Frame Relay Networks

1 Abstract

Frame relay networks reduce the cost of transmission lines and equipment and improve network performance and response time. Designed for transmission lines with a low error rate, frame relay networks provide minimal internal checking, and consequently, error detection and recovery is implemented in the attached user systems. The Frame Relay Bearer Service was developed specifically as a data service to handle high-volume, bursty traffic by means of high-speed packet transmission, minimal network delay, and efficient use of network bandwidth. The frame protocol supports the data transfer phase of the Service; the frame relay header and the local management interface are sources of congestion avoidance mechanisms. Current implementations include the StrataCom IPX FastPacket digital networking system, which provides the frame relay network, and Digital's DECNIS 500/600 and DEC WANrouter 100/500 software for attaching user equipment.

Today's communications networks are built using high-speed digital trunks that inherently provide high throughput, minimal delay, and a very low error rate. Such transmission networks supply highly reliable service without the overhead of error control functions. Frame relay is a packet-mode transmission network service that exploits these network characteristics by minimizing the amount of error detection and recovery performed inside the network.

This paper explains the nature of the Frame Relay Bearer Service (FRBS) and provides details of the interface defined for attaching user equipment. The implications for higher-layer protocols in the user equipment are also considered.

Following this tutorial, the paper introduces some current implementations. As an example of equipment used to construct a frame relay network, the technology deployed by the StrataCom integrated packet exchange (IPX) FastPacket range of equipment is described. Access to a frame relay network is typically via a router, as is illustrated in the discussion of two Digital products:

- The DECNIS V2.1 software, i.e., network integration server, for either the DECNIS 500 or the DECNIS 600 hardware units (abbreviated as DECNIS 500/600)

- The DEC WANrouter V1.0 software for either the DEMSB or the DEMSA hardware units (subsequently referred to as the WANrouter 100/500)

The paper concludes with a brief discussion of activities related to the further development of frame relay technology.

2 The Frame Relay Bearer Service

The FRBS was developed specifically as a data service to handle high-volume, bursty traffic. The service was designed to provide high-speed packet transmission, minimal network delay, and efficient use of network bandwidth.[1] Local area network (LAN)-to-LAN wide area internetworking is a typical application.

The packet-based frame relay technology uses a combination of features from existing standards for X.25 packet switching and time division multiplexed (TDM) circuit switching.[2] Frame relay provides an X.25-like statistical interface but with lower functionality (in terms of error correction and flow control) and hence higher throughput, because most processing requirements have been removed. At the same time, frame relay has the higher speed and lower delay qualities of TDM circuit switching without the need for dedicated full-time devices and circuits and wasted time slots when no data is being transmitted. The fact that the FRBS need not provide error detection/correction and flow control relies on the existence of intelligent end-user devices, the use of controlling protocol layers (CPLs), and high-speed and reliable communication systems. Access to the FRBS is via a frame relay interface defined between data circuit-terminating equipment (DCE) on the network side and data terminal equipment (DTE) on the user side. A typical frame relay configuration is shown in Figure 1.

In 1990, four vendors-StrataCom, Digital Equipment Corporation, Cisco Systems, and Northern Telecom-collaborated on developing a specification called the Frame Relay Specification with Extensions.[3] This document, edited by StrataCom, introduced a local management interface (LMI) to provide control procedures for permanent virtual circuits (PVCs). The LMI was structured into a basic, mandatory part and a number of optional extensions. It focused on PVCs for frame relay point-to-point connections rather than on switched virtual connections (SVCs), because SVCs are not well suited for LAN interconnection.

Subsequently, standards have emerged in this area that adopt the basic form of the LMI, without the optional extensions, as an annex for PVC control procedures. These standards do differ, however, in some respects. First, the recent standards have specified primary rate access (PRA) for the physical interface rather than Comité Consultatif International de Télégraphique et Téléphonique (CCITT) Recommendation V.35 for wideband electrical signaling, which was adopted in the original joint document.[4] Second, the standards include signaling support for SVCs. The frame relay service is being standardized by both the American National Standards Institute (ANSI) committee, ANSI T1S1, and the CCITT.

3 Frame Protocol

ANSI used the earlier work as a basis for developing the frame protocol to support the data transfer phase of the FRBS.[5] This protocol operates at the lowest sublayer of the data link layer of the International Organization for Standardization/Open Systems Interconnection (ISO/OSI) seven-layer reference model. The protocol is based on a core subset of link access protocol D (LAP-D), which is used in the Integrated Services Digital Network (ISDN). The frame protocol specifies the following characteristics of the frame relay protocol data unit (PDU) or frame:

- Frame delimiting, alignment, and transparency, provided by high-level data link control (HDLC) flags and zero-bit insertion/extraction.

