Technology Overview: Frame Relay

In this report:

Editor's Note

The demand for switched high-speed data services is now emerging as a major user requirement. Frame relay is a higher performance service compared to traditional X.25 packet switching and is positioned by the four key vendors in this market (Northern Telecom, Digital Equipment Corp., StrataCom, and cisco) as a LAN interconnection service. Frame relay technology is an evolutionary step beyond X.25 for improving packet network efficiency and accommodating more efficient applications, such as wide area interconnection of LANs at 56K bps and 1.544M bps rates. Networks based on frame relay provide communications at up to 2.048M bps (for Europe), bandwidth on demand, and multiple data sessions over a single access line. For information on Sonet and related standards, see Report CA40-010-801, "An Overview of Sonet-Based Systems."

Report Highlights

Frame relay is positioned to improve communication performance through reduced delays, more efficient bandwidth utilization, and decreased communications equipment cost. An X.25 public data network introduces a 200-ms. delay or more, whereas frame relay can reduce that delay to about 20 ms. This implies that frame relay can improve performance for existing communication resources, such as asynchronous terminals. To get maximum benefit from frame relay without incurring large equipment or communication charges (i.e., for dedicated T1 links between sites), the service must be tariffed by a carrier. Only one carrier has announced such a service to date.

This report was prepared exclusively for Datapro by Daniel Minoli, adjunct assistant professor at New York University's Information Technology Institute and full-time data communications strategic planner. He recently published *Telecommunication Technology Handbook,* highlighting key communication technologies affecting the scene of the 1990s.

Description

Traditional packet switching is a technology of the mid-1960s for solving the networking problems of that era. At that time, bandwidth was scarce and networks maximized the efficiency in transport. With the widespread availability of fiber cable, whose intrinsic traffic-carrying capacity has been doubling every two to three years, the need to maximize efficiency at the expense of end-to-end delay and switching node complexity is no longer imperative. The 1960's bit error rate also left a lot to be desired, with 10^{-6} stretching the technical limit. Fiber now routinely provides 10^{-8} to 10^{-9}, and Forward Error Correction (FEC) can improve that further. Error-prone circuits necessitated complex error checking and recovery procedures at each network node.

Hence, X.25 packet standards assume that the transmission media is error prone. To guarantee an acceptable level of end-to-end quality, error management is performed at every link by a fairly sophisticated but resource-intensive link protocol of the High-level Data Link Control (HDLC) family. HDLC provides core functions, including frame delimiting, bit transparency, error checking (with Cyclic Redundancy Checking), and error recovery, and other functions.

In frame relay, error correction and flow control are handled at network end points. Frame relay accelerates the process of routing packets through a series of switches to a remote location by eliminating the need for each switch to check each packet it receives for errors before relaying it to the next switch. Instead, only the switch that receives a packet from a sending device and the one passing it to the receiving device check for errors.[1] Alternatively, that function can be relegated in its entirety to the end-users' customer premises equipment (CPE). This error treatment increases performance and reduces bandwidth requirements, which in turn can—in principle—reduce communications costs and decrease the number of packet handling devices needed in a network.[2]

Frame relay is a multiplexed data networking service supporting connectivity between CPE (such as routers and bridges) and, eventually (when the service is tariffed), between CPE and carrier networking equipment. Switch manufacturers are now in a selling mode to invest in the technology.[3,4] Frame relay will be implemented on such products as LAN bridges, routers, and T1 multiplexers. Frame relay is a connection-oriented technology supporting packets of variable length.

Frame relay can be considered as the successor to X.25 and, like X.25, it specifies the interface between customer computers and a network, whether public or private. This interface specification is described in principle in CCITT Recommendation I.122 (1988), although that specification is only a framework document. It must be further codified by a series of implementors' agreements, particularly for interoperability testing. Obviously, the early mistakes of noninterconnecting "X.25-conformant" networks must be avoided for the mainstream data comm community to accept the technology.

The cost of dedicated T1 facilities is still fairly high. A typical interexchange carrier (IXC) tariff is $1,800 plus $10 per mile for metropolitan range distances (0 to 50 miles), and $2,025 plus $7.50 per mile for regional or national distances (101 miles or more). In addition, one must add the cost of the local loops at both ends of the circuit. The cost of a local loop varies depending on location: it consists of a fixed portion plus a mileage portion. The charge typically ranges from $300 to $800, depending on carrier and mileage.

Therefore, using a frame relay-configured bridge or router over a dedicated T1 link is not advantageous. On the other hand, such an interface used in conjunction with a fast-packet backbone multiplexer could be cost effective, since the user can obtain from the backbone needed bandwidth on a demand, rather than on a preallocated (and inefficient), basis. StrataCom (Campbell, CA) has the market lead in frame relay.

Frame Relay versus Available Technologies

Initially, frame relay was developed by the CCITT as an ISDN packet mode bearer service with logically separate control plane (C-plane) and user plane (U-plane) information. In the C-plane, all signaling capabilities for call control, parameter negotiation, etc., were contemplated to be based on a set of protocols common to all ISDN telecommunication services. In the U-plane, the basic bearer service provided in I.122 is the unacknowledged order-preserving transfer of data units from the network side of one user-network interface to the network side of the other user-network interface. The frame relay frame format is based on CCITT Q.921 (which corresponds to ANSI T1.602). New enhancements of I.122 are under way, as discussed later.

According to U.S. spec T1.606,[5] the following frame relay applications are conceivable:

1. Block interactive data applications. An example is high-resolution graphics (e.g., high-resolution videotex and CAD/CAM). The main characteristics of this type of application are low delays and high throughput.
2. File transfer, intended for very large file transfers. Transit delay is not as critical for this application as it is for the first application. High throughput might be necessary to produce reasonable transfer times for large files.
3. Multiplexed low bit rate. The multiplexed low bit rate application exploits the multiplexing capabilities of the Layer 2 protocol to provide an economical access arrangement for a large number of low bit rate applications. The low bit rate sources may be multiplexed onto a channel by a Network Termination (ISDN) function.
4. Character-interactive traffic. An example of character-interactive traffic is text editing. The main characteristics of this type of application are short frames, low delays, and low throughput.

Of these four applications, the first two are feasible because of frame relay's channel speed. The last two, although also feasible in conjunction with the speed, are somewhat more poised to exploit frame relay's intrinsic features. Figure 1 depicts the positioning of frame relay in the larger context of available internetworking technologies. As shown, frame relay supports bursty traffic at medium speeds. Technologies competing with frame relay include T1 links, fractional T1, fractional T3, ISDN Primary Rate service, H0, H11, Switched DS1, and, finally, Switched Multimegabit Data Service (SMDS).

The frame relay interface is a mix of evolving ANSI and CCITT ISDN data link layer standards that could eventually eclipse the X.25 standards. Frame relay reduces the protocol processing overhead inherent in X.25 and is reasonably well suited for routing LAN traffic. Frame relay provides both a private and a switched virtual circuit. Permanent virtual circuits establish a fixed path through the network so the receiving end can quickly reassemble a file. X.25 virtual circuits can route packets over different circuits, meaning packets can be received out of order. This forces the receiving end to reorder the packets before it can reassemble the file.

