

11

Frame Relay

Frame Relay is a protocol standard for LAN internetworking which provides a fast and efficient method of transmitting information from a user device to LAN bridges and routers.

The Frame Relay protocol uses a frame structure similar to that of LAPD, except that the frame header is replaced by a 2-byte Frame Relay header field. The Frame Relay header contains the user-specified DLCI field, which is the destination address of the frame. It also contains congestion and status signals which the network sends to the user.

Virtual Circuits

The Frame Relay frame is transmitted to its destination by way of virtual circuits (logical paths from an originating point in the network) to a destination point. Virtual circuits may be permanent (PVCs) or switched (SVCs). PVCs are set up administratively by the network manager for a dedicated point-to-point connection; SVCs are set up on a call-by-call basis.

Advantages of Frame Relay

Frame Relay offers an attractive alternative to both dedicated lines and X.25 networks for connecting LANs to bridges and routers. The success of the Frame Relay protocol is based on the following two underlying factors:


- Because virtual circuits consume bandwidth only when they transport data, many virtual circuits can exist simultaneously across a given transmission line. In addition, each device can use more of the bandwidth as necessary, and thus operate at higher speeds.
- The improved reliability of communication lines and increased error-handling sophistication at end stations allows the Frame Relay protocol to discard erroneous frames and thus eliminate time-consuming error-handling processing.

These two factors make Frame Relay a desirable choice for data transmission; however, they also necessitate testing to determine that the system works properly and that data is not lost.

Frame Relay Structure

Standards for the Frame Relay protocol have been developed by ANSI and CCITT simultaneously. The separate LMI specification has basically been incorporated into the ANSI specification. The following discussion of the protocol structure includes the major points from these specifications.

The Frame Relay frame structure is based on the LAPD protocol. In the Frame Relay structure, the frame header is altered slightly to contain the Data Link Connection Identifier (DLCI) and congestion bits, in place of the normal address and control fields. This new Frame Relay header is 2 bytes in length and has the following format:



Frame Relay header structure

DLCI

10-bit DLCI field represents the address of the frame and corresponds to a PVC.

C/R

Designates whether the frame is a command or response.

EA

Extended Address field signifies up to two additional bytes in the Frame Relay header, thus greatly expanding the number of possible addresses.

FECN

Forward Explicit Congestion Notification (see ECN below).

BECN

Backward Explicit Congestion Notification (see ECN below).

DE

Discard Eligibility (see DE below).

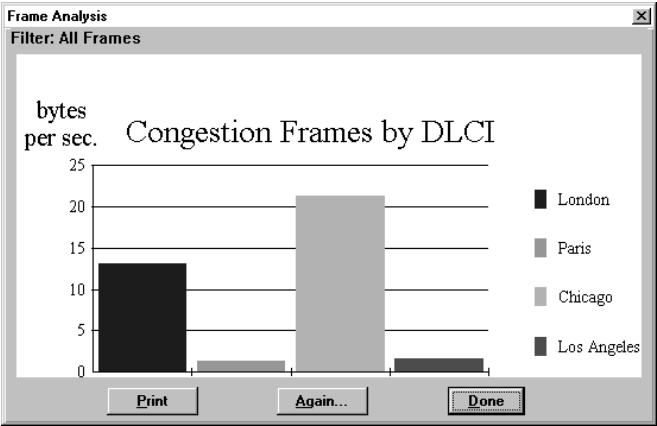
Information

The Information field may include other protocols within it, such as an X.25, IP or SDLC (SNA) packet.

Explicit Congestion Notification (ECN) Bits

When the network becomes congested to the point that it cannot process new data transmissions, it begins to discard frames. These discarded frames are retransmitted, thus causing more congestion. In an effort to prevent this situation, several mechanisms have been developed to notify user devices at the onset of congestion, so that the offered load may be reduced.

Two bits in the Frame Relay header are used to signal the user device that congestion is occurring on the line: They are the Forward Explicit Congestion Notification (FECN) bit and the Backward Explicit Congestion Notification (BECN) bit. The FECN is changed to 1 as a frame is sent downstream toward the destination location when congestion occurs during data transmission. In this way, all downstream nodes and the attached user device learn about congestion on the line. The BECN is changed to 1 in a frame traveling back toward the source of data transmission on a path where congestion is occurring. Thus the source node is notified to slow down transmission until congestion subsides.