- Framed integrity verification, provided by a frame check sequence (FCS). The FCS is generated using the standard 16-bit CCITT cyclic redundancy check (CRC) polynomial.

- Frame relay addressing, using headers of 2, 3, or 4 octets in length. Figure 2 shows the frame relay header formats. An extended address (E/A) bit is reserved in each octet to indicate whether or not the octet is the last one in the header.

  Most of the header represents the data link connection identifier (DLCI), which identifies the frame's virtual circuit. The header may also contain a DLCI or control indicator (D/C) to indicate whether the remaining six bits are to be interpreted as lower DLCI bits or as control bits. For alignment with LAP-D, the header also contains a bit to discriminate between commands and responses (C/R). This bit is not used for supporting frame relay access.

  The DLCI influences the routing of the frame to the desired destination. The DLCI is also used to multiplex PVCs onto the physical link and enables each endpoint to communicate with multiple destinations by means of a single network access. DLCIs may have either global or local significance in the network. In the global case, the scope of the DLCI extends throughout the network such that a particular DLCI always identifies the same destination, thus making the frame relay network look more like a LAN. In the local case, the scope of the DLCI is limited to the particular interface. When local DLCIs are used, the same DLCI can be reused at another interface to represent a different connection.

- Congestion control and avoidance information. The frame relay header also contains the forward explicit congestion notification (FECN) bit, the backward explicit congestion notification (BECN) bit, and the discard eligibility (DE) indicator, which are discussed in the

Congestion Avoidance section.

4 Permanent Virtual Circuit Control Procedures

Frame relay PVCs provide point-to-point connections between users. Although the PVCs are set up for long periods of time, they can still be considered virtual connections because network resources (i.e., buffers and bandwidth) are not consumed unless data is being transferred.

For interface management purposes, the frame relay interface includes control procedures based on the LMI definition contained in the original multivendor specification. These procedures use messages carried over a separate PVC identified by an in-channel signaling DLCI. The management messages are transferred across the interface using data link unnumbered information frames, as defined in CCITT Recommendation Q.922.[6] The messages use a format similar to that defined in CCITT Recommendation Q.931 for ISDN signaling in support of call control and feature invocation.[7] Each message is formed from a set of standardized information elements defining the message type and associated parameters. The control procedures perform three main functions:

- Link integrity verification initiated by the user device and maintained on a continuous basis. This function allows each entity to be confident that the other is operational and that the physical link is intact.

- When requested by the user, full status network report providing details of all PVCs. The user would normally request such a report at start-up and then periodically.

- Notification by the network of changes in individual PVC status, including the addition of a PVC and a change in PVC state (active /inactive).

The management protocol is defined in Annex D of ANSI T1.617, with equivalent functionality also defined in CCITT Recommendation Q.933, Annex A.[8,9]

5 Effect on Higher-level Protocols

Frame relay provides a multiplexed PVC interface and, with regard to routing software, can be modeled as a set of point-to-point links. However, the characteristics of the frame relay service differ from normal point-to-point links. The major differences are as follows:

- Round-trip delay across a frame relay network is normally longer than the delay across a dedicated point-to-point link.

- PVC throughput can be as high as 2 megabits per second (Mb/s), whereas many existing leased lines operate at lower rates.

- A single frame relay interface can have multiple virtual connections (each one going to a different destination) as compared to the traditional point-to-point link, which supports a single connection.

Given the specific characteristics just described, a frame relay interface may have many more packets in transit than a conventional point-to-point link. Consequently, an acknowledged data link protocol whose procedures include retransmission of data frames is of limited use in this environment. For a large number of virtual connections, the memory required to store the data frames pending acknowledgment would be prohibitive. In addition, if frames are being discarded due to congestion in the frame relay subnetwork, the retransmission policy would increase, rather than recover from, this congestion. Instead, an unacknowledged data link layer should be used.

Using an unacknowledged data link protocol has implications for the routing layer operating over frame relay. In particular, the data link can no longer be considered reliable, and the routing protocol must accommodate this characteristic.

6 Congestion Avoidance

When a frame relay network becomes congested, network devices have no option but to drop frames once their buffers become full. With an unacknowledged data link layer, the user device will not be informed if a data frame is lost. This lack of explicit signaling when operating over frame relay networks places a requirement on the higher protocol layers in the end-system equipment. The OSI transport layer protocol demonstrates how to deal with this type of characteristic. The destination end system's transport implementation detects data loss and requests the source to retransmit the frame. The implementation reduces the source's credit to one, thus closing the source's transmit window and, in effect, reducing traffic through the congested path.