To achieve frame relay's benefits, one should employ a fast-packet switch. Fast-packet switches employ statistical multiplexing techniques allowing a channel's entire bandwidth to be applied to the transmission, but they do not allocate bandwidths to users who do not require it. These switches can be located on customers' premises as T1 multiplexers that support a private fast-packet network. The switches can also be located at a carrier's central office. Although fast-packet switches can support

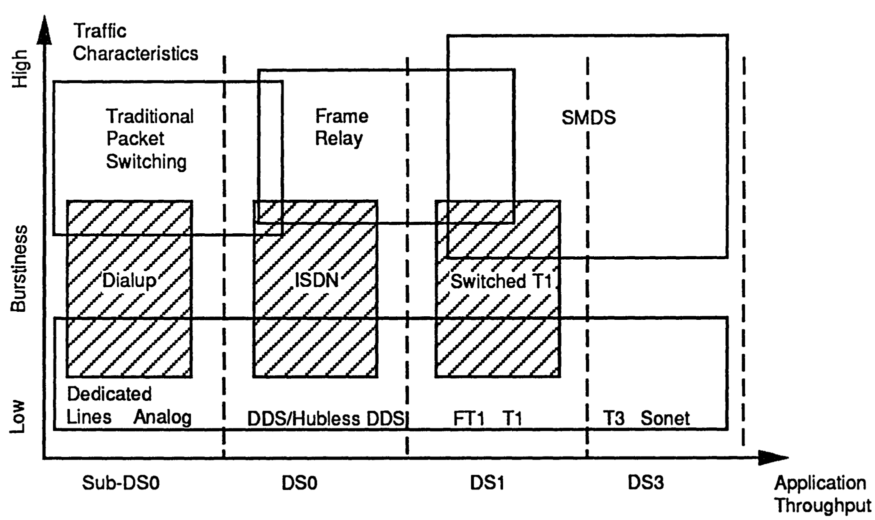

Figure 1.
Frame Relay in Context

non-frame relay traffic, together the two technologies can increase throughput between locations containing large amounts of bursty traffic.[5] Users with bursty traffic can find it advantageous to upgrade T1 equipment that uses time-division multiplexing with frame relay and fast-packet technologies. The drawback to TDM is that users must allocate the T1 circuit into individual channels, each supporting transmission of a specific data source. Since that bandwidth is allocated to only one user, it remains idle when the user does not need it.

A debate among interested parties now centers on frame relay versus cell relay, which is the underlying technology of Broadband ISDN. Some vendors have committed to frame relay; others have a program to advance cell relay; and others are pursuing both technologies. Some view frame relay as complementary, others as competitive.[4] Frame relay and cell relay are designed to meet different objectives and have therefore evolved in different directions. A simple categorization follows:[6]

Frame relay is a medium- to high-speed data interface for private networks which will be implemented in the next few years. Frame relay is being standardized at the DS1 rate.

Cell relay is a high- or very high speed switching system that supports public networks; these systems will emerge in the mid-1990s. Cell relay is discussed in the context of Sonet, which transmits at bit rates from 51.840M bps to 2.4G bps (or even 13G bps, eventually). Metropolitan Area Network-based services, however, such as Switched Multimegabit Data Service (SMDS), will be available in the same general time frame as frame relay.

SMDS is a Bellcore-proposed, high-speed (DS1 to DS3), connectionless packet switched service providing LAN-like performance and features over a metropolitan area. SMDS provides for the exchange of variable-length data units up to a maximum of 9,188 octets. It is defined as a technology-independent service, whose early availability via MAN technology is envisioned for the 1991 to 1993 time frame. SMDS is expected to be one of the first switched broadband offerings and eventually will be supported by BISDN. SMDS's Subscriber-Network Interface (SNI) complies with IEEE 802.6.

Speed of Contemporary User Applications

Customer premises equipment such as bridges and routers have operated in the 1M to 3M bps range. The maximum Ethernet throughput can be 10M bps; the effective throughput, considering protocols, can be in the 1M to 3M bps range. And while bridges and routers are quoted as transmitting in the 10,000 to 12,000 packet-per-second neighborhood, these are usually packets with no data. Therefore, while frame relay can be adequate for some LAN internetworking applications, other applications—such as CAD/CAM, medical imaging, heavy-use desktop publishing, animation, etc.—could need the higher speeds provided by SMDS and cell relay technology. The next generation of bridges and routers will process 50,000 packets per second in 1991; by 1992, many companies will require 100,000 packets per second.[7]

IBM has introduced 16M bps token-ring systems, for which a DS1-rate bridge/router may not be adequate for some applications. FDDI routers and bridges must operate at much higher rates. These 100M bps systems could become more prevalent, since the FDDI standards are practically complete, and FDDI might be deliverable over twisted pair. In addition, FDDI now interworks with Sonet, implying that there may be an impetus to introducing FDDI routers/bridges. This, in turn, could require high-throughput internetworking. Users may wonder how well a 1M bps frame relay service can bridge LANs operating at 100M bps. Moreover, FDDI rates are too low for some users (for example, in supercomputer environments); routers capable of sustained speeds of 800M bps are already available. Cell relay must operate at a high speed to support public network traffic, which when aggregated can increase to several orders of magnitude greater than any individual customer or one or more bridge/router(s).

IBM users now remotely extend the mainframe channel, another major high-speed application (see Table 1). The IBM channel has traditionally operated at 3M bps per second; it was raised to 4.5M bps in the late 1980s, and again to 10M bps in 1990. It appears unlikely that frame relay could effectively support these needs; BISDN/ATM would be better suited.

Frame Relay Limits

Frame relay is considered part of narrowband ISDN. Many observers have already pinned their

Table 1. Applications of Channel Extender Technology

Application	Explanation
Distributed Data Processing	Different portions of the data can reside in different hosts' disk drives, spread throughout the country. With channel extenders(*), all hosts have quick access to all the information.
Instantaneous Data Backup	Large amounts of data can be dumped in realtime to a secure location mirroring all database operations. This eliminates the need to generate, manage, and ship tapes, which is a labor-intensive activity. This eliminates the current bottlenecks of occasional dumps of small selected data sets (or portions thereof) over slow communications lines.
Dual Center Transaction Processing	With appropriate programming, each transaction can be maintained in two or more hosts for fully mirrored operation with negligible delays.
Darkened Data Center	All the primary mainframe equipment is placed in a remote, secure, and darkened facility, possibly in a suburban area. Channel extenders can be used to connect remote users. Minimum personnel would be needed.

(*) Channels operate at 3-, 4.5-, and 10M bytes per second.

hopes for data on Broadband ISDN, after a 10-year search for data services applications over narrowband ISDN failed to identify major new opportunities.

Frame relay is often touted as an "ideal" way to link LANs over wide area backbone networks.[8] At a purely technical level, however, since frame relay is a connection-oriented technology and LANs are connectionless, a connectionless service is the ideal way to interconnect them. One also should avoid developing entire technologies and deploying networks catering to a single application, like LAN interconnection.

Frame relay improves traditional X.25 packet switching for data communications. Cell relay supports the sophisticated services likely present in a 1990's organization, including data, voice, video, facsimile, high-quality image and graphics, and integrated messaging. Network managers will probably have to provide an integrated corporate communications infrastructure because of the two major business "discoveries" by senior executives in the 1980s:

1. Managing the network accounts for a large portion of the telecommunications expense. There is a strong desire to reduce network management complexity and financial burden. Integrated networks are easier to maintain and cheaper to manage. A smaller number of management systems should suffice, particularly with the commercialization of ISO-based network management standards.
2. Integrated networks cost less in terms of transmission facilities' recurring expenses because of the intrinsic economies of transport components. For example, for the cost of three to five DDS lines, one can already acquire a T1 facility. For the cost of three to five T1 lines, one can obtain a T3 facility.

