INTERNATIONAL NETWORK SERVICES

The knowledge behind the network.®

Frame Relay Encoding, Encapsulation, Framing, and Management Methods

Frame Relay Encoding, Encapsulation, Framing, and Management Methods

Introduction

The use of frame relay has grown significantly over the past several years as traditional host-based systems with dedicated point-to-point links have been replaced with distributed client/server networks. Frame relay tends to offer a cost-effective, reliable, and flexible wide-area network service. Frame relay costs are typically lower than dedicated point-to-point circuits since customers can aggregate multiple "virtual" circuits on a single physical circuit and customer premise equipment is relatively inexpensive. As the use of frame relay has grown, several physical and data link layer components have evolved, which has increased the efficiency and interoperability of frame relay devices.

This is a technical overview of frame relay and it discusses the various configuration options for data encoding, encapsulation, framing, and link management. As the use of frame relay continues to grow it is often configured with the parameters of signal encoding: B8ZS, signal framing: ESF, encapsulation frame format: IETF, point-to-point permanent virtual circuit (PVC) and lmi-type: ANSI. However, many IT professionals do not have a complete understanding of these parameters, their functionality, history, or other configuration options.

Frame Relay Background

The emergence of frame relay over the past several years is directly tied to changes in telephone carrier networks and the growing complexity of wide-area networks that are cost-effective, flexible, and scalable. Until the last several decades, telephone carrier networks consisted mainly of copper lines built to carry analog voice signals, which limited the bandwidth that could be provided for customers and caused high bit error rates on data transmission. To handle error conditions, the X.25 protocol was widely adopted in the 1970s because it had built-in error detection and correction capabilities. X.25 references the first three layers of the OSI model:

- Physical layer. RS232, X.21bis, and V.35.
- Data link layer. Link access procedure balanced (LAPB), a variant of high-level data link Control (HDLC).
- Network layer. Packet layer protocol (PLP).

The LAPB frame format is as follows:

Flag	Address	Control	Data	*Frame Check Sequence*	Flag
1 byte	1 byte	1 byte	Variable	2 bytes	1 byte

The data field of an X.25 frame is composed of a packet header and user data. The packet header includes a general format identifier (GFI), a logical channel identifier (LFI), and a packet type identifier (PTI). The GFI identifies the packet header format, the LCI identifies the locally significant virtual circuit, and the PTI identifies the type of X.25 packet.

X.25 packets can range from 64 to 4,096 bytes. Since the physical layer specification is based on RS-232, X.25 connections typically support a maximum transmission rate of 19.2 kbps (although higher speeds are possible). The performance of X.25 is also impacted by error-checking which happens at every intermediate switch rather than at the final destination. This can add substantial time to the latency of X.25 packets since each switch needs to process the entire packet and check it for errors prior to sending it on to the next switch.

Frame relay is a direct descendant of X.25 and is defined by the International Telecommunications Union (ITU-T) Q.922 and Q.933 standards. With the deployment of high-quality carrier networks and high-speed synchronous communication equipment, the built-in error detection and correction features of X.25 became unnecessary. To optimize data transmission, error handling is left to the destination devices. Unlike X.25 which defines the first three layers of the OSI model (physical, data link, and network) frame relay only defines the first two (physical and data link), which allows for greater flexibility of upper-layer protocols to be run across frame relay. Later sections in this article will discuss frame formats. In terms of performance, many router vendors support frame relay speeds in excess of four megabits (mbps) per second, while almost all digital service units (DSUs) support full T1 speed of 1.544 mbps. The frame relay specification also allows for a larger packet size than X.25. Frame relay can carry up to 8,193 bytes of data, although most providers only support a maximum packet size of 4,096 bytes.

The following terms, as defined in RFC 1604, obsoleted by 2954 in 2000, need to be understood to effectively configure and use frame relay:

- **Access channel.** An access channel generically refers to the DS1/E1 or DS3/E3-based UNI access channel or NNI access channel across which frame relay data travels. An access channel is the access pathway for a single stream of user data.
  - **T1**s are generally used in North America and have a bandwidth of 1.5 mbps. Within a given T1 line, an access channel can denote any one of the following:
    - *Unchannelized T1*. The entire T1 line is considered an access channel. Each access channel is comprised of 24 T1 time slots.
    - *Channelized T.* An access channel is any one of 24 channels. Each access channel is comprised of a single T1 time slot.
    - *Fractional T1*. An access channel is a grouping of *N* T1 time slots (*N*x56/64 Kbps, where *N* = 1-23 T1 time slots per FT1 access channel) that may be assigned in consecutive or non-consecutive order.