Frame relay networks are prone to congestion. Consider the scenario shown in Figure 3. Note that the committed information rate (CIR) represents minimum guaranteed throughput. In the configuration shown, the network device can support two PVCs: one running at 64 kilobits per second (kb /s) and the other at 128 kb/s. With no back pressure applied across the frame relay interface, in the worse case, the network device will become congested. The router can send frames into the network or a particular PVC at 1 Mb/s that will then be forwarded at a much slower rate. Once the network device's buffers are full, it will discard frames. As a result, routing and bridging control messages may be lost, thus causing the routing topology to become unstable. Since this, in turn, will likely lead to looping packets, a network meltdown could result.

In addition, if data frames are lost, the higher-layer protocols in the end system (e.g., the OSI transport layer) discover this situation and retransmit the lost frames. Repeated transmission of the same data causes the effective end-to-end throughput to drop well below the minimum guaranteed throughput.

The frame relay header has several mechanisms that can be used to apply the appropriate back pressure to prevent congestion.

- The FECN bit is set by the network when a frame experiences congestion as it traverses the network. In OSI and DECnet Phase V environments, this bit can be mapped onto the congestion-experienced bit in the header of the network layer PDU. This PDU, when subsequently delivered to the destination, allows the destination to discover that the path is congested and to notify the source transport to decrease its window and thus place less demand on the network. Standardization work is currently under way to add similar support to the transmission control protocol /internet protocol (TCP/IP).

- The BECN bit is set by the network when a frame traverses a congested virtual circuit in the opposite direction. This indicator is not perfect, because there is no guarantee that traffic will be generated in this direction on the virtual circuit. A source that detects it is transmitting on a congested path is expected to reduce its offered load.

- The DE bit, if set, indicates that during congestion the frame should be the first discarded. The procedures for deciding to set this bit are not clearly defined. This bit could be set by (1) the entry node of the network, e.g., when the input offered load is too high, or (2) the source user equipment, e.g., to discriminate data frames from the more important routing control messages.

Other methods can be used to avoid the consequences of congestion and hence frame loss. The LMI defined in the multivendor frame relay specification contained an optional extension that included a threshold notification bit in the PVC status information element of one of the messages. The threshold notification bit provided a means of allowing a network device to asynchronously inform a user device that a particular PVC connection was congested. The user device could then stop transmitting data on the connection until the network device informed it that the congestion was alleviated.

Since the loss of routing control messages can cause network instability, an alternative approach is to adopt manual configuration. Static network configurations use reachable addresses to provide routing information such that the transmission of routing control traffic is not required. Consequently, the routing behavior is independent of the performance of the network.

In addition, the user device could implement rate-based transmission to ensure that virtual circuits are not congested. However, a means of notifying the user device of the CIR of a virtual circuit was included only as an optional extension in the LMI specification, and use of such a method would destroy one of the major benefits of frame relay, i.e., the capability to allocate bandwidth on demand.

In practice, network devices have limited internal buffering to store frames; this is reflected in the CIR assigned to PVCs. Consequently, data loss occurs if user devices consistently transmit data on a PVC faster than its associated CIR. Adequate procedures and CPLs that cope with congested situations have yet to be developed and standardized. As a result, such situations may lead to unfairness in a multivendor environment where those users who support congestion avoidance will lose bandwidth to those who do not.

7 Products

Below we describe examples of frame relay products: the StrataCom IPX FastPacket equipment, which provides the frame relay network; and Digital's DECNIS 500/600 and WANrouter 100/500, which support the frame relay service by accessing the interface as user equipment.

The StrataCom IPX FastPacket Product Family

The StrataCom IPX FastPacket product family can be used to build networks that support both circuit-mode voice and data as well as frame relay. Within the network, the StrataCom IPX FastPacket nodes communicate using a technique based on cell switching, which involves the transmission of small, fixed-length cells. Additional, high-level functions provide services on top of the basic transmission network. StrataCom uses a hardware-based switching technique resulting in very high-speed switching (100,000 to 1,000,000 cells per second). With such high throughputs and low delays, these networks have been used for carrying voice, video, and data traffic.

The StrataCom IPX FastPacket network is configured by network management to provide the required virtual circuits between users. The StrataCom cell switching mechanism adopts a single-cell format for the transmission of all types of information, with each cell containing addressing information. Routing tables within the network nodes use this addressing information to forward the traffic along the desired virtual circuit. Since in any particular connection the path used for the sequence of cells is always the same, cell ordering is maintained. Intelligent interfaces at the edge of the network provide the functions required for specific services such as voice and data.

