



Network General Corporation


SNIFFER® NETWORK ANALYZER NETWORK AND PROTOCOL REFERENCE

**SNIFFER® NETWORK ANALYZER
NETWORK AND PROTOCOL REFERENCE**





Network General Corporation



SNIFFER[®] NETWORK ANALYZER

NETWORK AND PROTOCOL REFERENCE

DISCLAIMER OF WARRANTIES

The information in this document has been reviewed and is believed to be reliable; nevertheless, Network General Corporation makes no warranties, either expressed or implied, with respect to this manual or with respect to the software and hardware described in this manual, its quality, performance, merchantability, or fitness for any particular purpose. The entire risk as to its quality and performance is with the buyer. The software herein is transferred "AS IS."

Network General Corporation reserves the right to make changes to any products described herein to improve their function or design.

In no event will Network General Corporation be liable for direct, indirect, incidental or consequential damages at law or in equity resulting from any defect in the software, even if Network General Corporation has been advised of the possibility of such damages. Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you.

This document is copyrighted and all rights are reserved. This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from Network General Corporation.

*Distributed Sniffer System and SniffMaster are trademarks of Network General Corporation.
All other registered and unregistered trademarks in this document are the sole property of their respective companies.*

*© Copyright 1986 – 1992 by Network General Corporation.
All rights reserved. Present copyright law protects not only the actual text, but also the "look and feel" of the product screens, as upheld in the Atari and Broderbund cases.*

Document prepared by Ed Dalton, Jim Mar, Bryan Embrey, and Arlene Brenner.

November 1992

P/N 20026-004

Table of Contents

Preface	vii
About This Manual	vii
Organization of This Manual	vii
Manuals for the Sniffer Network Analyzer	vii
Other Sources of Information	viii
On-Line Help	viii
Tutorial	viii
Technical Support	viii
Chapter 1. Network Architectures	1-1
Overview	1-3
Ethernet® Network Architecture	1-3
Physical Interconnection and Speed	1-3
Access Control	1-6
Other Transceiver Functions	1-6
Format of a Frame	1-7
Assignment of Network Addresses	1-11
Token Ring Network Architecture	1-12
Physical Interconnection and Speed	1-12
Logical Interconnection	1-13
Access Control	1-13
Format of a Token Frame	1-14
Format of a Data Frame	1-15
Assignment of Network Addresses	1-16
Fiber Distributed Data Interface Network Architecture	1-16
Interconnection and Speed	1-18
Access Control	1-18
Format of a Token Frame	1-19
Format of a Data Frame	1-19
Station Management	1-19
WAN/Synchronous Architecture	1-20
Interconnection and Speed	1-20
Format of a Frame	1-21
Phases of Link Control	1-24
StarLAN™ Network Architecture	1-25
Physical Interconnection and Speed	1-25
Access Control	1-26
Format of a Frame	1-26

Format of the Data	1-26
Assignment of Network Addresses	1-26
ARCNET® Network Architecture	1-26
Physical Interconnection and Speed	1-26
Logical Interconnection and Access Control	1-27
Establishing and Maintaining the Token Sequence	1-28
Transmitting Data Frames	1-29
Format of a Frame	1-30
Format of the Data	1-31
IBM PC Network™ (Broadband) Architecture	1-31
Speed and Transmission Technique	1-32
Physical Interconnection	1-33
Access Control	1-33
Format of a Frame	1-34
Format of the Data	1-34
Assignment of Network Addresses	1-34
LocalTalk® Network Architecture	1-34
Physical Interconnection and Speed	1-34
Access Control	1-35
Format of a Frame	1-36
Format of the Data	1-37

Chapter 2. Major Protocol Suites **2-1**

Overview	2-3
Introduction	2-3
IBM® Protocol Interpreter Suite	2-5
Protocols Interpreted	2-5
Novell® NetWare® Protocol Interpreter Suite	2-7
Protocols Interpreted	2-8
XNS™ Protocol Interpreter Suite	2-10
Protocols Interpreted	2-11
TCP/IP Protocol Interpreter Suite	2-12
Protocols Interpreted	2-13
SUN® Protocol Interpreter Suite	2-16
Protocols Interpreted	2-17
ISO Protocol Interpreter Suite	2-18
Protocols Interpreted	2-18
TCP/IP Frames	2-20
DECnet® Protocol Interpreter Suite	2-21
Protocols Interpreted	2-22

Banyan® VINES™ Protocol Interpreter Suite	2-24
Protocols Interpreted	2-25
AppleTalk® Protocol Interpreter Suite	2-27
Protocols Interpreted	2-27
X Windows™ Protocol Interpreter Suite	2-30
Features of the Interpreter	2-30
X.25 Protocol Interpreter Suite	2-32
Protocols Interpreted	2-33
Frame Relay Protocol Interpreter Suite	2-34
Protocols Interpreted	2-35
SMT Protocol Interpreter Suite	2-36
Features of the Interpreter	2-37
Appendix A. Glossary	A-1
Appendix B. Bibliography	B-1
General References	B-3
Network References	B-3
Protocol References	B-5

List of Figures

Chapter 1. Network Architectures	1-1
Figure 1-1. Frame format for Ethernet, StarLAN, and PC Network.	1-7
Figure 1-2. Original "Ethernet" data format defined by Xerox.	1-8
Figure 1-3. IEEE 802.3 data format defined by IEEE/ANSI/ISO standards organizations.	1-8
Figure 1-4. The original data format used by IBM and Sytek (applies only to PC Network).	1-9
Figure 1-5. Indications used in LLC supervisory frames.	1-10
Figure 1-6. Commands and responses in LLC unnumbered frames.	1-10
Figure 1-7. Frame format for a token frame.	1-14
Figure 1-8. Frame format for a token ring data frame.	1-15
Figure 1-9. LLC frame format for Sub-Network Access Protocol (SNAP).	1-16
Figure 1-10. Relationship of FDDI standards.	1-17
Figure 1-11. Token frame format for FDDI.	1-19
Figure 1-12. Data frame format for FDDI.	1-19
Figure 1-13. The general HDLC (including LAPB and SDLC) frame format.	1-21
Figure 1-14. Commands and responses used in HDLC (SDLC and LAPB) supervisory frames.	1-22
Figure 1-15. Commands and responses used in HDLC (SDLC and LAPB) unnumbered frames.	1-23
Figure 1-16. Five ARCNET messages.	1-30
Figure 1-17. A data frame as recorded and "normalized" by the Sniffer Network Analyzer.	1-31
Figure 1-18. Frequency assignments for different versions of the PC Network adapter.	1-33
Figure 1-19. LLAP frame format.	1-36
Figure 1-20. Indications used in LLAP control frames.	1-36
Figure 1-21. LLAP data frame format.	1-37
Chapter 2. Major Protocol Suites	2-1
Figure 2-1. Layers of the OSI Network Model.	2-4
Figure 2-2. The IBM PI suite related to the OSI reference model.	2-5
Figure 2-3. The Novell NetWare PI suite related to the OSI reference model.	2-8
Figure 2-4. The XNS PI suite related to the OSI reference model.	2-10
Figure 2-5. The TCP/IP PI suite related to the OSI reference model.	2-12
Figure 2-6. The SUN PI suite related to the OSI reference model.	2-16