Integration requires engineering simplification: Why establish parallel networks for each new service when they can all be supported on a common platform? Why go to all the trouble of designing, planning, installing, operating, and managing different networks like SNA, X.25, LANs, voice, etc.?

Consider a need to connect different user groups in two separate locations. For instance, a company located in Washington, DC, and Boston must connect similar islands of interest between the two locations: sales, engineering, and finance. Traffic between these three groups consists of four erlangs each at 0.01 blocking.

Employing the Erlang B model (for simplicity), the company would require 10 channels/trunks for each community, or 30 trunks between Washington and Boston. Assuming the two sites were 501 miles apart, the cost of 30 analog voice grade lines is $14,527 per month (for the IXC component); using two T1 lines would cost $11,564 per month. If the company can aggregate these three user communities, however, it needs only 20 channels for 12 erlangs of traffic at 0.01 blocking. This implies that it can use a single T1 line for only $5,782 per month, or 39 percent of the nonintegrated solution's cost.

Ultimately, such a choice hinges not merely on technologies, but on what users want, when they want it, and with what applications. As this report implies, frame relay institutionalizes another drawback of X.25 packet switching: it describes an interface specification. The equipment vendors can still utilize proprietary internal protocols, such as internal transport, routing, and flow control protocols, forcing a private network user to furnish the entire network with equipment from the same vendor. By contrast, the cell relay technology specified in the BISDN Asynchronous Transfer Mode is open by design.

Retrofitting a circuit switched T1 mux with frame relay interfaces does not deliver frame relay's intrinsic benefits. With circuit switching, the user must preallocate some (or, if desired, all) bandwidth to the frame relay service, whether the service uses that bandwidth or not. An efficient utilization of the technology over a private backbone network would require an internal fast-packet technology. By letting all applications compete for the backbone bandwidth, a frame relay application can access the entire bandwidth when it has data to transmit; a frame relay application on a circuit switched multiplexer can only access a fraction of the total bandwidth.[9]

Attempts to improve on X.25 have been under way since the early 1980s, with no apparent success. Several vendors pursued a concept called "burst switching." Users may be left wondering whether frame relay will follow the route of burst switching, which never materialzed.

A Protocol View of Frame Relay

Recommendations CCITT I.232 and I.462 (X.31) describe packet mode bearer services supported by an ISDN; I.462 (X.31) specifies the procedures for virtual call and permanent virtual circuit bearer services. ISDN also considers new packet-switching technologies in addition to these traditional X.25 packet modes. Three potential services proposed for standardization by the CCITT in Recommendation I.122 ("Framework for providing additional packet mode bearer services") are the following:

1. Frame relaying 1 (FR-1—no functions above core Data Link functions are terminated by the network; if needed, such functions are terminated only end to end);
2. Frame relaying 2 (FR-2—no functions above the core Data Link functions are terminated by the network; I.441 upper functions are terminated only at the end points); and
3. Frame switching (the full Recommendation I.441 protocol is terminated by the network)

FR-1 can be provided over permanent virtual circuits (PVCs) or switched virtual circuits. With PVCs, no call setup establishment is needed on a per-packet or session basis, since the address fields are agreed upon when the user subscribes to the service. In FR-1 the network has no knowledge of the end-to-end protocol. The LAP-D core functions include the following:

- Frame delimiting, alignment, and transparency
- Frame multiplexing/demultiplexing using the address field
- Inspection of the frame to ensure that it consists of an integer number of octets prior to zero bit insertion or following zero bit extraction.
- Inspection of the frame to ensure that it is neither too long nor too short.
- Detection of transmission errors.

Under ISDN, the new packet mode bearer services described above possess these characteristics:

- All C-plane procedures, if needed, are performed in a logically separate manner using protocol procedures integrated across all telecommunications services. Namely, Q.931 will be used to set up and tear down the service; the C-plane is used to establish global address mapping.
- The U-plane procedures share the same Layer 1 functions based on Recommendations I.430/I.431. Moreover, they share the same core procedures.

On the user side, Recommendation I.430 or I.431 provides Layer 1 protocol for the user (U-) control

(C-) planes. The C-plane uses the D-channel with Recommendations I.441 and I.451 extended as Layer 2 and 3 protocols, respectively. In the case of permanent virtual circuits (PVCs), no realtime call establishment is necessary, and any parameters are negotiated at subscription time. The U-plane may use any channel on which the user implements at least the lower part (the core functions) of Recommendation I.441.

Frame Relaying Service

The term *relay* implies that the Layer 2 data frame is not terminated and/or processed at the end points of each network link, but is relayed to the destination, such as with a LAN. In contrast with CCITT's X.25-based packet switching, in Frame Relaying Service (FRS) the physical line between nodes contains multiple data links, identified by the address in the data link frame. FRS' major characteristics are out-of-band call control and link layer multiplexing. Unlike the (X.25-based) X.31 packet mode services, FRS integrates more completely with ISDN circuit mode services because of the connection control's out-of-band procedures.

It is expected that FRS, described in I.122, will become an ANSI standard, but additional supporting standardization is needed before the service can be offered in a carrier/vendor-independent fashion.

FRS is based on the ISDN D-channel LAP-D's frame structure, providing statistical multiplexing of different user datastreams within the Data Link Layer (Layer 2). In contrast, X.25 multiplexing occurs at the packet layer (Layer 3). In other words, FRS features the virtual circuit identifier, currently implemented in the Network Layer (Layer 3) of X.25, at the Link Layer (Layer 2) so that fast-packet switching can be accomplished more easily. In the X.25 environment, when a data call is established, the virtual circuit indicator is negotiated and used to route packets through the network for the duration of the call. The indicator is enveloped within the Layer 2 header/trailers, which must be processed before it can be exposed. This processing involves error detection and correction in addition to stripping the header/trailer. In the OSIRM environment layer, n+1 protocol information is enveloped inside layer information.

In LANs, actual packet routing is accomplished directly at Layer 2: the data packets are supplied with a 48-bit destination address, which is readily available and used to physically route the data to the intended destination. Also, no error recovery occurs in a LAN as a packet flows by a station on its way along the bus or ring. In FRS, only Layer 2's lower sublayer—with the core functions frame delimiting, multiplexing, and error detection—is terminated by a network at the user-network interface. Layer 2's upper procedural sublayer, which includes error recovery and flow control, operates between users on an end-to-end basis. In this sense, a user's data transfer protocol is transparent to a network.

Thus, FRS proposes to implement only the "core" functions of LAP-D on a link-by-link basis; the other functions, particularly error recovery, are performed on an end-to-end basis. Indeed, the capabilities provided by the Transport layer protocol (CCITT X.214/X.224, ISO 8072/8073) accommodate this transfer of responsibilities to the network boundaries. Frame relay service is a connection-oriented service, since routing is based on establishing virtual circuit indicators.

Figure 2 shows the partition of the Data Link Layer. For both FR-1 and FR-2, the network terminates only the core aspects of the data link protocol (I.441*). The terminals in FR-2 terminate the full data link protocol, whereas terminals in FR-1 terminate the "core" aspects of the data link protocol. What they terminate above core aspects is a user's option. See table associated with Figure 2.

