

Frame Relay

Background

Frame Relay was originally conceived as a protocol for use over ISDN interfaces. Initial proposals to this effect were submitted to the International Telecommunication Union Telecommunication Standardization Sector (ITU-T) (formerly the Consultative Committee for International Telegraph and Telephone [CCITT]) in 1984. Work on Frame Relay was also undertaken in the American National Standards Institute (ANSI)-accredited T1S1 standards committee in the United States.

There was a major development in Frame Relay's history in 1990 when Cisco Systems, StrataCom, Northern Telecom, and Digital Equipment Corporation formed a consortium to focus Frame Relay technology development and accelerate the introduction of interoperable Frame Relay products. This consortium developed a specification conforming to the basic Frame Relay protocol being discussed in T1S1 and ITU-T, but extended it with features that provide additional capabilities for complex internetworking environments. These Frame Relay extensions are referred to collectively as the *local management interface* (LMI).

Technology Basics

Frame Relay provides a packet-switching data communications capability that is used across the interface between user devices (for example, routers, bridges, host machines) and network equipment (for example, switching nodes). User devices are often referred to as *data terminal equipment* (DTE), while network equipment that interfaces to DTE is often referred to as *data circuit-terminating equipment* (DCE). The network providing the Frame Relay interface can be either a carrier-provided public network or a network of privately owned equipment serving a single enterprise.

As an interface to a network, Frame Relay is the same type of protocol as X.25 (see Chapter 12, "X.25"). However, Frame Relay differs significantly from X.25 in its functionality and format. In particular, Frame Relay is a more streamlined protocol, facilitating higher performance and greater efficiency.

As an interface between user and network equipment, Frame Relay provides a means for statistically multiplexing many logical data conversations (referred to as *virtual circuits*) over a single physical transmission link. This contrasts with systems that use only *time-division-multiplexing* (TDM) techniques for supporting multiple data streams. Frame Relay's statistical multiplexing provides more flexible and efficient use of available bandwidth. It can be used without TDM techniques or on top of channels provided by TDM systems.

Another important characteristic of Frame Relay is that it exploits the recent advances in wide-area network (WAN) transmission technology. Earlier WAN protocols such as X.25 were developed when analog transmission systems and copper media were predominant. These links are much less reliable than the fiber media/digital transmission links available today. Over links such as these, link-layer

protocols can forego time-consuming error correction algorithms, leaving these to be performed at higher protocol layers. Greater performance and efficiency is therefore possible without sacrificing data integrity. Frame Relay is designed with this approach in mind. It includes a *cyclic redundancy check* (CRC) algorithm for detecting corrupted bits (so the data can be discarded), but it does not include any protocol mechanisms for correcting bad data (for example, by retransmitting it at this level of protocol).

Another difference between Frame Relay and X.25 is the absence of explicit, per-virtual-circuit flow control in Frame Relay. Now that many upper-layer protocols are effectively executing their own flow control algorithms, the need for this functionality at the link layer has diminished. Frame Relay, therefore, does not include explicit flow control procedures that duplicate those in higher layers. Instead, very simple congestion notification mechanisms are provided to allow a network to inform a user device that the network resources are close to a congested state. This notification can alert higher-layer protocols that flow control may be needed.

Current Frame Relay standards address *permanent virtual circuits* (PVCs) that are administratively configured and managed in a Frame Relay network. Another type, *switched virtual circuits* (SVCs), has also been proposed. The *Integrated Services Digital Network* (ISDN) signaling protocol is proposed as the means by which DTE and DCE will communicate to establish, terminate, and manage SVCs dynamically. For more information on ISDN, see Chapter 10, “Integrated Services Digital Network.” Both T1S1 and ITU-T have work in progress to include SVCs in Frame Relay standards.