- **E1**s are generally used in Europe, South and Central America, and Asia and have a bandwidth of 2 mbps. Within a given E1 line, a channel can denote any one of the following:
  - *Unchannelized E1.* The entire E1 line is considered a single access channel. Each access channel is comprised of 31 E1 time slots.
  - *Channelized E1*. An access channel is any one of 31 channels. Each access channel is comprised of a single E1 time slot.
  - *Fractional E1.* An access channel is a grouping of *N* E1 time slots (*N*x64 Kbps, where *N* = 1-30 E1 time slots per fractional E1 access channel) that may be assigned in consecutive or non-consecutive order.
- **Access rate**. The data rate of the access channel, expressed in bits per second (bps). The speed of the user access channel determines how rapidly the end user can inject data into the network.
- **Committed burst size (Bc)**. The maximum amount of subscriber data (expressed in bits) that the network agrees to transfer, under normal conditions, during a time interval Tc.
- **Excess burst size (Be)**. The maximum amount of subscriber data (expressed in bits) in excess of Bc that the network will attempt to deliver during the time interval Tc. This data (Be) is delivered in general with a lower probability than Bc.
- **Committed information rate (CIR)**. The subscriber data rate (expressed in bps) that the network commits to deliver under normal network conditions. CIR is averaged over the time interval Tc (CIR = Bc/Tc).
- **DLCI**. Data Link Connection Identifier
- **Logical port**. This term is used to model the frame relay "interface" on a device.
- **Permanent virtual connection (PVC)**. A virtual connection that has its end points and bearer capabilities defined at subscription time.
- **Time slot (E1)**. An octet within the 256-bit information field in each E1 frame is defined as a time slot. Time slots are position-sensitive within the 256-bit information field. Fractional E1 service is provided in contiguous or non-contiguous time slot increments.
- **Time slot (T1)**. An octet within the 192-bit information field in each T1 frame is defined as a time slot. Time slots are position-sensitive within the 192-bit information field. Fractional T1 service is provided in contiguous or non-contiguous time slot increments.

Encoding

At the physical layer, all signals in frame relay networks are binary and so are represented by a one or a zero. The DSU is responsible for receiving these binary signals from the telephone carrier. To transmit binary signals on electronic media, voltage is raised or lowered. Most encoding methods are bipolar; they transmit successive values by alternating the voltage from positive to negative. In addition, since frame relay devices are typically synchronous (they do not use start and stop bits as with asynchronous devices), timing between devices is maintained by monitoring the transition of voltage from positive to negative or negative to positive. This can present problems, however, when the data being transmitted contains long strings of successive ones or zeros. To address this issue, several encoding mechanisms have been developed. Some of the most widely used methods include alternate mark inversion (AMI), high density bipolar 3 (HDB3), and bipolar 8-zero substitution (B8ZS).

Alternate Mark Inversion

AMI is one of the earlier means of encoding data on digital networks and is still used on many digital voice systems. Binary signals are represented by positive, negative, or no voltage. Ones are transmitted by positive or negative voltage while a zero is transmitted through the absence of voltage. With AMI, consecutive ones are transmitted by successive positive to negative voltage changes. While this eliminates timing problems that can be caused due to a long string of successive ones, it does not address the problems caused by a long string of successive zeros. To handle this situation, AMI performs a bipolar violation in every seventh bit. A bipolar violation is a successive, but not necessarily consecutive, one that is transmitted at the same voltage as the previous one. While this ensures that the devices stay in synchronization, despite a long string of consecutive zeros, it reduces the effective bandwidth since this “timing” bit takes the place of data, either a zero or an actual one, that could be transmitted.

High Density Bipolar 3

HDB3 is a bipolar encoding method based on AMI but provides efficiency through the use of a bipolar violation only when four or more consecutive zeros are transmitted. A bipolar violation is two successive ones transmitted at the same voltage; both are either positive or negative. Since successive ones are transmitted by changing the voltage from positive to negative, consecutive ones with both positive or negative voltage is a violation. This encoding technique keeps the devices synchronized so that if a long string of consecutive zeros are transmitted without using bandwidth to maintain timing as does AMI where every seventh bit is used to transmit a bipolar violation.