Frame Relay uses routing based on the address carried in the frame; the frame itself is defined by CCITT I. 441. Figure 3 depicts the lower core LAP-D frame.[9] The "remainder" of the data link layer functions above the core functions must be defined into a peer-to-peer protocol. This protocol is indicated as I.441* in Figure 2.

Work is currently under way in the standards bodies (CCITT SG XI and ECSA T1S1). The addendum to T1.606 defines congestion management strategies; it covers both network and end-user mechanisms and responsibilities to avoid or recover from periods of congestion.[10]

Frame Relaying 1 Service Description

FR-1 data units are frames as defined in Recommendation I.441. The basic bearer service provided is the unacknowledged transfer of frames from S/T to S/T reference point. More specifically, in the U-plane: it preserves their order as given at

Figure 2.
User/Network Interface Protocol Architecture

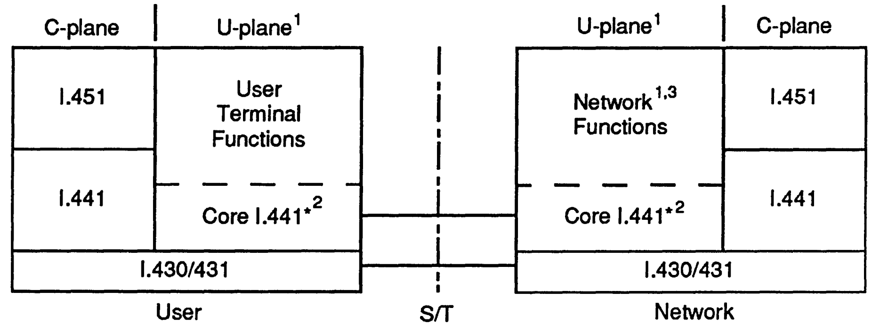

[1] The U-plane functions applicable to each bearer service are given in Table below.
[2] The core functions of Recommendation I.441 are described in text.
[3] The U-plane functions provided by the network at the S/T reference point are determined by the network after negotiation with the user, based on the requested bearer service and associated parameters. These functions are user-selectable for each call. A network may choose not to implement the full set of options. These functions may not be available one by one. So far only three groupings have been identified:
a) the null set,
b) the upper part of Rec. I.441, and
c) the upper part of Rec. I.441 and the data transfer of X.25 PLP.

U-plane Functions Applicable to Each Bearer Service

Bearer Service	User Terminal (Note 1')	Network
Frame Relaying 1	I.441* Core (Note 2')	I.441* Core
Frame Relaying 2	I.441*	I.441* Core
Frame Switching	I.441*	I.441*
X.25-based Additional Packet Mode	I.441* X.25 DTP	I.441* X.25 DTP

Note 1' - Additional user-selectable functions may be implemented.
Note 2' - I.441* is I.441 with appropriate extensions. The use of the extensions may depend on each bearer service and is for further study according to the Blue Book.

one S/T reference point if and when they are delivered at the other end (since the network does not terminate the upper part of I.441, sequence numbers are not kept by the network; networks should be implemented in such a way that, in principle, frame order is preserved); it detects transmission, format, and operational errors; frames are transported transparently (in the network), only the address and FCS fields may be modified (some bits being defined in the address field for congestion control can also be modified); and it does not acknowledge frames (within the network).

Frame Relaying 2 Service Description

The frame relaying data units are defined in Recommendation I.441. The basic bearer service provided is an unacknowledged transfer of frames from S/T to S/T reference point. More specifically, in the U-plane it preserves their order as given at one S/T reference point if and when they are delivered at the other end (since the network does not terminate the upper part of I.441, sequence numbers are not kept by the network. Networks should be implemented in a way that, in principle, frame order is preserved); it detects transmission, format, and operational errors; frames are transported transparently in the network, only the address and

Figure 3.
Address Field Format

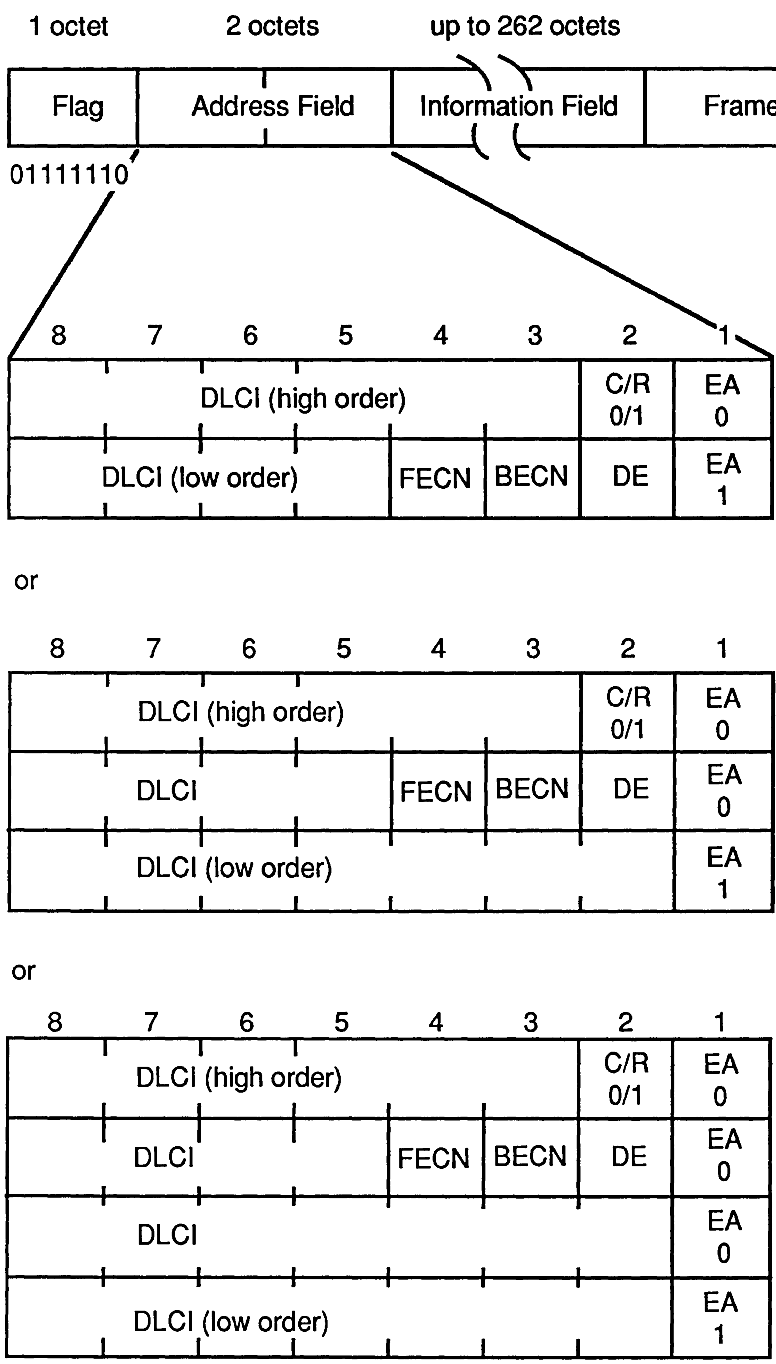

C/R	Bit intended to support a command/response indication. The use of this field is application specific
EA	Address field extension bit
DE	Discard eligibility indicator
BECN	Backward explicit congestion notification
FECN	Forward explicit congestion notification
DLCI	Data link connection identifier

FCS fields may be modified; it does not acknowledge frames (within the network); and normally, the only frames received by a user are those sent by the distant user.