LMI Extensions


In addition to the basic Frame Relay protocol functions for transferring data, the consortium Frame Relay specification includes LMI extensions that make supporting large, complex internetworks easier. Some LMI extensions are referred to as “common” and are expected to be implemented by everyone who adopts the specification. Other LMI functions are referred to as “optional.” A summary of the LMI extensions follows:

- *Virtual circuit status messages* (common)—Provide communication and synchronization between the network and the user device, periodically reporting the existence of new PVCs and the deletion of already existing PVCs, and generally providing information about PVC integrity. Virtual circuit status messages prevent the sending of data into *black holes*, that is, over PVCs that no longer exist.
- *Multicasting* (optional)—Allows a sender to transmit a single frame but have it delivered by the network to multiple recipients. Thus, multicasting supports the efficient conveyance of routing protocol messages and address resolution procedures that typically must be sent to many destinations simultaneously.
- *Global addressing* (optional)—Gives connection identifiers global rather than local significance, allowing them to be used to identify a specific interface to the Frame Relay network. Global addressing makes the Frame Relay network resemble a local-area network (LAN) in terms of addressing; address resolution protocols therefore perform over Frame Relay exactly as they do over a LAN.
- *Simple flow control* (optional)—Provides for an XON/XOFF flow control mechanism that applies to the entire Frame Relay interface. It is intended for those devices whose higher layers cannot use the congestion notification bits and that need some level of flow control.

Frame Format


The Frame Relay frame is shown in Figure 13-1. The *flags* fields delimit the beginning and end of the frame. Following the leading flags field are two bytes of *address* information. Ten bits of these two bytes make up the actual circuit ID (called the *DLCI*, for *data link connection identifier*).

Figure 13-1 Frame Relay Frame



The 10-bit DLCI value is the heart of the Frame Relay header. It identifies the logical connection that is multiplexed into the physical channel. In the basic (that is, not extended by the LMI) mode of addressing, DLCIs have local significance; that is, the end devices at two different ends of a connection may use a different DLCI to refer to that same connection. Figure 13-2 provides an example of the use of DLCIs in nonextended Frame Relay addressing.

Figure 13-2 Frame Relay Addressing



In Figure 13-2, assume two PVCs, one between Atlanta and Los Angeles, and one between San Jose and Pittsburgh. Los Angeles uses DLCI 12 to refer to its PVC with Atlanta, while Atlanta refers to the same PVC as DLCI 82. Similarly, San Jose uses DLCI 12 to refer to its PVC with Pittsburgh. The network uses internal proprietary mechanisms to keep the two locally significant PVC identifiers distinct.

At the end of each DLCI byte is an extended address (EA) bit. If this bit is one, the current byte is the last DLCI byte. All implementations currently use a 2-byte DLCI, but the presence of the EA bits means that longer DLCIs may be agreed upon and used in the future.

The bit marked “C/R” following the most significant DLCI byte is currently not used.


Finally, 3 bits in the 2-byte DLCI provide congestion control. The forward explicit congestion notification (FECN) bit is set by the Frame Relay network in a frame to tell the DTE receiving that frame that congestion was experienced in the path from source to destination. The backward explicit congestion notification (BECN) bit is set by the Frame Relay network in frames traveling in the opposite direction from frames encountering a congested path. The notion behind both of these bits is that the FECN or BECN indication can be promoted to a higher-level protocol that can take flow control action as appropriate. (FECN bits are useful to higher-layer protocols that use receiver-controlled flow control, while BECN bits are significant to those that depend on “emitter-controlled” flow control.)

The discard eligibility (DE) bit is set by the DTE to tell the Frame Relay network that a frame has lower importance than other frames and should be discarded before other frames if the network becomes short of resources. Thus, it represents a very simple priority mechanism. This bit is usually set only when the network is congested.

LMI Message Format

The previous section described the basic Frame Relay protocol format for carrying user data frames. The consortium Frame Relay specification also includes the LMI procedures. LMI messages are sent in frames distinguished by an LMI-specific DLCI (defined in the consortium specification as DLCI = 1023). The LMI message format is shown in Figure 13-3.

Figure 13-3 LMI Message Format



In LMI messages, the basic protocol header is the same as in normal data frames. The actual LMI message begins with four mandatory bytes, followed by a variable number of *information elements* (IEs). The format and encoding of LMI messages is based on the ANSI T1S1 standard.

The first of the mandatory bytes (*unnumbered information indicator*) has the same format as the LAPB *unnumbered information* (UI) frame indicator with the poll/final bit set to zero. For more information about LAPB, see the section “Layer 2” in Chapter 12, “X.25.” The next byte is referred to as the *protocol discriminator*, which is set to a value that indicates “LMI.” The third mandatory byte (*call reference*) is always filled with zeros.

The final mandatory byte is the *message type* field. Two message types have been defined. *Status-enquiry* messages allow the user device to inquire about network status. *Status* messages respond to status-enquiry messages. *Keepalives* (messages sent through a connection to ensure that both sides will continue to regard the connection as active) and PVC status messages are examples of these messages and are the common LMI features that are expected to be a part of every implementation that conforms to the consortium specification.