Bipolar 8 Zero Substitution

B8ZS, also referred to as “clear channel,” is a bipolar encoding method also based on AMI that is the predominate means of encoding used on frame relay data networks. The scheme inserts a bipolar violation into the signal encoding any time eight consecutive zeros are transmitted. A bipolar violation is two successive ones transmitted at the same voltage; both are either positive or negative. Since successive ones are transmitted by changing the voltage from positive to negative, consecutive ones with both positive or negative voltage is a violation. B8ZS inserts this violation into any string of zeros longer than seven bits so the receiving device can interpret it as a type of timing mark. This keeps the devices in synchronization despite the fact that the voltage is not changing (due to no ones being transmitted during the long string of zeros) without having to induce the overhead of sending a one just to maintain synchronization as is done with AMI.

B8ZS and HDB3 provide physical layer efficiency through their ability to maintain device synchronization without having to utilize bandwidth to ensure that a bipolar violation is sent on a regular basis. The selection of encoding on customer premise equipment needs to match the encoding used on the carrier’s frame relay device. B8ZS, along with AMI, is typically available in North America while HDB3 is often available in Europe.

Encapsulation

Frame relay encapsulation is configured on WAN devices such as routers and switches. Almost all frame relay devices support RFC-based frame formatting. A few vendors, such as Cisco, have also implemented their own encapsulation techniques. These proprietary encapsulation methods typically were developed prior to the publication of official industry standards.

Internet Engineering Task Force

Industry standards for frame relay encapsulation were published first in 1992 in RFC 1294 and subsequently replaced in 1998 by RFC 2427. The Internet engineering task force (IETF) frame format is as follows:

Flag	**Address**	**Control**	**Pad**	**NLPID**	**Data**	**Frame Check Sequence**	**Flag**
1 byte	2-4 bytes	1 byte	Up to 1 byte	1 byte	Variable	2 bytes	1 byte

The 2-byte address field is defined by ITU-T Q.922 that is 2 octets and includes a 10-bit DLCI. This address field can be increased to 4 bytes to allow for large DLCI addressing. The 2-byte address field allows for a 10-bit DLCI address that provides for 1,024 DLCIs (0-1023). The address field format is as follows:

8	**7**	**6**	**5**	**4**	**3**	**2**	**1**
DLCI, high order bits						CR	EA0
DLCI, low order bits				FECN	BECN	DE	EA1

A 3-byte address field allows for a 16-bit DLCI address that provides for 65,537 DLCIs (0-65,536). The address field format is as follows:

8	**7**	**6**	**5**	**4**	**3**	**2**	**1**
DLCI, high order bits						CR	EA0
DLCI				FECN	BECN	DE	EA1
DLCI, low order bits, or DL-Core						D/C	EA3

A 4-byte address field allows for a 23-bit DLCI address that provides for 8,388,609 DLCIs (0-8,388,609). The address field format is as follows:

8	**7**	**6**	**5**	**4**	**3**	**2**	**1**
DLCI, high order bits						CR	EA0
DLCI, low order bits				FECN	BECN	DE	EA1
DLCI							EA3
DLCI, low order bits, or DL-Core						D/C	EA4

- The **CR** bit is the command or response bit used to transmit information between stations such as connection management (i.e., session establishment) and frame acknowledgement.
- The **FECN** and **BECN** bits are used for congestion notification. The FECN bit is set to alert receiving devices if the frame experiences congestion. The BECN bit is set on frames traveling in the opposite direction of frames that experience congestion to notify those devices that congestion has occurred in that path.
- **DE** is the discard eligible bit. A circuit's committed information rate (CIR) determines the minimum rate that a carrier will receive data without making a frame discard eligible. If a frame is received by the carrier network at a rate greater than the CIR, the provider has the right to discard the frame. In addition, many routers can set the DE bit to perform traffic prioritization. On Cisco routers the command "frame-relay de-list" can be used to identify which types of packets can be marked as DE. The command "frame-relay de-group" can be used to apply that list to a particular DLCI.
- **EA** bits are known as extension address. These are used by frame relay to allow for a variable length DLCI. By setting the EA bit to one, the address field is extended by one byte in which six more bits are available for a DLCI.
- **D/C** is the DLCI core indicator. The three and four-byte extended addresses use the D/C to determine if there is additional DLCI information in the last byte or whether the information in the low-order DLCI field should be interpreted as DL-core information. However, the use of DL-core information has not yet been defined.

The Control field is similar to other HDLC-based formats. This field can be used for information (transfer of information between higher layer protocols), supervisory (flow control), and unnumbered (session establishment).

The Pad field is optional and is used to align the remainder of the frame to a two-octet boundary.