Figure 2-7. The ISO PI suite related to the OSI reference model.	2-18
Figure 2-8. The DECnet PI suite related to the OSI reference model.	2-21
Figure 2-9. The Banyan VINES PI suite related to the OSI reference model.	2-24
Figure 2-10. The AppleTalk PI suite related to the OSI reference model. . .	2-27
Figure 2-11. The X Windows PI suite related to the OSI reference model. . .	2-31
Figure 2-12. The X.25 PI suite related to the OSI reference model.	2-32
Figure 2-13. The Frame Relay PI suite related to the OSI reference model. .	2-34
Figure 2-14. The SMT PI suite related to the OSI reference model.	2-36

Preface

About This Manual

The *Network and Protocol Reference Manual* provides background information on a broad spectrum of network types and communication protocols. You will want to refer to it from time to time to help you get the most out of your Sniffer Network Analyzer.

Organization of This Manual

The following table describes the organization of this manual.

Chapter/Appendix	Contents
Chapter 1, "Network Architectures"	Covers basic knowledge about local and wide area network architectures.
Chapter 2, "Major Protocol Suites"	Covers basic knowledge about network protocols and protocol interpreters.
Appendix A, "Glossary"	Provides a glossary of data communications terminology.
Appendix B, "Bibliography"	Provides a bibliography of source material.

Manuals for the Sniffer Network Analyzer

Two types of manuals accompany the Sniffer Network Analyzer. The primary manuals, which include this manual, describe the system's normal operations; the supplementary manuals describe the programs that configure and test the system's various hardware and software components for troubleshooting. The actual manuals in your shipment depend on the system configuration.

The following table describes the primary manuals.

For Information On...	Read...
Installing and configuring the Sniffer Network Analyzer.	<i>Sniffer Network Analyzer: Installation Guide</i>
Operating the analysis functions on an Ethernet or Token-Ring network.	<i>Sniffer Network Analyzer Operations</i>
Operating the monitor functions on an Ethernet network. Using the monitor features effectively to detect network abnormalities.	<i>Sniffer Network Analyzer: Ethernet Monitor Operations</i>

For Information On...	Read...
Operating the monitor functions on a token ring network. Using the monitor features effectively to detect network abnormalities.	<i>Sniffer Network Analyzer: Token-Ring Monitor Operations</i>
Various network and protocol types.	<i>Sniffer Network Analyzer: Network and Protocol Reference</i>

Other Sources of Information

Network General Corporation (NGC) provides other sources of information that can help you get familiar with the Sniffer Network Analyzer.

On-Line Help

After highlighting an item in a console, analyzer, or monitor menu, you can see a phrase or sentence in a panel near the bottom of the screen. It explains the meaning of the highlighted item.

If you want to obtain general information on a particular feature of the Sniffer Network Analyzer, press F1 at any time. A window containing a list of topics opens. If you are displaying a monitor statistics screen, pressing F1 gives you information on the current screen.

Tutorial

NGC makes available a booklet with accompanying diskette entitled *Real Networks. Real Problems*. It presents case studies based on data captured with a Sniffer network analyzer from four different networks. The Sniffer analyzer and the server's analysis application have different capabilities, but the case studies allow you to see how investigation of a network problem proceeds.

You can obtain the tutorial free of charge from any of the company's sales representatives or directly from NGC.

Technical Support

A toll-free number is available to obtain technical support for the Sniffer Network Analyzer:

Phone for Network General's Technical Support Department:	(800) 395-3151
FAX:	(415) 327-9436



Network Architectures

Overview

The network reference material in this chapter covers the physical and data link layers of various network types. Each section summarizes the features of each particular network and the variations related to its use and understanding.

Ethernet[®] Network Architecture

Ethernet is a type of local area network (LAN) suitable for high-speed interconnection of computers and computer-controlled devices over moderate distances. The Ethernet architecture is defined by implementations from many manufacturers as ANSI/IEEE standard 802.3, ISO/DIS standard 8802/3, and FIPS standard 107.

There are several variations permitted by the standards and other variations that are not official but are, nonetheless, popular. To say that a network is an “Ethernet” means only that it is one of several closely-related, but not necessarily compatible, LANs. Even use of the word “Ethernet” may be confusing, since it is sometimes reserved to refer to the earlier DEC/Intel/Xerox (DIX) network, leaving 802.3 networks to designate the others.

Some system implementations of the Ethernet network also use at least a subset of a similarly standardized protocol for Logical Link Control (LLC) as defined by ANSI/IEEE standard 802.2 and ISO/DIS standard 8802/2.

Physical Interconnection and Speed

Stations connected to a conventional Ethernet network are all connected to the same bus, so that every station hears what any station transmits. The delay between transmission and reception depends only on the propagation delays through the wires and attaching devices. In this way Ethernet is similar to networks like StarLAN and IBM PC Network (see “StarLAN™ Network Architecture” on page 1–25, and “IBM PC Network™ (Broadband) Architecture” on page 1–31). It differs fundamentally, however, from networks like the IEEE 802.5 token ring, where stations are wired in a *logical ring* so that each station only hears what its upstream neighbor transmits (see “Token Ring Network Architecture” on page 1–12).

The most common Ethernet transmission speed is 10 megabits per second (Mbps), or 1,250,000 bytes per second, sent in baseband (non-modulated, non-RF) form.¹ The ANSI/IEEE documents refer to this as the 10BASE5 (10 Mbps, baseband, and 500 meters per segment) standard. Another common standard is 1BASE5 (1 Mbps, baseband, and 500 meters per segment) over twisted-pair cable that is based on AT&T's 1 Mbps StarLAN. A third standard

1. Because of the access control mechanism, the effective network throughput is significantly less than the transmission speed.

is 10BASE2 (10 Mbps, baseband, and 185 meters per segment). A fourth standard is 10BASE-T (10 Mbps, baseband, over twisted pair). The Ethernet Sniffer analyzer supports 10BASE5 networks, 10BASE2 networks, and some variants of them.