Bytes with congestion bits set according to DLCI value

Consolidated Link Layer Management (CLLM)

It may occur that there are no frames traveling back to the source node which is causing the congestion. In this case, the network will want to send its own message to the problematic source node. The standard, however, does not allow the network to send its own frames with the DLCI of the desired virtual circuit.

To address this problem, ANSI defined the Consolidated Link Layer Management (CLLM). With CLLM, a separate DLCI (number 1023) is reserved for sending link layer control messages from the network to the user device. The ANSI standard (T1.618) defines the format of the CLLM message. It contains a code for the cause of the congestion and a listing of all DLCIs that should act to reduce their data transmission to lower congestion.

Status of Connections (LMI)

Each DLCI corresponds to a PVC (Permanent Virtual Circuit). It is sometimes necessary to transmit information about this connection (e.g., whether the interface is still active) the valid DLCIs for the interface and the status of each PVC. This information is transmitted using the reserved DLCI 1023 or DLCI 0, depending on the standard used.

A multicast status may also be sent with the LMI. Multicasting is where a router sends a frame on a reserved DLCI known as a multicast group. The network then replicates the frame and delivers it to a predefined list of DLCIs, thus broadcasting a single frame to a collection of destinations.

Discard Eligibility (DE)

When there is congestion on the line, the network must decide which frames to discard in order to free the line. Discard Eligibility provides the network with a signal to determine which frames to discard. The network will discard frames with a DE value of 1 before discarding other frames.

The DE bit may be set by the user on some of its lower-priority frames. Alternatively, the network may set the DE bit to indicate to other nodes that a frame should be preferentially selected for discard, if necessary.

Frame Relay Standards

ANSI T1.618

T1.618 describes the protocol supporting the data transfer phase of the Frame Relay bearer service, as defined in ANSI T1.606. T1.618 is based on a subset of ANSI T1.602 (LAPD) called the “Core Aspects” and is used by both switched and permanent virtual calls.

T1.618 also includes the Consolidated Link Layer Mechanism (CLLM). The generation and transport of CLLM is optional. With CLLM, DLCI 1023 is reserved for sending link layer control messages.

T1.618 issues implicit congestion notification from the network to the user device. The congestion notification contains a code indicating the cause of the congestion and lists all DLCIs that have to reduce their traffic in order to lower congestion.

ANSI T1.617

To establish a Switched Virtual Circuit (SVC) connection, Frame Relay users must establish a dialog with the network using the signalling specification in T1.617. The procedure results in the assignment of a DLCI. Once the dialog is established, the T1.618 procedures apply.

To establish a Permanent Virtual Circuit (PVC), a setup protocol is used which is identical to the ISDN D-channel protocol and defined in T1.617.

With ISDN, users can use the D-channel for setup. For non-ISDN callers, there is no D-channel, so the dialog between the user and the network must be separated from the ordinary data transfer procedures. In T1.617, DLCI 0 is reserved.

T1.617 also contains specifications on how Frame Relay service parameters are negotiated.

ANSI LMI

ANSI LMI is a Permanent Virtual Circuit (PVC) management system defined in Annex D of T1.617. ANSI LMI is virtually identical to the Manufacturers’ LMI, without the optional extensions. ANSI LMI uses DLCI 0.