The network layer protocol identifier (NLPID) field is used to identify the type of upper-layer protocol transmitted in the frame. Not all protocols, however, have a predefined NLPID. In this case, the NLPID is set to 0x80 that identifies the protocol as Subnetwork Attachment Point (SNAP). A standard SNAP header then follows in the data field. This takes the following format:

Organizationally Unique Identifier	Protocol Identifier
12 bits (3 octets)	8 bits (2 octets)

The organizationally unique identifier (OUI) identifies an organization that administers the values of the protocol identifier (PID). Together the OUI and PID identify the protocol. As examples, 0x00-00-00-00-01 is 802.3 (Ethernet) and 0x00-00-00-00-03 is 802.5 (Token Ring).

Cisco

The Cisco-proprietary frame relay encapsulation method predates the IETF standard. While it is similar to IETF, the two methods are not compatible. The difference between the two formats is that the Cisco format does not provide for the use of an NLPID but rather uses the Ethertype to identify the type of protocol being encapsulated. The Ethertype is an IEEE code used to identify the protocol type found at the thirteenth and fourteenth octets of an Ethernet frame. For protocols without an Ethertype field, Cisco utilizes a proprietary set of values that is consistent with those values defined by Cisco for other router features where the Ethertype field is used. The Cisco frame format is as follows:

Flag	Address	Ethertype	Data	Frame Check Sequence	Flag
1 byte	2-4 bytes	2 bytes	Variable	2 bytes	1 byte

This Cisco format is only available on Cisco products and is not licensed to other vendors. As such, when using Cisco equipment in a heterogeneous environment (i.e., a frame relay network with Bay routers), the IETF frame format must be used. However, the Cisco format is the default frame relay encapsulation method. When configuring a serial interface for frame relay, the command to set the encapsulation is:

encapsulation frame-relay [cisco | ietf]

If neither "cisco" nor "ietf" is specified, the interface will default to "cisco." Additionally, different DLCIs can be set for encapsulation types with the following command:

frame-relay interface-dlci *dlci* [cisco | ietf]

Framing

Framing technology is used to synchronize the channels on the frame relay circuit. A frame relay T1 can be split into 24 independent 56 or 64 kbps transmission channels often referred to as DS0s. Each channel carries information in eight-bit bytes, so a frame consists of 192 information bits (eight bits multiplied by 24 channels). Bit 193, known as the framing bit, indicates to receiving equipment when a frame starts and finishes. Without it, sending and receiving equipment would not be able to synchronize their frames. Each byte is updated 8000 times a second; therefore, transmission speed is determined by multiplying 193 bits by 8000 cycles, which equals 1.544 mbps.

A framing pattern consists of a specified number of frame bits that are repeated. The number of frame bits and the functionality of the framing pattern vary among predominate framing formats. There are two main framing formats for T1s: super frame (SF) and extended super frame (ESF). SF uses a total of 12 frame bits; therefore, 12 frames need to be sent in order to transmit a framing pattern. ESF repeats every 24 frames instead of 12 frames; therefore, 24 frames need to be sent in order to transmit an ESF framing pattern. Unlike SF, ESF can be used for more than just frame synchronization. ESF also provides cyclic redundancy check and data link control.

The choice of either SF or ESF will not affect bandwidth utilization. However, the decision to use SF or ESF should be based on the types of equipment in the network and must be the same for pieces of equipment that communicate with each other.

Management

Frame relay management was initially designed to provide a signaling mechanism to exchange information regarding the status of PVCs and to ensure that the link is operating correctly in the absence of data traffic. In 1990, a consortium of frame relay device vendors published the local management interface (LMI) protocol. This consortium was comprised of Digital Equipment Corporation, Cisco, Nortel, and Stratacom. Subsequent variations to the LMI were developed and are known as Annex-D (ANSI) and Annex-A (Q.933A). Link management runs on a dedicated PVC; therefore, its operation does not impact actual user data. On Cisco routers, the following command is used to configure the management type:

frame-relay lmi-type [ansi | cisco | q933a]

The keyword "cisco" assigns the management type of link management. If a management type is not specified, *cisco (LMI)* will be used as the default. The most popular link management protocols are LMI and Annex-D since they provide the greatest level of function without causing additional overhead. Annex-D can be considered more desirable than LMI because Annex-D can report on the status of individual PVCs while LMI reports on the status of the entire frame relay circuit.