Frame Relay Network Interworking
In the future, FRS may be provided by different carriers. There will then be a demand to interconnect these FR networks with existing Packet Switched Public Data Networks (PSPDNs). Any

interconnection should allow users on different networks to communicate with each other in a uniform way, as though they were all on a single network.

To achieve these goals, standardized procedures must be established for the interworking between 1) FR networks (FR to FR) and 2) FR networks and X.25-based networks (FR to X.25). An FR connection may pass through multiple networks, including both private and public networks. For public networks, both local exchange and interexchange carriers may be involved. In a multivendor environment, the network of a given carrier or provider may consist of equipment from different vendors and/or different types of equipment from the the same vendor.

Two approaches for interworking can be used: network-to-network (NW-NW) interworking and node-to-node (ND-ND) interworking. In the NW-NW method, the signaling and transport procedures at an internetwork interface are standardized. Within a given network, internodal signaling is an internal matter.

Traditionally, telephone exchanges in a public switched telephone network are interconnected using the ND-ND method. This is the philosophy behind Signaling System No. 7 (SS7). In contrast, the NW-NW approach has been used for the interworking of data communication networks.

ANSI Frame Relay Standardization Efforts

The American National Standards Institute (ANSI) has recently issued three draft documents in reference to frame relay in the U.S.:

1. T1.606: Frame Relaying Bearer Service—Architectural Framework and service description (T1S1/88-225)
2. Addendum to T1.606: Frame Relaying Bearer Service—Architectural Framework and service description (T1S1/90-175)
3. T1.6ca: Core Aspects of Frame Protocol for use with Frame Relay Bearer Service (T1S1/90-214)

The data transfer phase of the frame relay bearer service is defined in T1.606. This document specifies a framework for frame relaying service in user-network interface requirements and

Figure 4.
Frame Relay Frame Format

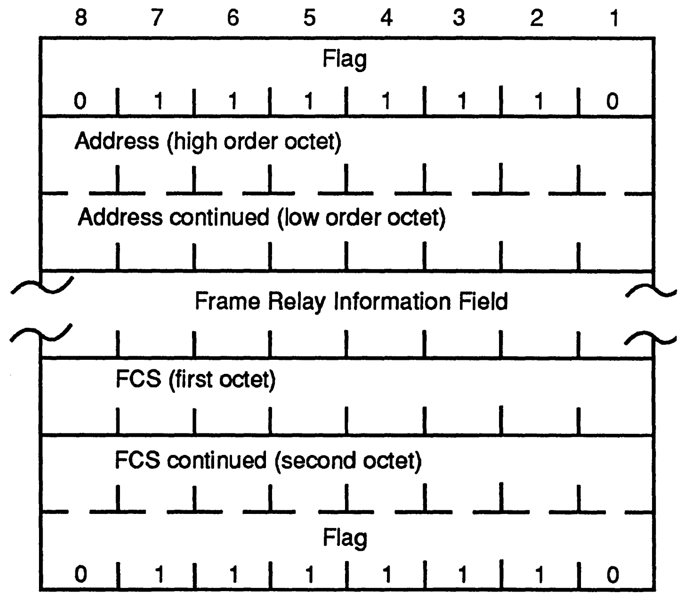

Note: The default address field length is 2 octets. It may be extended to either 3 or 4 octets by subscription.

internetworking requirements. This document includes both interworking with X.25 and interworking between frame relaying service.

The protocol needed to support frame relay is defined in T1.6ca; the protocol operates at the lowest sublayer of the data link layer of the OSI reference model and is based on a subset of T1.602 (LAP-D) called "core aspects."

The frame relay data transfer protocol defined in T1.6ca is intended to support multiple, simultaneous, end-user protocols within a single physical channel. This protocol provides transparent transfer of user data and does not restrict the contents, format, or coding of the information or interpret the structure. This standard is applicable to Frame Relay Bearer Service (FRBS). It is intended for use on any bearer channels and operates on the D-channel concurrently with ANSI Standard T1.602.

Frame Relay Frame Structure

The frame relay frame format is shown in Figure 4. The field shown in the figure is described below.

Flag Sequence

All frames start and end with the flag sequence, consisting of one 0 bit followed by six contiguous 1 bits and one 0 bit. The flag preceding the address field is defined as the opening flag. The flag following the frame check sequence (FCS) field is defined as the closing flag. The closing flag can also serve as the opening and must be capable of receiving one or more consecutive flags.

Address Field
The address field consists of at least two octets, as illustrated in Figure 3, but may optionally be extended up to 4 octets.

Control Field
There is no control field for Frame Relay core services.

Frame Relay Information Field
The Frame Relay information field follows the address field and precedes the frame check sequence. The contents of the user data field consists of an integral number of octets (no partial octets). Networks can support a default maximum information field size of 262 octets. All other maximum values are negotiated between users and networks and between networks. T1S1 strongly recommends that networks support a negotiated maximum value of at least 1,600 octets for applications such as LAN interconnect, to prevent the need for segmentation and reassembly of user equipment.

Transparency
A transmitting data link layer entity must examine the frame content between the opening and closing flag sequences (address, Frame Relay information, and FCS fields), and must insert a 0 bit after all sequences of five contiguous 1 bits (including the last five bits of the FCS) to ensure that a flag or an abort sequence is not simulated within the frame. A receiving data link layer entity must examine the frame contents between the opening and closing flag, which directly follows five contiguous 1 bits.

Frame Checking Sequence (FCS) Field
The FCS field is a 16-bit sequence.

Order of Bit Transmission
The octets are transmitted in ascending numerical order; inside an octet bit 1 is the first bit to be transmitted.

Invalid Frames
An invalid frame is one of the following:

- Is not properly bounded by two flags (e.g., a frame abort), or
- Has fewer than five octets between flags (Note: if there is no information field, the frame has four octets and the frame will be considered invalid), or
- Does not consist of an integral number of octets prior to zero bit insertion or following zero bit extraction, or
- Contains a frame check sequence error, or
- Contains a single octet address field, or
- Contains a Data Link Connection Identifier (DLCI) that is not supported by the receiver.

If the frame received by the network is too long, the network can do one of the following:

- Discard the frame,
- Send part of the frame toward the destination user, then abort the frame, or
- Send the frame toward the destination user with valid FCS.

Selecting one or more of these behaviors is an option for designers of frame relay network equipment, and is not subject to further standardization by T1S1. Users cannot make any assumption as to which of these actions the network will take. In addition, the network may optionally clear the frame relay call if the number or frequency of too-long frames exceeds a network-specified threshold. Invalid frames are discarded without notifying the sender. No action is taken as a result of those frames.

Frame Abort
Receipt of seven or more contiguous 1 bits is interpreted as an abort and the data link layer ignores the frame currently being received.