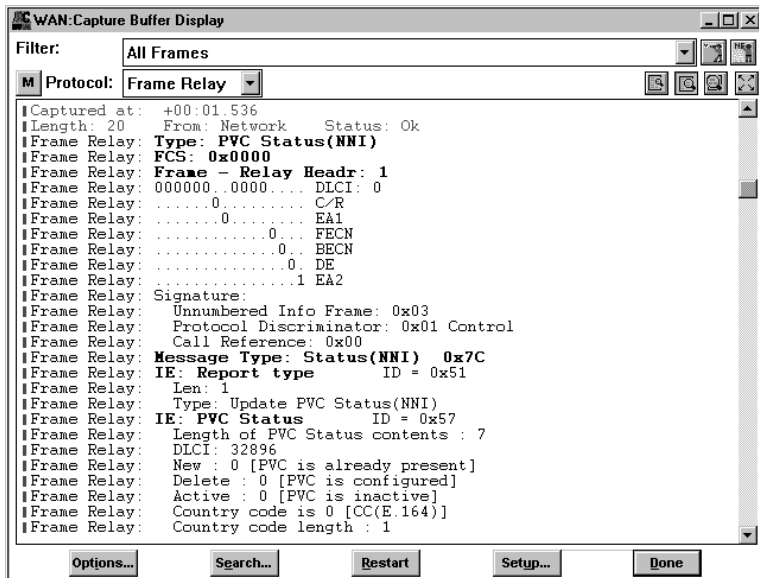
Manufacturers' LMI

Manufacturers' LMI is a Frame Relay specification with extension-document number 001-208966, September 18, 1990.

Manufacturers' LMI defines a generic Frame Relay service based on PVCs for interconnecting DTE devices with Frame Relay network equipment. In addition to the ANSI standard, Manufacturers' LMI includes extensions and LMI functions and procedures. Manufacturers' LMI uses DLCI 1023.

Frame Relay NNI PVC

Network-to-Network (NNI) PVC Frame Relay implementation is described in FRF.2. The NNI interface is concerned with the transfer of C-plane and U-plane information between two network nodes belonging to two different Frame Relay networks.



NNI PVC decode

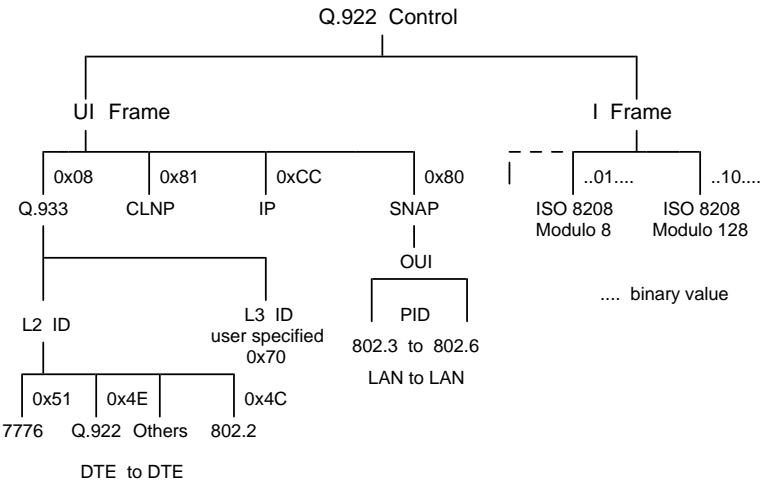
FRF.3

FRF.3 provides multiprotocol encapsulation over Frame Relay within an ANSI T1.618 frame. The structure of such a Frame Relay frame is as follows:

8	7	6	5	4	3	2	1	Octet
Flag (7E hex)								1
T1.618 Address (including 10-bit DLCI)								2 3
Q.922 Control (UI or I frame)								4
Optional Pad (0x00)								5
NLPID								6
Data . .								7
Frame Check Sequence								n-2 n-1
Flag (7E hex)								n

FRF.3 frame structure

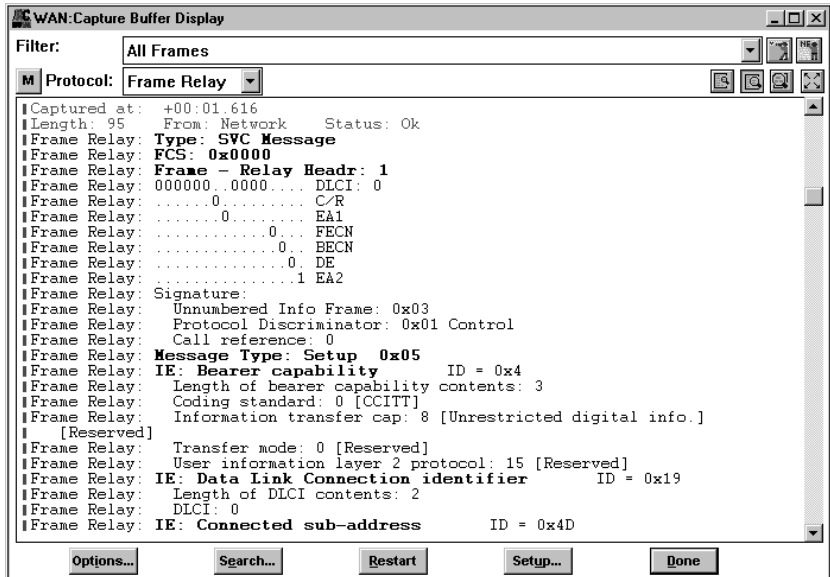
The NLPID (Network Level Protocol ID) field designates what encapsulation or what protocol follows. The following diagram details the possible values for NLPID and the protocols which are designated by each value. For example, a value of 0xCC indicates an encapsulated IP frame.