Figure 4 illustrates the concept employed by StrataCom of building service-specific functions on top of a common cell switching technology. The figure shows examples of various types of external interfaces.

For the frame relay interface, StrataCom supports the optional features defined to address congestion. The IPX FastPacket node provides the optional explicit congestion indicators defined in the frame header, which are set based on averaging queues that build up in the IPX FastPacket nodes in the network. Support is also provided for the optional threshold notification feature defined as part of the LMI; the actual threshold values, together with buffer configuration, can be configured by the network manager.

Frame Relay Support in Digital's Family of Multiprotocol Routers

Digital has provided frame relay support in its family of multiprotocol routers that employ the OSI intermediate system-to-intermediate system (IS-IS) routing protocol. Frame relay user device functionality is implemented in the DECNIS V2.1 software for either the DECNIS 500 or the DECNIS 600 hardware units, and in the DEC WANrouter V1.0 software for either the DEMSB or the DEMSA hardware units.

Part of the development of the frame relay support involved cooperating with StrataCom to produce a working frame relay specification. In particular, extensions were added to the LMI to provide appropriate congestion control procedures. Digital's software supports the Frame Relay Specification with Extensions, Revision 1.0, written by StrataCom and the relevant ANSI T1S1 standards.[3,1,5,8] The software has been tested and is known to be compatible with the StrataCom IPX FastPacket 16/32 equipment with Frame Relay Interface Card Software.

The DECNIS and WANrouter implementations use the point-to-point protocol (PPP) for the transmission of multiprotocol datagrams over point-to-point links. PPP is defined in Requests for Comment (RFCs) 1331 and 1332, with bridging extensions specified in RFC 1220; support for DECnet Phase IV is defined in RFC 1376 and for OSI in RFC 1377.[10-14] Congestion avoidance procedures include support for both the threshold notification signal in the LMI (when available) and the FECN. The threshold notification signal causes the end system to modify its rate of data transmission. Receipt of a frame with the FECN bit set causes the equivalent bit in the network layer PDU header to be set, which in turn causes the end systems to reduce their offered traffic. The BECN and DE bits are never set or examined.

8 Related Activities

Various committees are involved in activities related to the frame relay technology. These activities include standards work, specifications, and efforts to address technical issues such as interoperability.

Table_1:_Current_Status_of_Frame_Relay_Standardization

Standard	ANSI	Status	CCITT	Status	Remarks
Architecture and SVC Description	T1.606	Standard	I.233	Standard	Replaces I.222
Congestion Management Principles	Addendum to T1.606	Standard	I.370	Standard	
Data Transfer-Core Aspects	T1.618	Standard	Q.922 (Annex A corresponds to T1.618)	Standard	Most important frame relay standard
Access Signaling for SVCs	T1.617	Standard	Q.933	Standard	
Management Procedures for PVCs	Included in T1.617 Annex D	Standard	Included in Q.933 Annex A	Standard	Concepts accepted in CCITT Standards

The overall frame relay network architecture is defined in ANSI T1.606, Frame Relay Bearer Service-Architectural Framework and Service Description. [1] Access is provided by the frame relay interface, which is defined in various ANSI standards for both permanent and switched virtual circuits. ANSI T1.618, DSS1-Core Aspects of Frame Protocol for Use with Frame Relay Bearer Service contains a definition of the protocol for exchanging frames across the interface, as well as annexes concerned with local management (e.g., notification of PVC status).[5] Although all implementations to date have focused on a PVC-based interface, SVC access is defined in ANSI T1.617, DSS1-Signaling Specification for Frame Relay Bearer Service. [8] Each of these T1S1 standards has an equivalent CCITT recommendation, as shown in Table 1.

Other Current Activities

The Internet Engineering Task Force (IETF) is developing specifications for RFCs related to the frame relay technology. A specification called Multiprotocol Interconnect over Frame Relay defines an encapsulation mechanism for supporting multiple protocols over frame relay networks. To allow use of the simple network management protocol (SNMP), an experimental management information base (MIB) for frame relay DTEs is also under development.

To promote the frame relay technology, a Frame Relay Forum has been set up in both North America and Europe. A technical committee has been established to address issues related to the technology in terms of its interoperability and evolution in multivendor environments. This committee actively participates with the standards bodies and develops implementation agreements and interoperability test procedures. Work continues to define a network-to-network control interface, multicasting capabilities, multiple protocol encapsulation, and interworking with other technologies, such as the switched multimegabit data service (SMDS) defined by Bell Communications Research, Inc.[15]

The cell switching adopted by StrataCom within their network is expected to change over time to conform with emerging CCITT recommendations for broadband ISDN.[16] These recommendations cover asynchronous transfer mode (ATM), which defines a standard cell structure and ATM adaptation layers (AALs) for particular higher-level functions.