Address Field Format
The format of the default address field is shown in Figure 3. This field includes the address field extension bits, a bit reserved for use by end-user equipment intended to support a Command/Response indication bit, Forward and Backward explicit congestion indicator bits, a Discard eligibility indicator, and a Data Link Identification (DLCI) field. The minimum and default length of

the address field is two octets, and it can be extended to three or four octets. To support a larger DLCI address range, the three-octet or four-octet address fields can be supported at the user-network interface or the network-network interface based on bilateral agreement.

Interim Specification

The ANSI standards are expected to be approved in mid-1991, and the CCITT standard would be finalized six months later.[11] As a consequence, at the end of 1990 a number of vendors backed an interim joint frame relay specification to ensure some degree of interoperability of new products delivered before ECSA (T1S1) and CCITT standards are finalized. The joint specification is not as complete as the ECSA draft document, but provides a basic set of agreements to begin developing products.

The need to offer interoperable frame relay products is critical, and vendors realize that users may not be willing to deploy technologies that lock them in with systems that could become obsolete in a year or two. cisco Systems, Inc.; Digital Equipment; Northern Telecom, Inc.; and StrataCom, Inc. jointly developed a frame relay specification on which product development can be based until international standards become available.[12,13] The interim specification, announced on September 4, 1990, is based on the emerging ECSA standard but has additional management features and broadcasting.[13] For example, it includes a 16-bit header containing data for routing and congestion control. It will also support automatic reconfiguration of devices with a frame relay interface, and has the capability to detect faults in devices using frame relay interfaces. The interim spec also covers the physical layer.[11]

Extensions to the basic T1S1 draft documents include the following:[14]

1. Support for a global addressing convention, to identify a specific end device
2. Multicast capability, to send frames to all devices which belong to a "multicast group"
3. Asynchronous status updates
4. Flow control, for preventing congestion collapse in a frame relay network
5. Extensions to the Local Management Interface (LMI)

StrataCom had been working on its own specification consistent with the ECSA and CCITT standards under development, but it agreed to share its data with the other three vendors when it decided to enter the joint effort.[11]

Approximately 20 other vendors have announced commitments to support the interim specification, including Newbridge Networks, Telematics International, 3Com Corp., Timeplex, Wellfleet, and Vitalink Communications Corp. Several vendors have started announcing product delivery dates; bridge and router vendors are among the first to announce frame relay support. By forming strategic alliances, vendors can get frame relay products to market more quickly.

Agreement on frame relay implementation specifications will facilitate the emergence of equipment from various vendors, allowing flexibility in user choices.[15] The equipment should also be usable with ISDN. Vendors are trying to avoid the implementation problems experienced in the early 1980s when X.25 packet-switching products entered the market—incompatible X.25 implementations still abound to this day. For example, with the interim specification, users of Digital's routers and cisco's line of multiprotocol router/bridges can connect their equipment to a StrataCom IPX-based private network, or to a public network with a) a Northern Telecom SuperNode at both COs serving the two end points, and b) a tariffed frame relay service.

For some vendors, such as those offering internetworking products, adding frame relay support may require a simple software upgrade since bridges and routers are already based on packet architectures; it will also require the addition of appropriately configured termination cards on the data comm side of the devices. Note, however, that MAC-layer bridges employ connectionless protocols; hence, the connection to frame relay is not trivial. Routers, on the other hand, may already support a connection-oriented service, and so the move to frame relay may be simpler. Routers are typically more expensive, however, than bridges.

Vendors of T1 multiplexers based on circuit-switching time-division multiplexing (TDM) architectures need more work to accommodate frame relay because they do not have experience with