LMI (Cisco)

LMI provides for 992 virtual circuit addresses and uses DLCI 1023 as a management circuit. The DLCIs are assigned as follows:

DLCI	Function
0	Call control and signaling
1-15	Reserved for future use
16-1,007	Virtual circuit addresses
1,008-1,022	Reserved for future use (such as multicasting)
1,023	LMI

Annex-D (ANSI)

Annex-D provides for 976 virtual circuit addresses and uses DLCI 0 as the management circuit. The DLCIs are assigned as follows:

DLCI	Function
0	Management messages for link integrity and call signaling
1-15	Reserved for future use
16-991	Virtual circuit addresses
9921,007	Circuits for carrying consolidated link layer management (CLLM). These are used by frame relay services to carry information about the network.
1,008-1,022	Reserved for future use
1,023	CLLM used to send management messages regarding higher layer protocols

Annex-A (Q933A)

Annex-A is very similar to Annex-D. Both forms for management use DLCI 0 to communicate. Annex-A, however, does not implement several of the features seen in Annex-D and LMI, such as the D bit (deleted) and R bit (congestion or receiver-not-ready).

Message Format

Link management messages for LMI, Annex-D, and Annex-A have the same format. They contain a five-byte header, a one-byte message type identifier, one or more information elements (IEs) of variable length, and a two-byte CRC. A message header includes a two-byte frame relay header (DLCI, FECN, BECN, etc., as discussed earlier). The rest of the header has a one-byte control field (typically set to 03H), a one-byte protocol discriminator (set to 09H for LMI and 08H for Annex-D and Annex-A), and a one-byte call reference field (set to 00H as it is no longer used).

Control	**Protocol Discriminator**	**Call Reference**	**Message Type**
1 byte (03H)	1 byte (09H or 08H)	1 byte (00H)	1 byte (75H or 7DH)

The Message Type field will either be set to a value of 75H (for a status enquiry message) or 7DH (for a status message). Following the header there are three bytes which provide the IE identifier, the length of the IE, and the type of IE. The type of report is either full status of all PVCs or a sequence number exchange for link integrity verification; in addition, Annex-D and Annex-A can report on that status of a single PVC. The information elements follow these fields. The message format ends with a two-byte frame check sequence and a one-byte flag field.

Management Messages

Link management protocols are designed so frame relay subscribers originate all exchanges. The subscriber begins a message exchange by sending a status enquiry message that the provider responds with a status message. There are several resources available on the Internet which describe these messages, including those listed in the references section of this white paper. The information in the following subsections has been compiled from several of these resources, especially *A Tutorial of the Frame Relay Protocol* and the *AIT Workbook*.

An exchange of management messages can perform either of two functions:

- A "heartbeat" exchange that verifies that the link is operational
- A report regarding the status of each individual DLCI

Management frames contain a six-byte header followed by a list of IEs. The IEs convey heartbeat and status information.

Management Frame Information Elements

The basic LMI protocol supports three IEs:

- Report type
- Keep-alive
- PVC status

Management messages contain a report type and a keep-alive element. A full status message from the network also contains one PVC status element for each PVC on the link.

The keep-alive information element contains a pair of eight-bit sequence numbers, current and last received, through which the heartbeat process maintains a running check on the health of the link.

Between every 5 to 30 seconds (recommended default is 10 seconds) the heartbeat process sends a status enquiry message that contains a report type value of sequence number exchange and a keep-alive element. When the provider network receives the message, it records the current sequence number as its last received sequence number, increments it by one to produce its new current sequence number, and transmits a status message with a keep-alive element that contains the new numbers. If either side receives a heartbeat message in which the sequence numbers don't follow correctly, it may declare an LMI sequence error. LMI replies to a status enquiry with a keep alive identifier that reports the sequence number exchanges. Annex-D and Annex-A provide link integrity verification as well as the sequence number exchange.

After a specified number of sequence number exchanges, the subscriber sends a status enquiry with a value of full status in report type element. The network answers with a status message containing a PVC status information element for each DLCI currently defined for the link. These messages are:

- N (new) bit
- A (active) bit

The N bit is set to one when the PVC status element is reporting on a newly defined DLCI. The N bit will be reset to zero in all subsequent PVC status elements for that DLCI.

The A bit is set to one when the PVC to which the element refers is sending and receiving data. The A bit can be used by Cisco routers to determine if a backup circuit should be used in the event that a frame relay PVC fails. The command “backup interface” is used to configure a secondary or dial backup interface on Cisco routers. This feature uses the A bit to determine a PVC’s status to automatically bring up or shut down a backup interface such as an ISDN or analog circuit.