Thick Ethernet

There are two common schemes for the wiring of stations into an Ethernet. The original *thick Ethernet* scheme uses a backbone of yellow, semi-rigid, 0.4-inch diameter coaxial cable segments for up to 100 stations per segment. Each segment can be a maximum of 500 meters long, and segments may be connected with repeaters, subject to the restriction that there are no more than two repeaters in the path between any two stations. Thus the maximum total cable length between two stations is 1500 meters but is often less than that in practice because of the way the cable is routed. The cable impedance is 50 ohms, and segments must be terminated on each end with a 50-ohm terminator attached with an N-series coaxial connector. The shield conductor must be grounded to an earth reference at only one point and fully insulated everywhere, including at connectors and terminators, to prevent shock hazards.

A station is connected to a thick Ethernet cable by means of a *transceiver*, which is a signal converter that must be attached directly to the cable. Transceivers are called *medium attachment units* (MAU) by the IEEE Standards documents. Many transceivers clamp on to the cable and use a *vampire tap* to make contact with the outer shield and inner conductor without requiring that the cable be cut. The spacing between transceivers must be a multiple of 2.5 meters; the yellow cable has black marks to indicate possible transceiver attachment points.

A single station is connected to its transceiver with a more flexible 4-pair *drop cable* or *Attachment Unit Interface (AUI) cable* whose maximum length is 50 meters. In addition to carrying network data in each direction on separate pairs, the drop cable also supplies power for the transceiver electronics from the equipment being attached to the network. Both ends of the drop cable use DB-15 connectors: a male connector with locking posts for the equipment end and a female connector with a slide latch for the transceiver end. Note that the slide latch is often too big for the back panel of personal computers, so at least two variations of this attachment scheme are in common use.

To reduce the per-station cost of the transceiver and to circumvent the limit of 100 stations per segment, transceiver fan-out boxes are available. These typically allow four or eight drop cables to be connected to a single box that is connected to a standard transceiver clamped to the coaxial cable segment.

Thin Ethernet

Another approach to reducing the cost of an Ethernet installation is *thin Ethernet*, "cheapernet," or 10BASE2. It dispenses with the large, semi-rigid coaxial cable and separate external transceivers. Instead, it uses flexible RG-58/U coaxial cable and transceiver circuitry built into each attaching computer. For this wiring scheme, the overall topology is still a linear segment but the segment passes by the back of each computer. Attachment is made by splicing BNC connectors into the cable and inserting a "T" connector. The

segments are still terminated with 50-ohm terminators at each end. The maximum length of a segment is usually 185 meters, although some vendors support longer distances.

Cheapernet transceivers are generally part of the network adapter circuitry, and the exposed connector is a female BNC. A hybrid approach is also possible, since there are *thin Ethernet transceivers* available that attach with a standard drop cable to an attaching computer but have a female BNC instead of a *vampire tap* for connecting to the network. There are also *thin to thick adapters* that allow segments to be built out of both kinds of cable. Some devices provide both a BNC connector (for thin Ethernet) and a DB-15 (AUI) connector (for a thin- or thick-Ethernet transceiver). On the Sniffer analyzer, you determine which connector is active by the position of a jumper.

Twisted Pair Ethernets

Ethernet over twisted pair has been standardized by an IEEE 802.3 project, known as the 10BASE-T standard. The standard provides a means for attaching AUI-compatible devices to 24 gauge, unshielded twisted pair cable, instead of the usual coaxial media.

Stations connect their standard AUI connector to a special twisted pair transceiver (called a *Twisted Pair MAU* by the IEEE Standard documents) that has an RJ-45 phone jack for the other connection. The transceiver is then wired to one port of a *multiport repeater* (MPR) typically located in a wiring closet. The MPRs can in turn be wired to higher-level MPRs. The resulting network topology is a distributed star, where each twisted-pair cable is allowed to be up to 100 meters long. Note that this is physically unlike the original multidrop Ethernet but can be centrally managed more easily and has higher reliability.

There are other twisted pair Ethernet architectures that predate those previously mentioned; AT&T's StarLAN 10 and LattisNet from SynOptics Communications are two examples. Most computers (and network analyzers) with an AUI interface can be used interchangeably on any of these networks with the appropriate transceiver.

Other Ethernets

Ethernet variations proliferate both with and without benefit of official standardization. Most share at least a philosophical tradition with the Ethernets discussed above but are often incompatible in the sense that equipment designed for one species cannot be connected with equipment designed for another. Among these are:

- Broadband Ethernets, some of which run at 5 Mbps so as to fit within a single television channel assignment (see "IBM PC Network™ (Broadband) Architecture" on page 1–31).
- Fiber-optic Ethernets, most of which use two cables per station and a centralized optical coupler to rebroadcast the transmitted signal to all receivers.

Access Control

Ethernet and its variants, StarLAN and PC Network, use an algorithm to control transmission called *carrier sense multiple access with collision detection* (CSMA / CD). A similar architecture is *carrier sense multiple access with collision avoidance* (CSMA / CA). Apple Computer implemented CSMA / CA in its LocalTalk[®] network using the LocalTalk Link Access Protocol or LLAP (see "LocalTalk[®] Network Architecture" on page 1–34).

Before transmitting, a station waits until it hears no other station transmitting. It defers until it senses no carrier from another station and then sends its message. Since there is the possibility that another station may decide to transmit at roughly the same time, the sending station monitors the network to hear if its own message appears on the network ungarbled.¹ If it detects a collision with another message, it intentionally sends a few additional bytes (a process called *jamming*) to ensure propagation of the collision throughout the system to all other transmitting stations. The station then remains silent for a randomly determined amount of time (controlled by a *backoff algorithm*) before attempting to transmit again.

Collisions may be heard by receiving stations as badly-formatted frame fragments called *runts* that are typically shorter than the minimum frame size² and have incorrect check sums. These frames may also have an incorrect starting sequence and may not be seen as the start of a frame at all. Note that the propagation time through Ethernet and StarLAN networks is typically much longer than the transmission time for a bit. Therefore, different receivers will see different effects depending on their position relative to the various transmitters (and, for StarLAN, to hubs).

In addition to deferring to active traffic, stations wait before beginning transmission after the end of another transmission. The pause varies according to network type: 9.6 microseconds for Ethernet, 96 microseconds for StarLAN, and 48 microseconds for PC Network. This enforced interframe spacing allows time for receivers to recover from one frame and to prepare to receive the next.