Figure 5.
LAN Interconnection Options

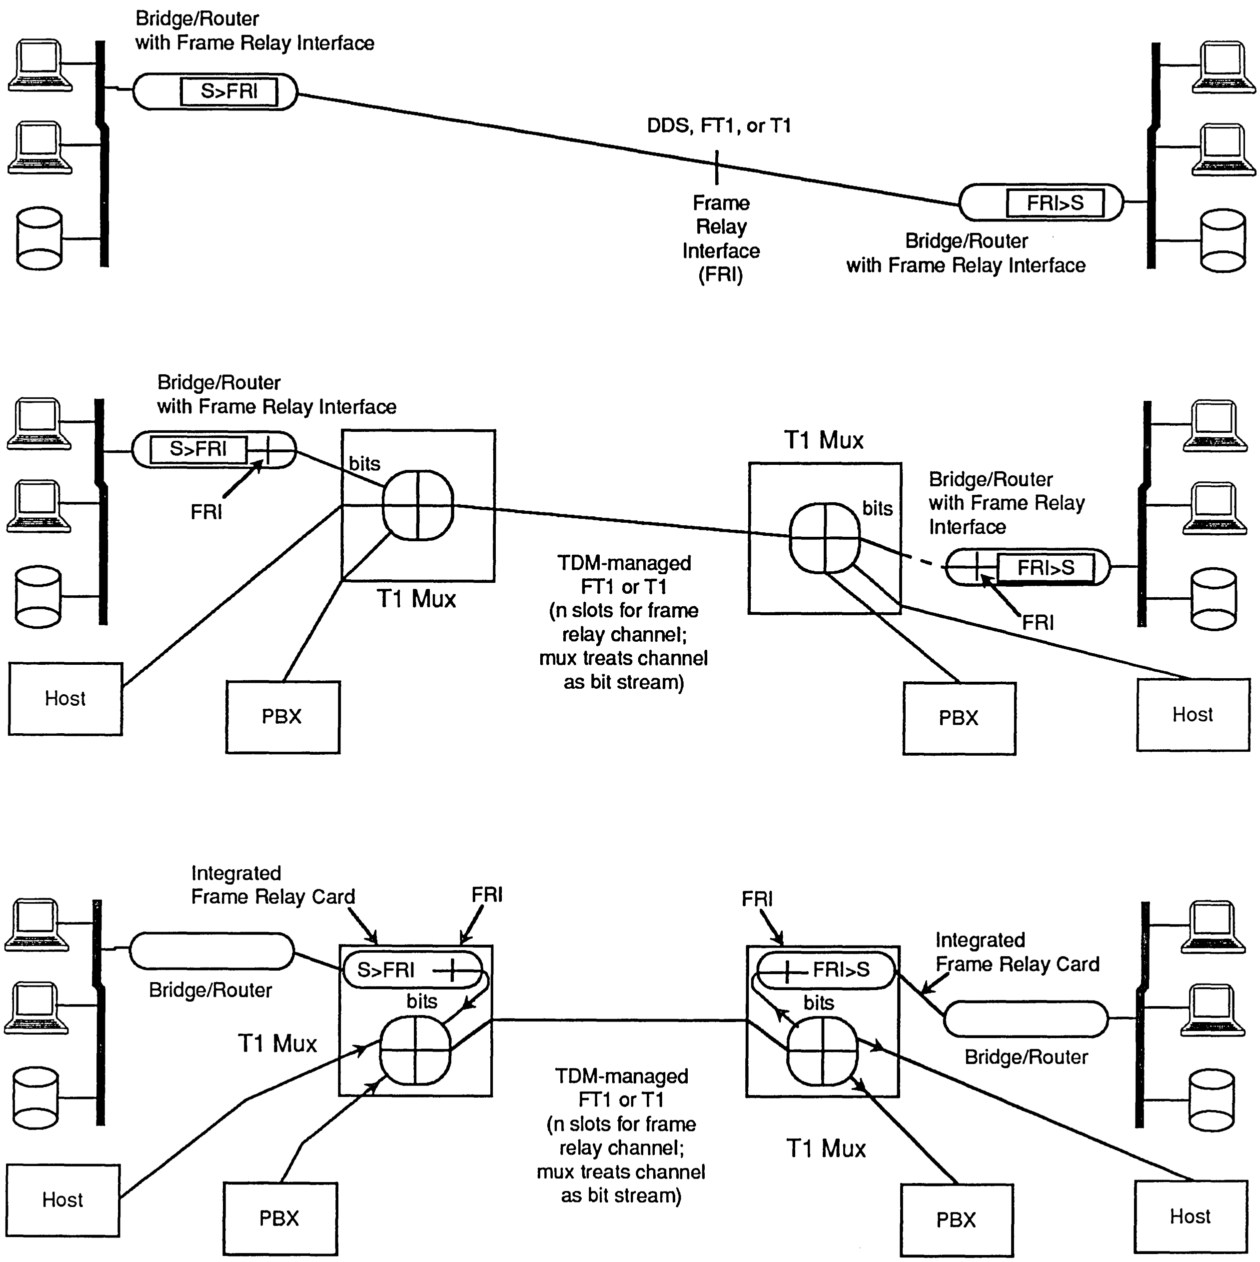

packet-switching technology. These vendors must add a packet engine to support frame relay.

Two T1 vendor approaches offer a short-term solution. The first approach is to develop frame relay modules, or boards, for existing circuit switched multiplexers. The second approach is to use a front-end frame relay developed by another vendor or strategic partner. With interim solutions, the T1 multiplexer may typically allocate only a finite amount of bandwidth for frame relay support, and performance and throughput problems may occur. In the long term, traditional T1 equipment may have to be redesigned to fully exploit the advantages of packet switching in general, and frame relay in particular.

Figure 6.
Additional LAN Interconnection Options

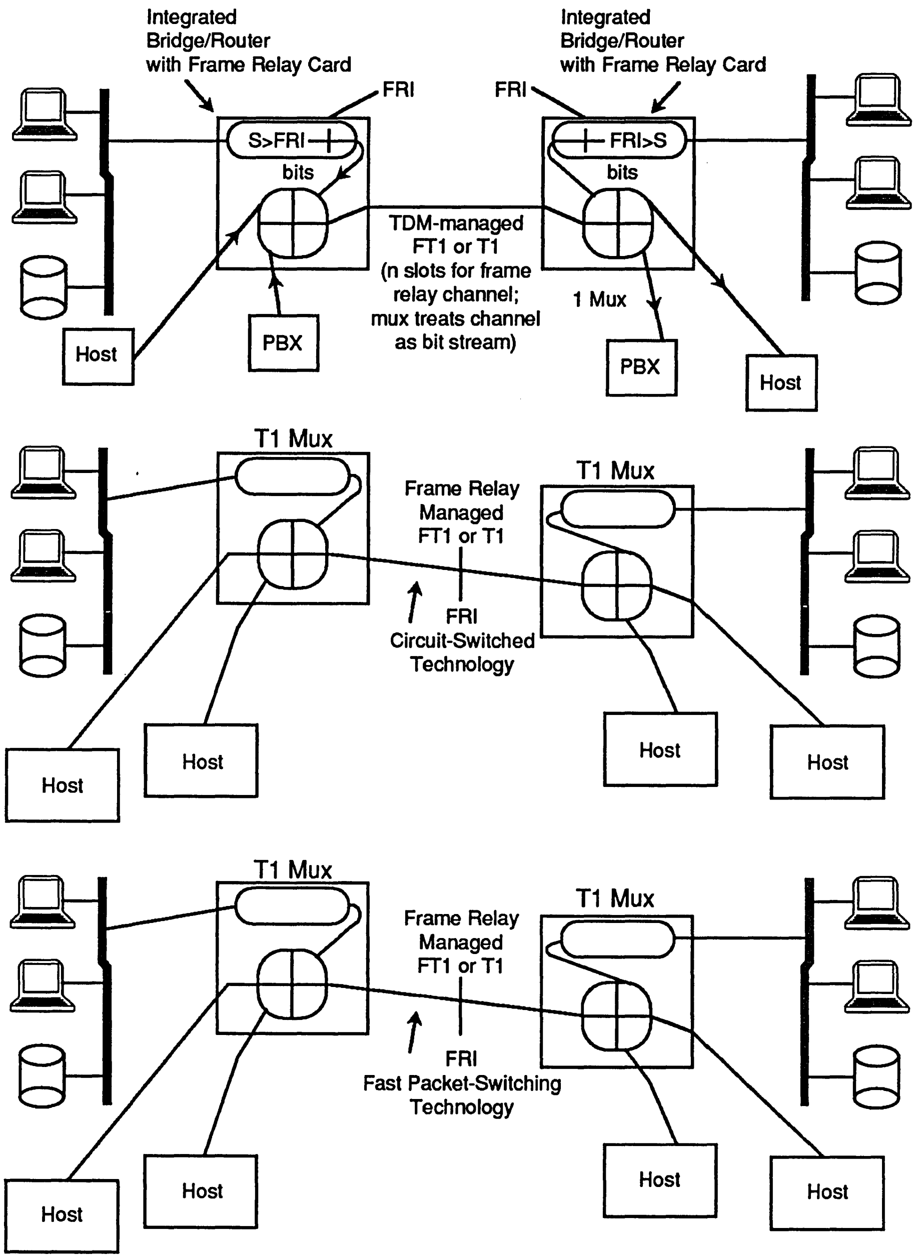

Some vendors are reportedly concerned about the interim specification. The concerns center on possible anticompetitive implications. Some vendors may continue to pursue bilateral agreements on frame relay, raising the possibility of incompatible products, although everybody recognizes the detrimental overall effect of such a posture. Vendors including Advanced Computer Communications, Hughes Network Systems, and Netrix—all of which are actively working on frame relay products—reportedly would welcome the opportunity to help develop the interim specification. As of fall 1990, another vendor group could possibly emerge with its own specification.

Other companies are calling on the interim specification architects to expand their group to bring in users and other vendors. Some vendors are also calling for a means of interoperability testing, an issue yet to be addressed. The standards bodies are pointing out that ANSI is the forum that users and vendors should use to work out a frame relay standard; the T1S1 version is already in draft form and in the process of final balloting.[15]