Functions of the D (deletion) and R (congestion or receiver-not-ready) bits are defined by optional LMI extensions.

Management Extensions

The LMI specification also includes several optional extensions that may or may not be implemented by frame relay providers. These include:

- Flow control
- Minimum bandwidth available
- Provider-initiated status updates

Flow Control

The capability provides a way for the network to report congestion to the subscriber. The flow control feature uses the optional R bit in the PVC status IE as a “receiver-not-ready” signal for the PVC whose status is being reported. A value of one in the R bit indicates congestion while zero indicates no congestion. The R bit is not supported in Annex-A.

Minimum Bandwidth Available

This feature communicates the minimum bandwidth available on the network to the PVC. In most implementations this number will be the PVC's CIR. However, some providers may use this feature to respond to changing traffic conditions by dynamically increasing or decreasing the bandwidth available to individual PVCs.

Global Addressing

Frame relay DLCI values are typically only of local significance. And thus, they do not provide a mechanism to identify remote network interfaces or devices attached to these interfaces. The use of global addressing allows devices to address each other through globally defined DLCI values.

Multicasting

The DLCI values of 1,019-1,022 have been reserved for multicasting. Frames sent on one of these DLCIs are replicated to all network exit points for a specified group of destinations.

Provider-Initiated Status Update

Provider-initiated status update allows the network to communicate changes in a PVC's status by means of a status update without first receiving a status enquiry from the subscriber. The status update contains only PVC status information elements so it cannot function in the heartbeat process. Further, it contains status elements for only those PVCs whose status has changed. Changes reported include:

- Deletion of a PVC (reported by setting the optional D bit of the status element)
- Changes in the minimum bandwidth allocated to a PVC
- Activation or deactivation of a PVC (indicated by setting or clearing the A bit)
- Flow control information (changes in congestion status, signaled by setting or resetting the R bit).

Consolidated Link Layer Management

CLLM is a signaling protocol for frame relay that predates LMI. The original frame relay specification defines an optional consolidated link layer management (CLLM) message. CLLM's major function is to augment the BECN system for the reporting of congestion by allowing the network to report in the “backward” direction, in the absence of user data traffic, explicit congestion notification. CLLM transmits cause codes that include the following: network congestion due to excessive traffic, facility or equipment failure, maintenance, and unknown.

About INS

INS (International Network Services Inc.) is a leading global provider of vendor-independent network consulting and security services. We offer a full range of consulting services to help companies build, optimize, manage, and secure their network infrastructures to enable their business initiatives and achieve a sustainable operating advantage. INS is a recognized leader in complex, multi-vendor network consulting, having helped more than 75% of the Fortune 500 and delivered more than 15,000 engagements over the past decade. Headquartered in Santa Clara, CA, INS has regional offices throughout the United States and internationally in the UK, the Netherlands, Germany, and France. For additional information, please contact INS at 1-888-767-2788 in the U.S., 44 (0) 1628 50300 in Europe, or 1-650-318-1020 worldwide, or visit www.ins.com.

References

DGS6910 Digital Telephony & Switching, Monash University,
http://www.dgs.monash.edu.au/~jwb/subjects/dgs6910/spc/spc_sw.html

HDB3, University of Aberdeen Electronics Research Group,
http://www.erg.abdn.ac.uk/users/gorry/course/phy-pages/hdb3.html

T1 Repeatered Line Transmission Engineering, Bob Gross, Larus Corporation,
http://www.laruscorp.com/t1tut.htm

Why you should consider T1, Certified Consultants and Systems, http://www.hb.quik.com/~ccs/t1.html

Cisco IOS Software Command Summary, Cisco Systems Inc.

Frame Relay Internetworking, revision 1.2, American Research Group, Stiefelmeyer International Ltd.

ITU-T Recommendation Q.933, International Telecommunications Union,
http://www.itu.int/itudoc/itu-t/com11/implgd/q933_ww7.doc

AIT Workbook, Frame Relay Traffic in an IPX Network, Cisco Systems Inc.

Frame Relay, author unknown, Cisco Systems Inc.

RFC 1604 obsoleted by 1596, obsoleted by 2954 Definitions of Managed Objects for Frame Relay Service
ftp://ftp.isi.edu/in-notes/rfc2954.txt

RFC 1294 obsoleted by 1490, obsoleted by 2427 Multiprotocol Interconnect over Frame Relay
ftp://ftp.isi.edu/in-notes/rfc2427.txt

Frame Relay Forum http://www.frforum.com/