Other Transceiver Functions

Besides moving serial data to and from the attached device and the network, the transceiver (or MAU) has several other functions:

- **Collision detection:** One wire pair in the drop cable transmits the *collision detect signal* (called *signal quality error*, SQE, in IEEE documents) from the transceiver to the device. In addition to its use to control transmission, this signal can also be used to detect some network and transceiver failures. A transceiver that supports the *signal quality error test* feature will generate a collision detect signal at the end of a transmitted frame to demonstrate that the collision detect circuitry is at least partially operational.

1. For PC Network, the sending station monitors the network on the frequency to which the head end translates the traffic.
2. PC Network permits a smaller frame size than Ethernet and StarLAN. A PC Network data field may be as small as four bytes.

- **Jabber detection:** The transceiver is required to disable transmission if the device transmits for longer than the longest possible frame. This prevents some kinds of device failures from halting network activity but is not foolproof.

Format of a Frame

All data in a frame is transmitted as a sequence of 8-bit bytes starting with the least significant bit. The format of each frame is as follows:

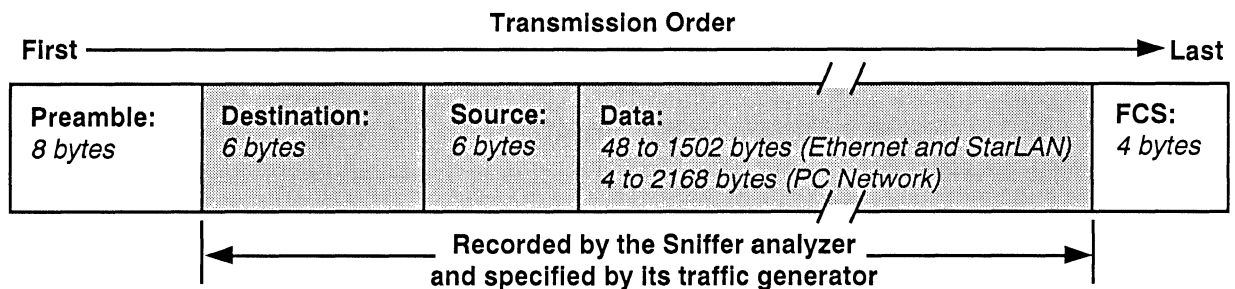


Figure 1–1. Frame format for Ethernet, StarLAN, and PC Network.

The **preamble** is a fixed data pattern used for receiver synchronization and recognition of the start of a frame. The **destination** and **source** addresses are each 6 bytes. Various kinds of multicast addresses are indicated if the first transmitted bit (the low-order bit of the first byte) of the destination address is a one. Note that the IEEE 802.3 standard also permits 2-byte source and destination addresses, but their use is rare and inconsistent with 10BASE5 and 1BASE5.

If the first transmitted bit of the source address is a one, it indicates that a *routing information* (RI) field is present before the data field. This is a convention inherited from token ring. Such frames are generally seen on an Ethernet only when forwarded over a MAC-level bridge from a token ring.

The **data** field must have a minimum number of bytes so that the duration of a frame is longer than the worst-case propagation delay (45 microseconds for Ethernet) through the network. For Ethernet and StarLAN, the minimum must be 48 bytes so that the frame is at least 60 bytes (excluding the **preamble** and the **FCS**). For PC Network, the minimum must be 4 bytes so that the frame is at least 28 bytes. If this were not true, then collisions might not be detected. The IEEE 802.3 standard specifies other limits on the frame size for non-10BASE5 and non-1BASE5 networks only.

Following the data is a 4-byte *frame check sequence* (**FCS**) that checks the validity of the source, destination, and data fields. The Sniffer analyzer does not record the **FCS** but will optionally collect and identify frames whose **FCS** is incorrect.

Format of the Data

To specify how the initial bytes of the data field are to be interpreted, there are several conventions in general use. The original format as defined by Xerox®, now often called the “Ethernet” format (in contrast to the IEEE 802.3 format),

contains no length field and begins with a 2-byte **Ethertype** field that indicates the major protocol type. For example, the IP protocol of TCP/IP is assigned the Etherbyte hex 0800.




Figure 1–2. Original “Ethernet” data format defined by Xerox.

The assignment of Etherbytes to protocols was originally done by Xerox and has now been taken over by the US Defense Communications Agency.

The second convention for interpretation of the data field is that supported by the IEEE/ANSI/ISO standards organizations for 802.3 networks and begins with a 2-byte **length** field (most significant byte first), followed by a *Logical Link Control* (LLC) header that conforms to the IEEE 802.2 standard.




Figure 1–3. IEEE 802.3 data format defined by IEEE/ANSI/ISO standards organizations.

Although both data formats can and do appear on the same network, there is no foolproof way to distinguish between the two by looking at the frame. In practice, though, almost all assigned Etherbytes are numerically larger than the maximum frame length of 1500. Thus, if the first two bytes of frame data are larger than 1500, it is probably an Ethernet/Etherbyte frame. (The only common exception is the PUP Etherbyte, hex 0200.)

The Sniffer analyzer decodes frames in either format and makes the format decision automatically unless instructed otherwise by the operator.

For frames of either type, the embedded **protocol data** represents information that is interpreted by higher level protocols and may contain many other such nested levels.

A third convention, the original format used by IBM® and Sytek, applies only to PC Network. It begins with a 2-byte **length** field that is then immediately followed by the **NetBIOS header**. There is no identification of the protocol type, so this format cannot easily be distinguished from other protocols on the network. Note that this is not the same NetBIOS format as is used on token ring networks or on PC Network using the IBM LAN Support Program.




Figure 1–4. The original data format used by IBM and Sytek (applies only to PC Network).

LLC Frames

A frame that conforms to the 802.2 standard is an LLC frame. LLC is a protocol that provides reliable connection-oriented virtual circuits or connectionless datagrams between processes. It is a subset of the International Standards Organization-defined superset, *High-level Data Link Control* (HDLC). See the subsection on HDLC and two other HDLC subsets, SDLC and LAPB, in the section, “WAN/Synchronous Architecture” on page 1–20.

The **protocol data** part of LLC frames (see Figure 1–3) begins with a 3 or 4-byte header, the first two bytes of which are the *destination service access point* (**DSAP**) and *source service access point* (**SSAP**). The SAP numbers are preassigned codes that indicate which subprotocol is used in the rest of the frame. For example, the NetBIOS protocol¹ has been allocated a single SAP (hex F0), and Systems Network Architecture (SNA) has been allocated four SAPs (hex 04, 05, 08, and 0C). The **SSAP** often equals the **DSAP**, except in frames that are establishing an initial SNA connection.