Together, status and status-enquiry messages help verify the integrity of logical and physical links. This information is critical in a routing environment because routing algorithms make decisions based on link integrity.

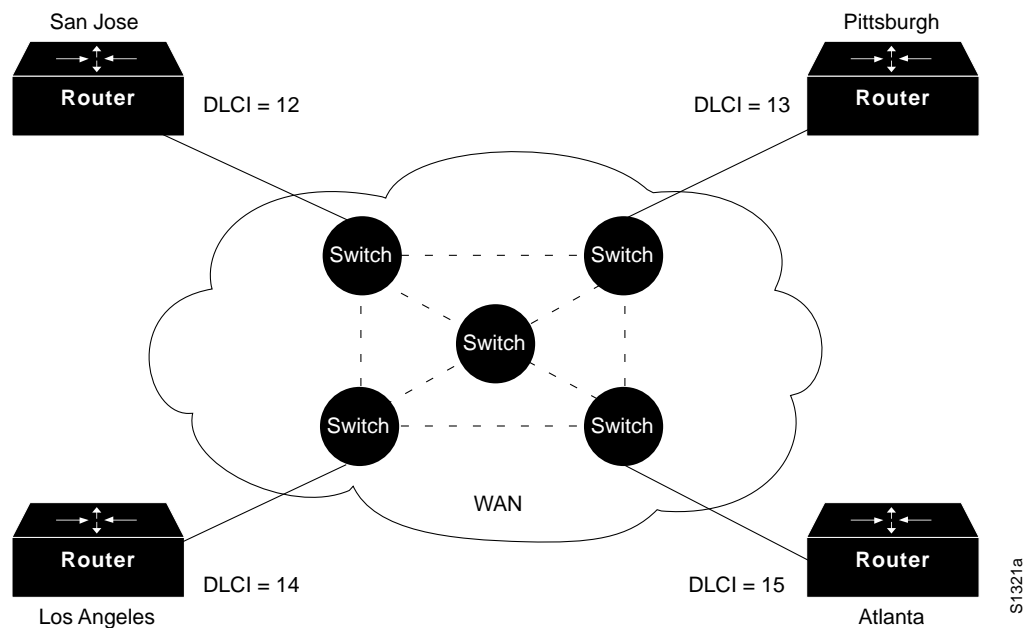
Following the message type field is some number of IEs. Each IE consists of a single-byte *IE identifier*, an *IE length* field, and one or more bytes containing actual data.

Global Addressing

In addition to the common LMI features, there are several optional LMI extensions that are extremely useful in an internetworking environment. The first important optional LMI extension is *global addressing*. As noted earlier, the basic (nonextended) Frame Relay specification only supports values of the DLCI field that identify PVCs with local significance. In this case, there are no addresses that identify network interfaces, or nodes attached to these interfaces. Because these addresses do not exist, they cannot be discovered by traditional address resolution and discovery techniques. This means that with normal Frame Relay addressing, static maps must be created to tell routers which DLCIs to use to find a remote device and its associated internetwork address.

The global addressing extension permits node identifiers. With this extension, the values inserted in the DLCI field of a frame are globally significant addresses of individual end-user devices (for example, routers). This is implemented as shown in Figure 13-4.

Figure 13-4 Global Addressing Exchange



In Figure 13-4, note that each interface has its own identifier. Suppose that Pittsburgh must send a frame to San Jose. The identifier for San Jose is 12, so Pittsburgh places the value 12 in the DLCI field and sends the frame into the Frame Relay network. At the exit point, the DLCI field contents are changed by the network to 13 to reflect the source node of the frame. Each router interface has a distinct value as its node identifier, so individual devices can be distinguished. This permits adaptive routing in complex environments.

Global addressing provides significant benefits in a large, complex internetwork. The Frame Relay network now appears to the routers on its periphery like any LAN. No changes to higher-layer protocols are needed to take full advantage of their capabilities.