Multiprotocol encapsulation over Frame Relay

UNI SVC

FRF.4 is Frame Relay switched virtual connection user-to-network interface agreement. It applies using equipment attached to a non-ISDN Frame Relay network, or to an ISDN network using only case A.



UNI SVC decode

The following is a list of valid SVC message types:

- Call proceeding.
- Connect.
- Connect Acknowledge.
- Disconnect.
- Progress.
- Release.
- Release complete.
- Setup.
- Status.
- Status enquiry.

NNI SVC

Network-to-Network (NNI) SVC Frame Relay implementation is described in FRF.10. This implementation agreement applies to SVCs over Frame Relay NNIs and to SPVCs. It is applicable at NNIs whether both networks are private, both are public, or one is private and the other public. Such frames are automatically recognized by the analyzer and correctly displayed.

FRF.11

Frame Relay is now a major component of many network designs. The protocol provides a minimal set of switching functions to forward variable sized data payloads through a network. The basic frame relay protocol, described in the Frame Relay Forum User to Network (UNI) and Network to Network (NNI) Implementation Agreements, has been augmented by additional agreements which detail techniques for structuring application data over the basic Frame Relay information field. These techniques enabled successful support for data applications such as LAN bridging, IP routing, and SNA.

FRF.11 extends Frame Relay application support to include the transport of digital voice payloads. Specifically FRF.11 addresses the following requirements:

- Transport of compressed voice within the payload of a Frame Relay frame.
- Support a diverse set of voice compression algorithms.
- Effective utilization of low-bit rate Frame Relay connections.
- Multiplexing of up to 255 sub-channels on a single frame relay DLCI.
- Support of multiple voice payloads on the same or different sub-channel within a single frame.
- Support of data sub-channels on a multiplexed Frame Relay DLCI.

FRF.12

FRF.12 is the Frame Relay Fragmentation Implementation Agreement. Fragmentation queuing reduces both delay and delay variation in Frame Relay networks by dividing large data packets into smaller packets and then reassembling the data into the original frames at the destination. This ability is particularly relevant to users who wish to combine voice and other time-sensitive applications, such as SNA mission-critical applications, with non-time-sensitive applications or other data on a single Permanent Virtual

Circuit (PVC). The main benefit of fragmentation is the ability to utilize common User to Network Interface (UNI) access lines or Network to Network Interface (NNI) lines and/or PVCs for communications combining large data frames and real-time protocols.

Fragmenting frames enhances the utility and uniformity of Frame Relay networks, reducing delay and delay variation while upgrading application responsiveness, quality and reliability. As a result, multiple types of traffic, such as voice, fax and data, can be transparently combined on a single UNI, NNI and/or PVC.

The Fragmentation Implementation Agreement provides for transmission of Frame Relay Data Terminal Equipment (DTE) and Data Communications Equipment (DCE) with the ability to fragment long data frames into a sequence of shorter frames that are then reassembled into the original frame by the receiving peer DTE or DCE. Frame fragmentation is necessary to control delay and delay variation when real-time traffic, such as voice, is carried across the same interfaces as data traffic. Fragmentation enables the interleaving of delay-sensitive traffic on one PVC with fragments of long data frames on another PVC using the same interface.

FRF.12 supports three fragmentation applications:

1. Locally across a Frame Relay UNI interface between the DTE/DCE peers.
2. Locally across a Frame Relay NNI interface between DCE peers.
3. End-to-End between two Frame Relay DTEs interconnected by one or more Frame Relay networks.

When used end-to-end, the fragmentation procedure is transparent to Frame Relay network(s) between the transmitting and receiving DTEs.

