transparently by the technology that implements the circuit bundle.

This paper defines bonding, provides an overview of the methods available for implementing bonding on a network, and explains how Speakeasy's bonding solution creates maximum value for our customers.

WHAT IS BONDING?

Bonding, or NxT1, is a link aggregation technology where two or more circuits are aggregated into one channel that functions as a single, higher-bandwidth logical link (see Fig. 1). Individual, or member, T1 circuits are bonded together on the customer's router and on the client aggregation router (ERX) using multilink point-to-point protocol (MLPPP) or another technology such as multilink frame relay (MLFR).

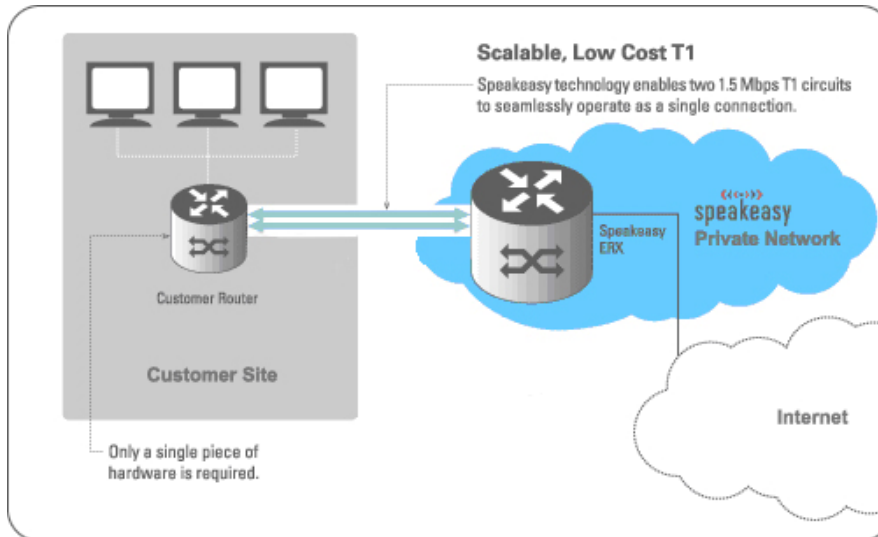


Figure 1: Individual T1 circuits are bonded together on the customer's router and on the ERX to function as a single connection.

INVERSE MULTIPLEXING

Bonding is also called inverse multiplexing. Inverse multiplexing (IMUX) speeds up data transmission by dividing a data stream into multiple concurrent streams that are transmitted at the same time across separate channels, such as T1 lines, and then re-sequenced at the other end into the original data stream (see Fig. 2). Just the reverse of ordinary multiplexing, which combines multiple signals into a single data stream, IMUX is commonly used where data in a high-speed local area network (LAN) flows back and forth into a wide area network (WAN) across the "bottleneck" of a slower line such as a 1.5Mbps T1. Using multiple T1 lines, the data stream can be distributed across all of the lines at the same time via a round-robin load-sharing technique. Packets are sent across successive T1 circuits as shown in Figure 3. At the receiving end, the data is recombined into a single stream.




Figure 2: Inverse multiplexing

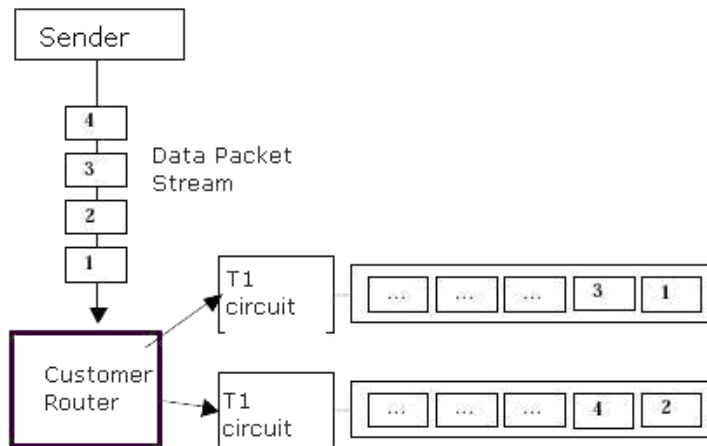


Figure 3: Round-robin load-sharing of packets

Some common methods, or protocols, for implementing inverse multiplexing include hardware (bit-based IMUX), over ATM (IMA), multi-link frame relay (MLFR), and MLPPP. Each method has particular benefits and drawbacks.