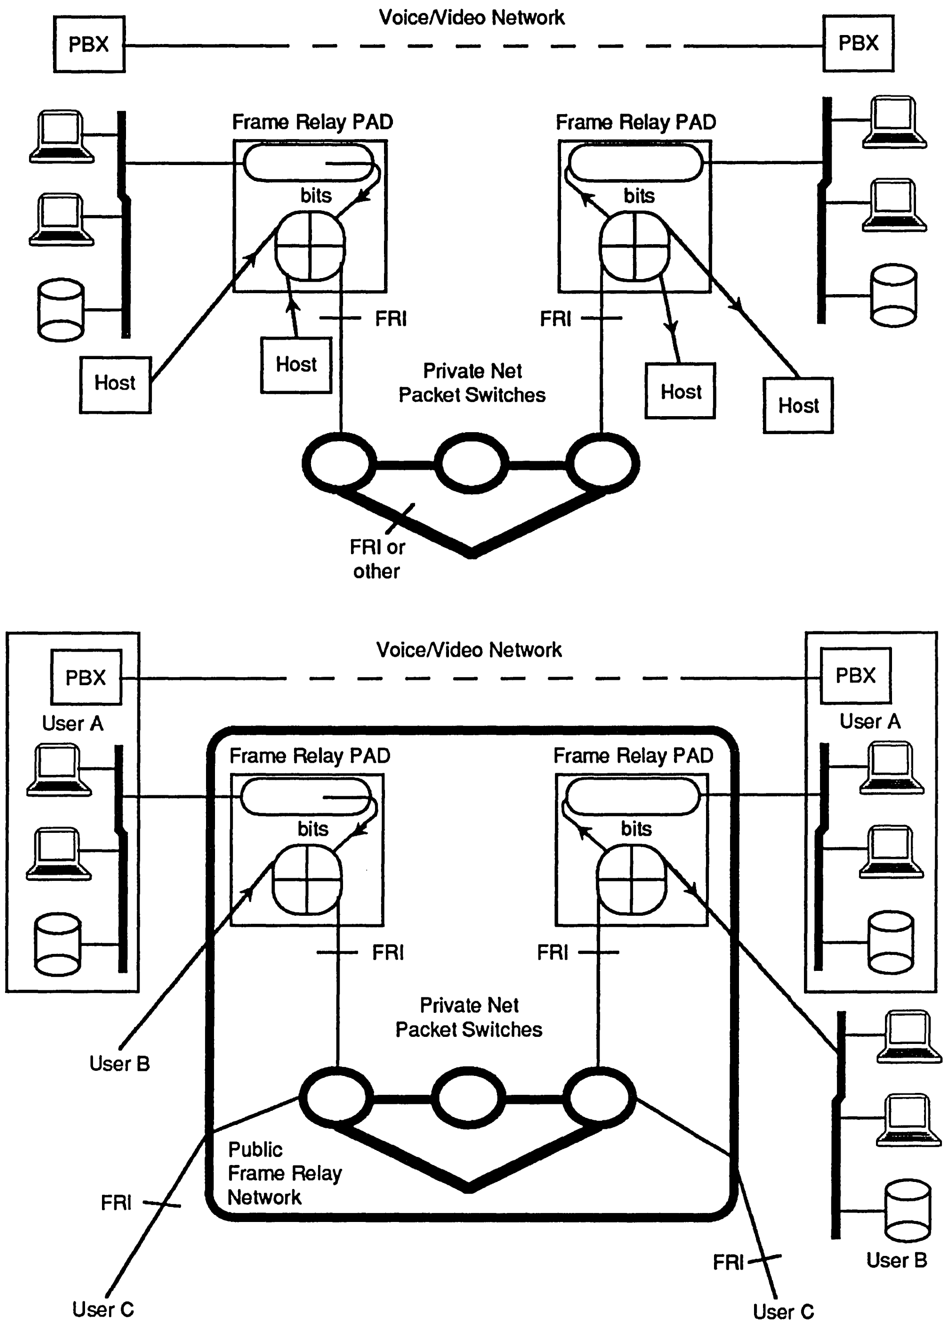

Figure 7.
Frame Relay Interconnection Options

Frame Relay Applications in Private Networking

Users with dedicated LAN links today must examine traffic loads to determine if frame relay and fast-packet will be beneficial. Users with little LAN interconnection traffic may be better off using a 64K bps circuit on a TDM-based T1 multiplexer, while those with higher volumes may want to replace TDM multiplexers with ones supporting fast-packet and frame relay.

Because frame relay-based equipment better utilizes backbone facilities than existing circuit switching T1 multiplexers (and also existing X.25 switches), frame relay and fast-packet switching benefit users who want to connect LANs over integrated backbones supporting a variety of other traffic, assuming that these applications can also share the frame relay bandwidth (to take advantage of resource sharing). But users who simply want to provide high-speed links between remote LANs may be better off using fractional T1, T1, fractional T3, or even T3 links, until SMDS becomes available.

In the short term, most users who need LAN internetworking may find the preferred path is using bridges or routers linked via private lines, according to several analysts.[5] This approach provides excellent response time; the drawback is the cost of the dedicated facilities. Unless frame relay communication service is tariffed, and at the right level, using a frame relay-configured bridge over a dedicated link does not appear to add value, and, in fact, affects response time and involves expenditures for new equipment.

According to some observers, most users must transport a mix of data, voice, and video; hence, they will find it difficult to cost justify building a network solely dedicated to LAN traffic.[2]

Figure 5 shows some examples of LAN interconnection options using frame relay technology. The top portion shows the use of a T1 line dedicated to a new bridge/router system incorporating frame relay. The middle portion demonstrates where a fixed portion of bandwidth from a T1 mux is employed to connect a bridge/router system which incorporates frame relay. The bottom diagram shows a T1 mux with an integrated frame relay card but not a bridge; a fixed portion of bandwidth from a T1 mux is employed. These three scenarios are likely to represent the early usage of FR technology.

Figure 6 shows other examples of possible interconnection options using frame relay. The top part shows a T1 mux with an integrated bridge that uses frame relay; a fixed portion of bandwidth from a T1 mux is employed. The middle part depicts the situation where various streams run into the frame relay-configured mux; the trunk side uses frame relay, but the trunk bandwidth is managed in circuit mode. The bottom part is the same as the previous case, but the trunk side uses frame relay and the trunk bandwidth is managed in fast-packet mode.

Figure 7 depicts a more long-term usage of frame relay. The top part demonstrates a private packet switched network utilizing frame relay network-wide to achieve efficiency; PADs may be required. A separate network for voice and video is required. The bottom part shows a public packet switched network that utilizes and offers frame relay to achieve efficiency; PADs may be required. Multiple users share the network. A separate network for voice and video is required.

References

[1]J. Brown, "Frame Relay, Fast Packet Not Suited for All LAN Users," *Network World*, August 13, 1990, p. 2.
[2]B. Wallace, "US Sprint Unleashes Frame Relay Service," *Network World*, October 1, 1990, p. 1.
[3]*Communications Week*, September 24, 1990, p. 67.
[4]T. Sweeney, "Nynex is Leading the Way," *Communications Week*.
[5]T1.606: Frame Relaying Bearer Service—Architectural Framework and service description (T1S1/88-225).
[6]J. McQuillen, "Frame Relay vs. Cell Relay," *Business Communications Review*, 1990.
[7]J.M. McQuillan, "Broadband Networks—The end of Distance," *Data Communications*, June 1990.
[8]DSS1—Core Aspects of Frame Protocol for use with Frame Relay Bearer Service (T1S1/90-214).
[9]P. Desmond, "Frame Relay Standards Get T1 Mux Makers' Backing," *Network World*, June 11, 1990, p. 2.
[10]Addendum to T1.606: Frame Relaying Bearer Service—Architectural Framework and service description (T1S1/90-175).
[11]*Network World*, September 10, 1990, p. 1.
[12]*Network World*, October 15, 1990, p. 68.
[13]Joint News Release: Northern Telecom, Digital, StrataCom, cisco Systems, September 4, 1990.
[14]Frame relay specification with extensions, Revision 1.0, DEC, Northern Telecom, StrataCom, cisco, September 18, 1990.
[15]P. Desmond, "Firms question goals of frame relay group," *Network World*, October 1, 1990, p. 9. ■