The **control** field defines the exact function of LLC frames. There are three LLC frame types:

Information

I format frames are used to send arbitrary sequenced data interpreted by the protocol that the SAPs designate. The LLC header contains a *send* sequence number N(S) for this frame and a *receive* sequence number N(R) of the next frame expected from the other station.

Supervisory

S format frames contain or consume an N(S) number but do not contain an N(R) number. In addition they can contain the following indications as either a command or a response:

1. This technique is also used on PC Network by the LAN Support Program.

Command/Response	Action
RR	<i>Receive Ready.</i> Transmission of Information frames can proceed.
RNR	<i>Receive Not Ready.</i> Transmission is temporarily blocked.
REJ	<i>Reject.</i> Retransmission starting with N(R) number is requested.

Figure 1–5. Indications used in LLC supervisory frames.

Unnumbered

U format frames contain neither N(S) nor N(R) numbers but may contain control information or data. The commands and responses in the unnumbered frame format are:

Command/Response	Action
SABME	<i>Set Asynchronous Balanced Mode (Extended).</i> Establish a virtual connection, also called a link.
DISC	<i>Disconnect.</i> Terminate a virtual connection.
DM	<i>Disconnected Mode.</i> The connection is broken.
UA	<i>Unnumbered Acknowledgment.</i> For SABME and DISC response.
FRMR	<i>FrameReject.</i> The format of a received frame was invalid. The protocol data field contains the reason.
XID	<i>Exchange Identification.</i> Used between two stations to exchange identification and the characteristics of the two stations.
TEST	<i>Test Probe.</i> Should be echoed by the receiving station.
UI	<i>Unnumbered Information.</i> Used by the SAP protocol for any purpose.

Figure 1–6. Commands and responses in LLC unnumbered frames.

The only LLC frames that are allowed to contain data after the LLC header are I, UI, TEST, XID, and FRMR types.

There are three types of LLC operation. The first is known as *unacknowledged connectionless service*. This is the minimal use of LLC in which every frame is a UI type. Thus, the data part of every frame viewed on a Sniffer analyzer begins with three bytes: the two SAP bytes and a UI (hex 03) control byte. This technique is typically used to implement higher level protocols that do connection control and sequencing themselves, and therefore do not need the services of the standard LLC, but that prefer compatibility with LLC formats.

The other two types carry more overhead. *Connection-oriented service* provides flow control, sequencing, and error recovery. *Acknowledged connectionless service* is also connectionless like the first type but provides for acknowledgment and relieves higher layers of that responsibility.

Assignment of Network Addresses

The modern convention for 6-byte node addresses is to divide them into an initial 3-byte code representing the manufacturer of the equipment, and a unique 3-byte serial or sequence number assigned to that piece of equipment. The responsibility to assign manufacturer codes was initially taken by Xerox but has now been assumed by the IEEE.

It is the intention of the IEEE that the same code be used for all networks supported by a manufacturer, including Ethernet/StarLAN (IEEE 802.3), PC Network, and token ring (IEEE 802.5). However, Ethernet/StarLAN/PC Network bytes are transmitted with the least significant bit first, and token ring bytes are transmitted with the most significant bit first. This has led to considerable confusion in the way that assigned addresses are used.

The IEEE's stated position is that the assigned code specifies how the address should appear on the network cable. The result is that a network address stored in a computer's memory will be different, depending on what the originating network was. Since network addresses appear at various places within the protocol levels of a frame, and since a frame may have been forwarded by various network types, this imposes a difficult burden on network software. In fact, observation of network software and hardware from a variety of vendors shows that some have chosen to use the same address as it appears in memory on both networks, thus using a code on one or the other network that, according to the IEEE interpretation, is not assigned to them.

Sytek/IBM 6-Byte Address

The original Sytek/IBM PC Network adapters used an entirely different convention for node addresses: 4 bytes of unique serial, followed by a 2-byte manufacturer's code. The only codes ever assigned were 0000 (for IBM) and 0100 (for Sytek).

Token Ring Network Architecture

The token ring is a type of LAN suitable for high-speed interconnection of computers and computer-controlled devices over moderate distances. The token ring architecture is defined by implementations from IBM, Texas Instruments, and others as ANSI/IEEE standard 802.5 and ISO/DIS standard 8802/5.

Most system implementations of the token ring network also use at least a subset of a similarly standardized protocol for LLC and defined as ANSI/IEEE standard 802.2 and ISO/DIS standard 8802/2.

Physical Interconnection and Speed

Stations connected to the token ring network are wired together physically in star-like fashion. Each station uses one cable to attach itself to a nearby *passive concentrator*, or *multiple access unit* (MAU).¹ The MAUs can be linked together and may be separated by moderate distances. The number of stations and the distance limitations depend on several variables, including cable type, but typically one or two hundred stations can be interconnected into a single network segment using cables between stations and MAUs up to 300 meters long. The distance limitations can be overcome by using special line drivers or fiber optic cables. Networks of many hundreds or thousands of stations over large distances can be created using bridges between network segments.

One type of connector used to attach to MAUs was designed by IBM and called the IBM Data Connector. They are hermaphroditic, so that any two may be joined; cable "extension cords" have the same connector on both ends. The cable that connects to a particular computer may use the hermaphroditic connector if there is enough panel space for the mating connector (about one inch by one inch) or may use a non-standard connector. The convention for personal computers is to use a DB-9 female connector on the backpanel and DB-9 male on a cable whose other end has the hermaphroditic connector. Other vendors have built MAUs with RJ-11 connectors that use less panel space.

There are two basic speeds for the network: 4 Mbps, or 500,000 bytes per second, and 16 Mbps, or 2,000,000 bytes per second. This does not include the many levels of overhead in a typical application, and throughput for the user will often be many times less than that. There is nothing in the hardware or software architecture that limits the network to 16 Mbps.

Stations operating at 4 and 16 Mbps may not be attached to the same token ring at the same time. The first station to become operational on the ring establishes the speed, and any subsequent station entering the ring at a different speed will cause transmission errors and will often result in other stations removing themselves from the ring. Unfortunately, current ring adapters have no provision for determining the speed of an established ring before attempting to join it.

1. Do not confuse this use of the MAU acronym with Medium Attachment Unit (MAU) defined in IEEE Standards documents for Ethernet.