Bit-Based Inverse Multiplexing

Bit-based IMUX, the simplest method of implementing inverse multiplexing, distributes individual bits in sequence to the member circuits on a round-robin basis. This method operates on Layer 1, and is therefore independent of the protocols used to transmit the data and the network equipment. While the IMUX equipment adds another piece of hardware to the network, this method requires relatively little overhead to re-sequence the received bits, provides visibility to the individual circuits as well as the bundle, and can be used with virtually any application. However, since there is no industry standard for bit-based IMUX, products using this method must use a proprietary protocol that is recognized by the IMUX equipment on both ends of the links.

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 (MLFR)

MLFR is an IMUX standard developed by the Frame Relay Forum for FR environments. In this method, multiple T1 circuits are bundled together to create a multi-megabit channel where the channel functions as a single interface at the data link layer and distributes a data stream across the member circuits on a packet-by-packet basis. Frame relay switches must be MLFR-compatible. MLFR provides support for variable frame sizes and fragmentation, low latency, and minimal overhead bandwidth.

Multilink PPP (MLPPP)

MLPPP is also a packet-based IMUX method. As standardized by the Internet Engineering Task Force (IETF) in RFC 1990, MLPPP bundles multiple point-to-point protocol (PPP) links together to create a single multi-megabit channel. This method was developed to alleviate some of the limitations of proprietary load sharing protocols used by many routers. MLPPP provides vendor-independence, efficient frame mapping, low overhead, packet fragmentation and a connectionless IP environment. It does not require switches that support a particular protocol, but it does require that both ends of the bundled link support MLPPP. This method is discussed in detail in the following section.

MULTILINK POINT-TO-POINT PROTOCOL (MLPPP)

Multilink Point-to-Point Protocol (MLPPP) is a method of splitting, recombining, and sequencing datagrams across multiple logical data links. MLPPP allows a router to join discrete circuits and makes it possible to divide a data stream between multiple links in such a way that the remote router can re-sequence the packets in the original order upon reception.

Traditionally, MLPPP technology has been used for Integrated Services Digital Network (ISDN). In ISDN, and other access technologies, MLPPP takes multiple PPP links and “bonds” them into a logical multi-link PPP bundle to derive more bandwidth than from a single physical connection. MLPPP technology provides the ability to have a single session consume more bandwidth than one member circuit can provide. For example, an individual host can utilize the full 3Mbps capacity of a bundle that is comprised of two 1.5Mbps T1 circuits.

An MLPPP-enabled CPE router takes incoming Ethernet frames from the LAN, strips them of their MAC addresses and adds an MLPPP header along with packet sequence numbers. Using the round-robin load-sharing technique described earlier, the MLPPP frames are then distributed across all of the access circuits that comprise the bundle. The MLPPP interface at the provider edge of the link then re-sequences the incoming packet stream.

The main advantage of MLPPP is that it does not require any specialized hardware such as a dedicated inverse multiplexer, frame-relay switch, or special IMA cards. Because they can implement MLPPP with software run over standard interfaces, businesses circumvent capital expenditures for dedicated equipment, as well as saving the space, power, and firmware updates that would be required to maintain new equipment. Less equipment also means fewer points of failure and fewer measures to maintain environmental control.

Using MLPPP to bond multiple links into a single connection can provide redundancy in the event that one or more of the T1 links in the bundle experience error, test or failure conditions and at least one link continues to function properly. These remaining link(s) will continue to pass traffic up to their aggregate capacity. Removing the problem circuit(s) from the bundle, which decreases the bundle bandwidth accordingly, and then reinstating the circuit(s) once the anomaly has cleared allows the bundle to continue to function. Standards-based MLPPP protocol provides a method to ensure bundle recovery when individual circuits have problems, increasing overall network reliability.

MLPPP Packet Transmission

Network Protocol packets are first encapsulated (but not framed) according to normal PPP procedures. With fragmentation enabled, large packets are broken up into multiple segments sized appropriately for the multiple physical links. Although it would otherwise be permitted by the PPP specification, implementations must not include the Address and Control Field in the logical entity to be fragmented. A new PPP header consisting of the Multilink Protocol Identifier and the Multilink header is inserted before each section. (Thus the first fragment of a multilink packet in PPP will have two headers, one for the fragment, followed by the header for the packet itself.) Systems implementing the multilink procedure are not required to fragment small packets. There is also no requirement that the segments be of equal size, or that packets must be broken up at all.