Multicasting

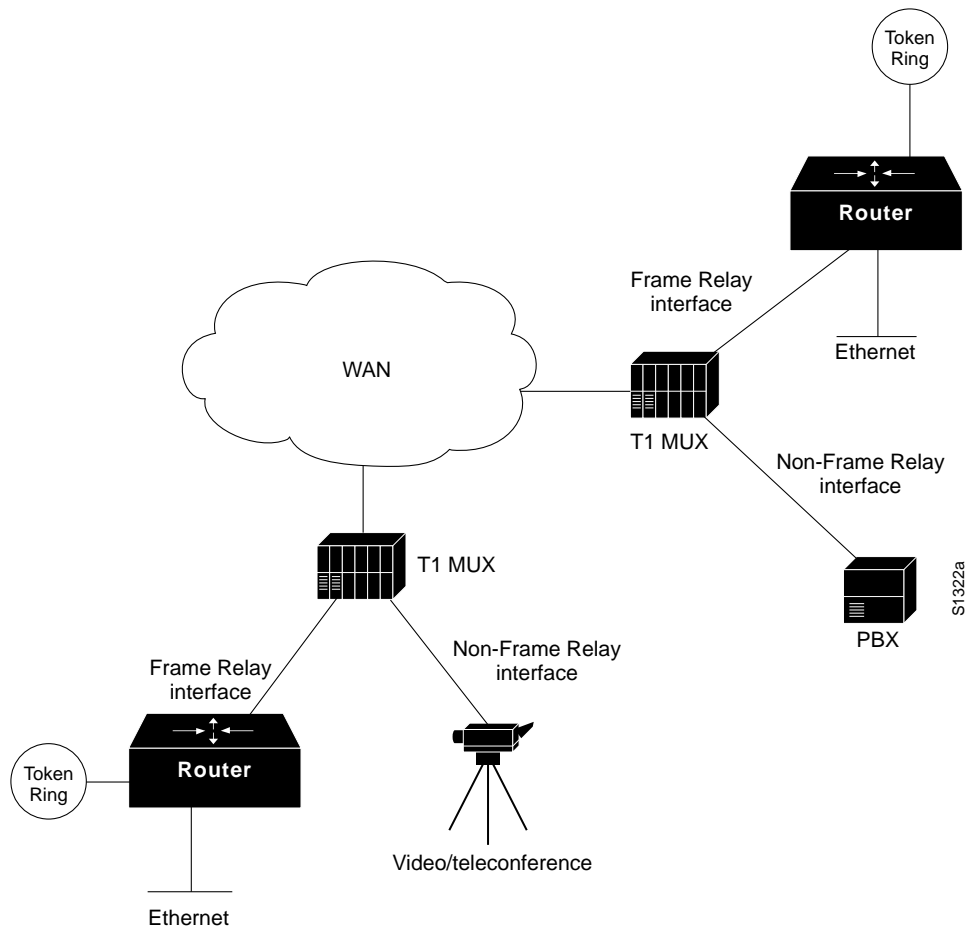
Multicasting is another valuable optional LMI feature. Multicast groups are designated by a series of four reserved DLCI values (1,019 to 1,022). Frames sent by a device using one of these reserved DLCIs are replicated by the network and sent to all exit points in the designated set. The multicasting extension also defines LMI messages that notify user devices of the addition, deletion, and presence of multicast groups.

In networks that take advantage of dynamic routing, routing information must be exchanged among many routers. Routing messages can be sent efficiently by using frames with a multicast DLCI. This allows messages to be sent to specific groups of routers.

Network Implementation

Frame Relay can be used as an interface to either a publicly available carrier-provided service or to a network of privately owned equipment. A typical means of private network implementation is to equip traditional T1 multiplexers with Frame Relay interfaces for data devices, as well as non-Frame Relay interfaces for other applications such as voice and video-teleconferencing. Figure 13-5 shows this configuration.

Figure 13-5 Hybrid Frame Relay Network



A public Frame Relay service is deployed by putting Frame Relay switching equipment in the central offices of a telecommunications carrier. In this case, users can realize economic benefits from traffic-sensitive charging rates, and are relieved from the work necessary to administer and maintain the network equipment and service.

In either type of network, the lines that connect user devices to the network equipment can operate at a speed selected from a broad range of data rates. Speeds between 56 kbps and 2 Mbps are typical, although Frame Relay can support lower and higher speeds. Implementations capable of operating over 45-Mbps (DS3) links are expected to be available soon.

Whether in a public or private network, the support of Frame Relay interfaces to user devices does not necessarily dictate that the Frame Relay protocol is used between the network devices. No standards for interconnecting equipment inside a Frame Relay network currently exist. Thus, traditional circuit-switching, packet-switching, or a hybrid approach combining these technologies can be used.