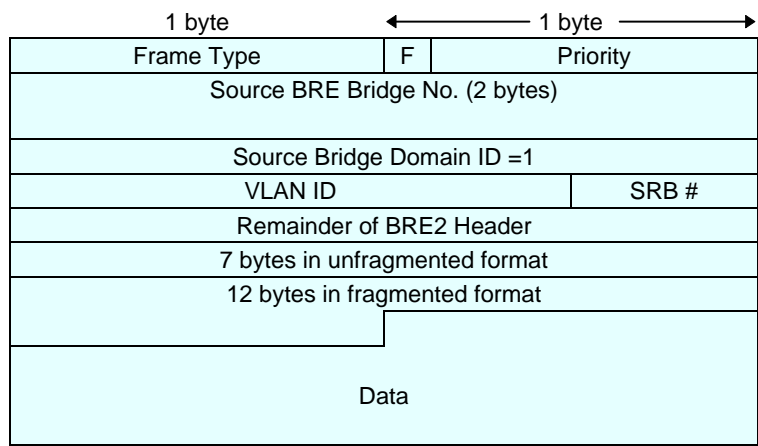
FREther

FREther is a variant of Frame Relay which is comprised of the Frame Relay header followed by the EtherType field. This is an additional form of encapsulation over Frame Relay which is used by some customers.

Timeplex (BRE2)

BRE (Bridge Relay Encapsulation) is a proprietary Ascom Timeplex protocol that extends bridging across WAN links by means of encapsulation. BRE2 is an improved form of the standard, providing better performance due to the fact that it sits directly on the link layer protocol, requires less configuration and provides its own routing protocol. BRE2 was deployed in 4.0 router software and is available in all 4.x and 5.x versions of the software.

The format of the BRE2 frame is as follows:



BRE2 frame format

If F is 0 than the frame is unfragmented and the BRE2 header is 17 bytes long. If F is 1, the frame is fragmented and the BRE2 header is 22 bytes long. SRB # is the Source Route Bridge Number (4 bits).

Cascade

In order to provide a Frame Relay service, Regional Bell Operating Companies (RBOCs) deploy Cascade switches in multiple LATAs and interconnect them to provide service to customers across the LATAs, as well as manage the switches in multiple LATAs from a single network management station.

The trunk header format for the Cascade STDX family of switches conform to ANSI T1.618-1991 ISDN Core Aspects of Frame Protocol for use with Frame Relay Bearer Service.

The Cascade trunk header format is as shown below:

1	2	3	4	8
R	C/R	Version		Reserved
ODE	DE	BECN	FECN	VC priority
Channel ID				
User Management data depending				

Cascade trunk header format

R
Reserved.

C/R
Command/Response field bit.

Version
Header version. Defines the version of the trunk header format for the Cascade STDX family of switches. This field is currently set to 0.

ODE
Set to 1 if the ingress rate is greater than the excess burst size.

DE
Discard eligibility indicator which is implemented based on the definitions specified in ANSI T1.618.

BECN

Backward explicit congestion notification according to ANSI T1.618.

FECN

Forward explicit congestion notification according to ANSI T1.618.

VC priority

Virtual circuit priority. Used to differentiate the sensitive traffic from traffic not sensitive to delays, such as file transfers or batch traffic. The priority may be 1, 2 or 3, 1 being the highest.

For management data, a 5th byte contains information concerning the type of PDU information to follow. Values are as follows:

- 0 Call request PDU
- 1 Confirmation PDU
- 2 Rejected PDU
- 3 Clear PDU
- 4 Disrupt PDU
- 5 Hello PDU
- 6 Hello Acknowledgment PDU
- 7 Defined Path Hello PDU
- 8 Defined Path Hello Acknowledgment PDU

LAPF

The purpose of LAPF is to convey data link service data units between DL-service users in the U-plane for frame mode bearer services across the ISDN user-network interface on B-, D- or H-channels. Frame mode bearer connections are established using either procedures specified in Recommendation Q.933 [3] or (for permanent virtual circuits) by subscription. LAPF uses a physical layer service, and allows for statistical multiplexing of one or more frame mode bearer connections over a single ISDN B-, D- or H-channel by use of LAPF and compatible HDLC procedures.