9 Summary

Frame relay is a simplified form of packet-mode switching that, at least in theory, provides access to high bandwidth on demand, direct connectivity to all other points in the network, and consumption of only the bandwidth actually used. Thus, to the customer, the frame relay technology offers a reduction in the cost of transmission lines and equipment and improved performance and response time.

Routers connected to a frame relay network can consider the multiplexed, PVC interface as a set of point-to-point links. The special characteristics of a frame relay network require that special care be taken in selecting the data link protocols and in handling congestion.

10 Acknowledgments

The authors thank StrataCom, Inc. for providing significant input on cell switching technology and its use in their IPX FastPacket equipment. The authors would also like to acknowledge Cliff Didcock of the Computer Integrated Telephony Development Group who consulted in Digital's initial frame relay implementation.

11 References

1. ANSI T1.606: Frame Relay Bearer Service-Architectural Framework and Service Description (New York: American National Standards Institute, Inc., 1990).

2. CCITT Recommendation X.25: Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for Terminals

Operating in the Packet Mode and Connected to Public Data Networks
by Dedicated Circuit (Geneva: International Telecommunications Union,
1988).

3. Frame Relay Specification with Extensions, Revision 1.0, Cisco Systems, Digital Equipment Corporation, Northern Telecom, Inc., and StrataCom, Inc. (September 1990).

4. CCITT Recommendation V.35: Data Transmission at 48 Kilobits per Second Using 60-108 kHz Group Band Circuits (Geneva: International Telecommunications Union, 1976).

5. ANSI T1.618: DSS1-Core Aspects of Frame Protocol for Use with Frame Relay Bearer Service (New York: American National Standards Institute, Inc., 1990).

6. CCITT Recommendation Q.922: ISDN User-Network Interface Layer 3 Specification for Basic Call Control (Geneva: International Telecommunications Union, 1991).

7. CCITT Recommendation Q.931: ISDN Data Link Layer Specification for Frame Mode Bearer Services (Geneva: International Telecommunications Union, 1991).

8. ANSI T1.617: DSS1-Signaling Specification for Frame Relay Bearer Service (New York: American National Standards Institute, Inc., 1990).

9. CCITT Draft Recommendation Q.933: ISDN Signalling Specification for Frame Mode Bearer Services (Geneva: International Telecommunications Union, 1991).

10. Point-to-Point Protocol for the Transmission of Multi-protocol Datagrams over Point-to-Point Links, Internet Engineering Task Force RFC 1331 (May 1992).

11. The PPP Internet Protocol Control Protocol (IPCP), Internet Engineering Task Force RFC 1332 (May 1992).

12. Point-to-Point Protocol Extensions for Bridging, Internet Engineering Task Force RFC 1220 (April 1991).

13. PPP DECnet Phase IV Control Protocol (DNCP), Internet Engineering Task Force RFC 1376 (November 1992).

14. PPP OSI Network Layer Control Protocol (OSI NLCP), Internet Engineering Task Force RFC 1377 (November 1992).

15. Bellcore TR-TSV-000772, Generic System Requirements in Support of Switched Multi-Megabit Data Service, Bell Communications Research, Inc. (May 1991).

16. CCITT Draft Recommendation I.121: Broadband Aspects of ISDN (Geneva: International Telecommunications Union, 1991).

12 Biographies

Robert J. Roden Robert Roden is a consulting engineer in Networks Engineering. Recently, he has been working on new transmission technologies such as frame relay and switched multimegabit data service. He has also worked on computer integrated telephony and chaired a group developing related standards. Robert joined Digital in 1986 from Racal Milgo, where he was responsible for local area networks and network management platforms. He received a B.Sc. (1971) in physics and a Ph.D. (1974) in materials science from the Imperial College in London.

Deborah Tayler Deborah Tayler, a principal software engineer in Networks Engineering Europe, is currently responsible for the design and implementation of frame relay and point-to-point protocol functionality on multiprotocol routers. She joined Digital in 1982 and has worked on DECtalk, ALL-IN-1, and computer integrated telephony projects. Deborah received a B.Sc. (1981) in economics from University College in London and an M.Sc. (1982) in the theory and applications of computation from Loughborough University of Technology in Loughborough, Leicestershire.

13 Trademarks