Logical Interconnection

Each interconnection cable contains two twisted pairs of wires. Although the stations and MAUs are cabled in a star-like fashion, the electrical effect of the token ring cables and connectors is to create a continuous ring from station to station. One twisted pair in the cable to each station is used to transmit to the next station in the ring, and the other pair is used to receive from the previous station. The ordering depends on how cables are plugged into the MAUs and how the MAUs are interconnected.

The operation of the ring depends on each station retransmitting data from its *receive pair* onto its *transmit pair*, regardless of whether that station is involved in the conversation. To insure that the ring is operational even when some stations are turned off, connecting a cable from a station to a MAU is not sufficient to cause that station to enter the ring; it also must send a DC voltage on its transmit pair to trigger a relay in the MAU. If power to the station fails, or if the cable to the station is disconnected at either end, the relay loses power and the ring bypasses that station.

When the relay is not powered, the cable to the station has its transmit pair connected to its own receive pair so that it may test the network adapter and cable. Prior to inserting itself onto the ring by supplying the relay voltage, the network adapter sends several thousand data frames to itself to verify correct operation. This process, plus the network adapter self-test, may take 15 seconds or more.

Access Control

Only one station on the entire ring is allowed to transmit data at a time. To control access, a 3-byte message giving permission to transmit, called the *free token*, continually circulates when there is no other traffic. Each idle station retransmits it as it is received. A station that wants to transmit data waits for the token and then sends its data instead of the token. When its data transmission is finished, it regenerates the token message. In addition to this simple rotational priority scheme, there are also ways to establish other priorities. Every message (including the token) contains both a 3-bit priority field for itself and a 3-bit reservation priority for a possible subsequent message.

A data message, called a *frame*, may be directed to a single destination station or to any of various groups of stations. In all cases, the addressee (the station that receives the message) does *not* remove it from the ring. It simply makes a local copy of the message and retransmits it. The originating station is responsible for removing the message from the ring when the message returns to it. The originator then replaces the message with the token.

In visualizing the traffic flow on the network, it is important to realize that most frames are much longer (in time) than the round-trip delay around the ring, particularly at 4 Mbps. Each station introduces a delay of less than 3 bit times when it is repeating data from its receive cable to its transmit cable, whether or not it is making a copy of the data. That delay for each station, plus the cable propagation delays, produce the total ring round-trip time. For a typical network of 50 stations, the round-trip time might be about 50 microseconds. A

1000-byte (8000 bit) frame takes 2000 microseconds to transmit at 4 Mbps, so the transmitter must be removing the beginning of the frame that has made the trip around the ring long before it has finished sending out the whole frame.

At 16 Mbps, small frames will take less time to transmit than the round-trip token timing, especially for large networks. To improve performance, the newer token ring adapters implement an *early token release algorithm* that allows them to transmit the free token to the next station before they have received a frame just transmitted.

In normal operation, the token is circulated and regenerated by the cooperative operation of all stations acting democratically. If the token is destroyed by transmission error or other fault (a station attaching or removing from the ring typically destroys the token because of electrical noise created by the relay operation) it is the responsibility of a station designated as the *active monitor* to notice the absence of the token and to regenerate it. There is only one active monitor on the network at a time, although every station is able to assume the role if needed.¹ If the active monitor is disabled or leaves the ring, a monitor contention process begins through which a new active monitor is elected by the remaining stations.

Format of a Token Frame

All data is transmitted as a sequence of 8-bit bytes sent serially and Manchester encoded. The minimum transmission is the 3-byte *token*.

SDEL 1 byte	AC 1 byte	EDEL 1 byte
-----------------------	---------------------	-----------------------

Figure 1–7. Frame format for a token frame.

Both *starting delimiter (SDEL)* and *ending delimiter (EDEL)* have intentional Manchester code violations in certain bit positions so that the start and end of a frame can never be accidentally recognized in the middle of other data. The *access control (AC)* byte contains a bit that indicates that this is a token, not a data frame, and contains priority information. Tokens are not recorded by the Sniffer analyzer.

1. The Sniffer analyzer cannot play the role of active monitor when in capture mode.

Format of a Data Frame

If a message is not a token, then it is a *data frame*.

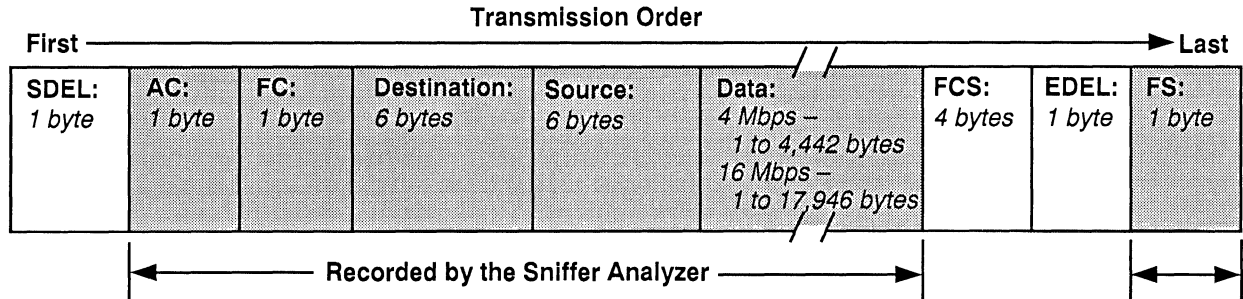


Figure 1–8. Frame format for a token ring data frame.

The **SDEL**, **AC**, and **EDEL** fields are as before, except the **AC** byte now says that this is a data frame and not a token. The **FC** byte contains frame information. The **destination** and **source** addresses are each 6 bytes, and various kinds of *multicast* or *broadcast* addresses are indicated if the first bit of the destination address is a one. Following the **data** field is a *frame check sequence (FCS)* that checks the validity of all previous data starting with the **AC** byte. The Sniffer analyzer records bytes starting with **AC** and ending with the last byte of the data field.

The last byte for data frames is the *frame status (FS)* byte containing bits that may be set on by the recipient of the frame: *address recognized*, if a station matched the destination address and *frame copied* if it was able to successfully make a local copy of the data as it passed by. Note that the **FS** is not covered by the frame check sequence (so that the **FCS** doesn't have to be changed by the recipient), but for greater reliability, the bits in the **FS** are each duplicated. The Sniffer analyzer records the **FS** byte and displays it in the detail view.

Optional Routing Information

Token ring uses a technique known as *source routing* when it does routing. Thus, there is an optional *routing information (RI)* field that may be present at the beginning of the data part of any frame. The **RI** field, which may be up to 32 bytes long, contains information about the path that the frame took if it was forwarded through multiple network segments by bridges. If the **RI** field is present, the first bit of the source address will be a one.