The format of the header is shown in the following illustration:

	8	7	6	5	4	3	2	1
Default address field format (2 octets)	Upper DLCI						C/R	EA 0
	Lower DLCI				FECN (Note)	BECN (Note)	DE (Note)	EA 1

LAPF address format

Control field bits (Modulo 128)	8	7	6	5	4	3	2	1	
I Format	N(S)							0	Octet 4 (Note)
	N(R)							P/F	
S Format	X	X	X	X	Su	Su	0	1	Octet 5
	N(R)							P/F	Octet 4
U Format	M	M	M	P/F	M	M	1	1	Octet 5
									Octet 4

LAPF format

EA
Address field extension bit.

C/R
Command Response bit.

FECN
Forward Explicit Congestion Notification.

BECN

Backward Explicit Congestion Notification.

DLCI

Data Link Connection Identifier.

DE

Discard Eligibility indicator.

D/C

DLCI or DL-CORE control indicator.

N(S)

Transmitter send sequence number.

N(R)

Transmitter receive sequence number.

P/F

Poll bit when used as a command, final bit when used as a response.

X

Reserved and set to 0.

Su

Supervisory function bit.

M

Modifier function bit.

Multiprotocol over Frame Relay

RFC 1490 <http://www.cis.ohio-state.edu/htbin/rfc/rfc1490.html>

RFC 2427 <http://www.cis.ohio-state.edu/htbin/rfc/rfc2427.html>

Multiprotocol over Frame Relay is a method of encapsulating various LAN protocols over Frame Relay. In this case, all protocols encapsulate their packets within a Q.922 Annex A frame. Additionally, frames must contain information necessary to identify the protocol carried within the protocol data unit (PDU), thus allowing the receiver to properly process the incoming packet. The format of such frames is as shown in the following illustration:

Flag (7E Hex)
Q.922 address
Control
Optional Pad (0x00)
NLPID
...
data
...
FCS
Flag (7E Hex)

Multiprotocol over Frame Relay frame

Q.922 address

2-octet address field containing the 10-bit DLCI field. On some networks the Q.922 address may optionally contain 3 or 4 octets.

Control

Q.922 control field. The UI value is of 0x03 is used unless negotiated otherwise. The use of XID (0xAF or 0xBF) is permitted.

Pad

Used to align the remainder of the frame to a two octet boundary. There may be 0 or 1 pad octet within the pad field. The value is always 0.

NLPID


Network Level Protocol ID, administered by ISO and CCITT. Identifies the encapsulated protocol.

FCS

2-byte frame check sequence.

There are two basic types of data packets that travel within the Frame Relay network: routed packets and bridged packets. These packets have distinct formats and therefore, must contain an indicator that the destination may use to correctly interpret the contents of the frame. This indicator is embedded within the NLPID and SNAP header information.

For those protocols that do not have a NLPID already assigned, it is necessary to provide a mechanism to allow easy protocol identification. There is a NLPID value defined indicating the presence of a SNAP header. The format of the SNAP header is as follows:



All stations must be able to accept and properly interpret both the NLPID encapsulation and the SNAP header encapsulation for a routed packet.

The three-octet Organizationally Unique Identifier (OUI) identifies an organization which administers the meaning of the Protocol Identifier (PID) which follows. Together they identify a distinct protocol. Note that OUI 0x00-00-00 specifies that the following PID is an Ethernet.

Some protocols have an assigned NLPID, but because the NLPID numbering space is so limited, not all protocols have specific NLPID values assigned to them. When packets of such protocols are routed over Frame Relay networks, they are sent using the NLPID 0x80 followed by SNAP. If the protocol has an Ethernet assigned, the OUI is 0x00-00-00 (which indicates an Ethernet follows), and PID is the Ethernet of the protocol in use. There is one pad octet to align the protocol data on a two octet boundary.

The second type of Frame Relay traffic is bridged packets. These are encapsulated using the NLPID value of 0x80 indicating SNAP. As with other SNAP encapsulated protocols, there is one pad octet to align the data portion of the encapsulated frame. The SNAP header which follows the

NLPID identifies the format of the bridged packet. The OUI value used for this encapsulation is the 802.1 organization code 0x00-80-C2. The PID portion of the SNAP header (the two bytes immediately following the OUI) specifies the form of the MAC header, which immediately follows the SNAP header. Additionally, the PID indicates whether the original FCS is preserved within the bridged frame.