The PPP encapsulation is used to disambiguate multi-protocol datagrams. This encapsulation requires framing to indicate the beginning and end of the encapsulation. The frame consists of the protocol field, the information field and the padding field. The protocol field is one or two octets, and its value identifies the datagram encapsulated in the information field of the packet. The information field is zero or more octets, and it contains the datagram for the protocol specified in the protocol field. The maximum length for the information field, including padding, but not including the protocol field, is termed the Maximum Receive Unit (MRU), which defaults to 1500 octets. By negotiation, consenting PPP implementations may use other values for the MRU. On transmission, the Information field may be padded with an arbitrary number of octets up to the MRU. It is the responsibility of each protocol to distinguish padding octets from real information.

MLPPPoA and MLPPPoFR

MLPPPoA allows MLPPP connections across an ATM or ATM/Frame Relay inter-working network over one or more ATM stations. Stations may be associated with one or more ATM ports over same or different interface types. MLPPPoA enables PVC link aggregation and bandwidth increase under MLPPP control. MLPPP over multiple high-speed VCs facilitates transporting both real-time and data traffic more efficiently by distributing data between member circuits.

MLPPPoFR allows MLPPP connections across a Frame Relay or Frame Relay/ATM inter-working network over one or more Frame Relay Bypass stations in accordance with RFC 1973. Bypass stations may be associated with one or more Frame Relay ports. Port physical interfaces may be the same or different types. This feature enables PVC link aggregation and bandwidth increase under MLPPP control. MLPPP Link Fragmentation and Interleaving (LFI) allows QoS support when FR mechanisms cannot be used.

A Permanent Virtual Circuit (PVC) based frame relay/ATM (FR-ATM) service Inter-working Function (IWF) provides flexibility for customers who use FR and ATM services. FR-ATM PVC Service IWF, as defined in the Frame Relay Forum's FRF.8.2, performs protocol conversion by unwrapping ingress bound variable length FR frames or fixed length ATM cells. The conversion maps the inter-working variables to appropriate ATM or FR parameters and then re-builds fixed length ATM cells or variable-length FR frames as appropriate.

THE SPEAKEASY ADVANTAGE

Our all-Juniper network, together with the Samsung router on the customer edge, enabled Speakeasy to implement a combination of MLFR- and MLPPP-based bonding solutions, allowing us to offer our customers cost-effective, multi-megabit connectivity with low latency, low overhead and no additional specialized equipment.

Because we own and manage our network, we can easily implement new technologies to improve the connectivity options we offer our customers. The flexibility of our platform allows us to combine technologies, such as Quality of Service (QoS), to enhance current voice and data products while maintaining the ability to enable new technologies such as IPv6 and multicasting in the future.

Speakeasy Voice customers are protected from many of the call quality issues common with other providers, such as jitter, echo, and packet loss resulting in

dropped calls. Voice traffic is carried on the Speakeasy [network](#) and our privately peered partner networks without ever touching the public Internet. Unlike other providers who don't have full control over the networks over which calls are placed, Speakeasy provides both voice service and connectivity over one network. This allows us to use VQ Technology to prioritize voice calls over other data for crystal-clear call quality.

FOR FURTHER READING

To read more about the subjects discussed in this paper, please see the links below.

The PPP Multilink Protocol <http://www.rfc-editor.org/rfc/rfc1990.txt>

PPP Multiplexing <http://www.rfc-editor.org/rfc/rfc3153.txt>

The Point-to-Point Protocol (PPP) <http://www.rfc-editor.org/rfc/rfc1661.txt>

Frame Relay/ATM PVC Service Interworking Implementation Agreement
http://www.mfaforum.org/tech/fr_ia.shtml

PPP in Frame Relay <http://www.rfc-editor.org/rfc/rfc1973.txt>

PPP over Asynchronous Transfer Mode Adaptation Layer 2 (AAL2)
<http://www.rfc-editor.org/rfc/rfc3336.txt>

Point-to-Point Protocol (PPP) Vendor Protocol <http://www.rfc-editor.org/rfc/rfc3772.txt>

The PPP Bandwidth Allocation Protocol (BAP) and The PPP Bandwidth Allocation Control Protocol (BACP) <http://www.rfc-editor.org/rfc/rfc2125.txt>

Multiprotocol Interconnect over Frame Relay <ftp://ftp.rfc-editor.org/in-notes/rfc2427.txt>