The **RI** field is not part of the IEEE 802.5 token ring standard, although it has been proposed for official adoption. IBM software currently uses this extension to the standard.

MAC Frames

A data frame may be a *medium access control* (MAC) frame that contains information used to control the token ring network itself. Most MAC frames are generated and processed by the computer's network adapter and are not of concern to software within the host. The type field in the FC byte indicates whether the frame is a MAC frame.

MAC frames are used for processes like *monitor contention*, *error reporting*, and *error recovery*. In addition, the *active monitor* announces its presence with a periodic MAC frame, and all stations that could become the monitor (for example, *standby monitors*), should the need arise, do likewise.

MAC frames contain a major type code followed by a variable number of variable-length fields called *subvectors* that give additional information.

LLC Frames

A data frame that is not a MAC frame is (in all non-proprietary uses of the token ring network) an LLC frame. See the discussion of the LLC frame format in "Format of a Frame" on page 1-7. Since all non-MAC frames are supposed to use 802.2 LLC, an alternative mechanism has been standardized for encapsulating the older Ethertype formats. It is known as the Sub-Network Access Protocol (SNAP), and you can see the frame format in Figure 1-9.




Figure 1-9. LLC frame format for Sub-Network Access Protocol (SNAP).

SNAP has been assigned SAP hex AA, and the data field following the **Control** field is a 5-byte *protocol identification header*. The **Agency code** is an assigned manufacturer code, but sometime 0 is used. The **Local code** is typically the Ethertype. Ethernet-style data follows the header.

Assignment of Network Addresses

A discussion of the assignment of network addresses can be found in the section, "Ethernet Network Architecture."

Fiber Distributed Data Interface Network Architecture

Fiber Distributed Data Interface (FDDI) is a LAN standard suitable for high-speed interconnection of multivendor LANs over moderate to long distances. The FDDI architecture is defined by implementations of the ANSI/ISO X3T9 standards.

The FDDI standards define a version of the Physical and Data Link layers of the Open Systems Interconnection (OSI) model. The FDDI standards consist of the following:

- Physical Layer Medium Dependent (PMD) defines the medium requirements, the connectors that attach the nodes to the medium, and the fiber-optic transceivers.
- Physical Layer Protocol (PHY) defines the symbols, line states, data framing, clocking requirements, and encoding/decoding techniques.
- Media Access Control (MAC) defines the data link packets required for token passing, frame formatting and checking, data link addressing, error detection and recovery, and bandwidth allocation among nodes.
- Station Management (SMT) defines how to manage the PMD, PHY, and MAC portions of FDDI. SMT facilities include connection management, node configuration, error recovery, and formatting of SMT frames. SMT also includes a Management Information Base (MIB).

Figure 1–10 shows the relationship among the FDDI standards and how they combine to form the FDDI specification. Although Logical Link Control (LLC) is not part of the FDDI standards, FDDI requires LLC for proper operation and data transmission.

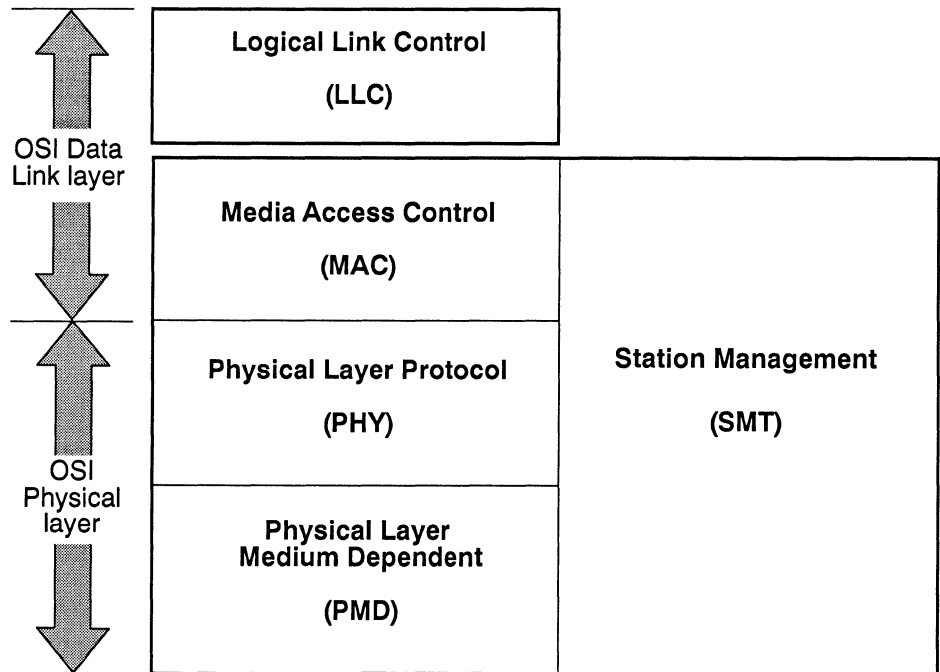


Figure 1–10. Relationship of FDDI standards.

Interconnection and Speed

Physically, an FDDI network can have a ring, star, or tree topology. An FDDI network can consist of a concentrator with attached stations, a tree of concentrators, a dual counter-rotating ring, or a dual ring of trees. The dual counter-rotating ring provides a fault-tolerant structure that ensures data flow in case of a cable or station failure on the ring.

Stations attached to an FDDI network are physically connected by a media interface connector (MIC). A MIC aligns the fiber with the transmit/receive optics in the station. The connector consists of a keyed plug and a keyed receptacle.

Logically, stations are connected to Media Access Control (MAC) entities. Each MAC has a unique data link address. The MACs are interconnected in a token ring configuration -- each MAC receives tokens and data from its upstream neighbor and transmits tokens and data to its downstream neighbor. A station can have multiple instances of PMD, PHY, and MAC entities, but only one instance of SMT.

FDDI operates at 100 MBps over a fiber-optic cable. The ANSI standard defines the fiber medium as 62.5/125 micrometer, multimode, graded-index, fiber optic cable.

Access Control

The FDDI standard specifies a timed-token passing method for controlling access to the ring. Stations compete for the right to initialize the ring through a mechanism known as the claim process. During the claim process each station continually transmits "claim frames" containing that station's required target token rotation time (TTRT). If a station receives a claim frame with a lower TTRT than its own, it stops transmitting its own claim frames and simply repeats the claim frames that it receives. When a station receives its own claim frame, it surmises that its value of TTRT is the lowest one and that, in accordance with the FDDI rules, it can therefore initialize the ring.

To initialize the ring, the station issues a token. The token passes around the ring without being captured by any station. After the token's third rotation, the ring is in a steady state operation and stations can begin transmitting data.

In normal operation, the token is circulated and regenerated by all active stations on the ring. If a station wants to transmit data, it waits until it detects the token. It then captures the token and stops the token passing process. The station can then send data frames until it has finished sending data or until its token holding time (THT) expires. Then the station releases the token onto the ring for use by another station.

Format of a Token Frame

The token format consists of four fields, as shown in Figure 1-11. The starting delimiter indicates the initial boundary of a transmission sequence. The frame control field defines the frame type, and the ending delimiter indicates the end of the transmission.




Figure 1-11. Token frame format for FDDI.

Format of a Data Frame

The FDDI frame format is similar to the IEEE 802 frame format, and is shown in Figure 1-12. The frame starts with a preamble and a unique starting delimiter octet, followed by the frame control (FC) octet. The FC octet determines whether the frame is an SMT frame, a MAC frame, a synchronous LLC frame, or an asynchronous LLC frame. Following the data field is a frame check sequence that checks the validity of all previous data, and an ending delimiter. The last octet for data frames is the frame status byte, containing bits that may be set by active stations on the ring.




Figure 1-12. Data frame format for FDDI.

Station Management

The Station Management (SMT) standard defines the station level controls necessary to manage a station on an FDDI network. SMT contains these components:

- Connection management (CMT) connects nodes to the FDDI network, manages physical connection between adjacent PHYs, configures PHY and MAC entities within a node, and coordinates trace functions.
- Ring management (RMT) receives status information from MAC and CMT and reports this status to SMT and higher-level processes.
- SMT frame services provide the means to control and observe the FDDI network.

See the discussion of the SMT protocol interpreter suite in Chapter 2 for more details concerning station management.

WAN/Synchronous Architecture

Synchronous data communications are used locally (as, for example, on a simple RS-232C link between computers in the same building) and also for communication over wide area networks (WANs). WANs in turn are being used increasingly in support of internetworking — connecting local area networks (LANs) together by means of WAN links. Most internetworking is currently done over leased lines using some variation of HDLC-like protocols. The Frame Relay protocol (described in Chapter 2) is frequently used to forward LAN frames over WAN links. Some implementations use the X.25 protocol (also described in Chapter 2) over HDLC.

At the network level, X.25 defines a general purpose interface between data terminal equipment (DTE) and data circuit-terminating (or communications) equipment (DCE). A DTE is an end-user machine. The entry point to the WAN is a DCE.

At the logical link level, the International Organization for Standardization (ISO) published the widely used standard, High-level Data Link Control (HDLC) protocol. There are several important protocol subsets of the HDLC superset. Link Access Procedure Balanced (LAPB) supports the widely accepted X.25 packet network protocol. Synchronous Data Link Control (SDLC) is IBM's version of HDLC and is used for its Systems Network Architecture (SNA) protocol. Finally, LLC is the standard released by the IEEE 802 standards committee for LANs (See the discussion of "LLC Frames" in the section on "Ethernet® Network Architecture" on page 1–3).

Several standards have been implemented at the lower layers of the synchronous architecture and allow users a wide range of options to interface to a packet-switched network. Physical level standards include the Electrical Industries Association's familiar RS-232C. The International Consultative Committee for Telephony and Telegraphy (CCITT) defined two well-known sets of standards, the V-Series (V.35, for example) and the X-Series (X.21, for example). Several of these other standards are related to RS-232C: CCITT V.24 and V.28, CCITT X.21bis, and ISO 2110.

Interconnection and Speed

Typically, RS-232C specifies a 25-pin connector, so that up to 25 wires can be used to connect two devices. The electrical characteristics of RS-232C place a limit of about 50 feet on the distance between, for example, a DTE and its modem. These same characteristics limit the data transmission rate across the interface to a maximum of about 19.2 Kbits/s.

V.35 allows data transmission at higher speeds (56 Kbits/s to 1.54 Mbits/s in the US; 64Kbits/s to 2.048 Mbits/s internationally). It is implemented using both the frequency modulation and amplitude modulation techniques. The Sniffer analyzer connection to V.35 is baseband.

Format of a Frame

Frames captured from a synchronous line and interpreted by the Sniffer Internetwork Analyzer generally conform to the HDLC standard defined by ISO. In particular, two subsets of HDLC are relevant to Sniffer Internetwork Analyzer users — SDLC and LAPB. The Sniffer analyzer will also recognize proprietary versions of HDLC framing implemented by Wellfleet, Cisco, Vitalink, Proteon, IBM, and Point-to-Point (PPP), as well as Frame Relay formats conforming to the ANSI standard or to RFC 1294.

SDLC is IBM's version of the HDLC superset. It is the logical link layer protocol for IBM's network layer services of its SNA protocol. LAPB is the link layer protocol for the network layer protocol X.25. Sniffer menus and screens refer to LAPB as HDLC.

One major difference between LAPB and SDLC involves the stations' responsibilities and the rules that the stations follow when transferring information. LAPB stations, like LLC stations, allow any station to initiate transmissions without prior permission from any other station (referred to as *Asynchronous Balanced Mode* or ABM). SDLC, on the other hand, requires a subservient *secondary* station to receive explicit permission from a dominant *primary* station before transmitting (referred to as *Normal Response Mode* or NRM). Unlike the *balanced* configuration of LAPB stations where stations are equally responsible for the link, the SDLC station configuration is *unbalanced*—the prime responsibility for the link depends upon the primary station.

Both LAPB and SDLC conform to the general HDLC frame format. The format of each frame is as follows:




Figure 1-13. The general HDLC (including LAPB and SDLC) frame format.

The **flag** is a special bit sequence, 01111110, that indicates either the beginning or the end of a frame. A technique known as bit stuffing is used to ensure that occurrences of this bit pattern in the frame data are not misinterpreted as flags. The transmitting machine checks between the opening and closing flags, and it inserts a 0 bit when it encounters five continuous 1 bits. After the frame has been "stuffed" and flags placed, the transmitting machine sends the frame to the receiver. The receiver monitors the bit stream. Whenever the receiver encounters 6 consecutive 1 bits, it knows that it has received a flag. Otherwise, when it sees 5 consecutive 1 bits followed by a 0, it removes that 0.

The **address** field identifies the primary or the secondary station transmitting a particular frame. Each station has a unique address. In unbalanced

configurations, address fields contain addresses of secondary stations in both commands and responses. In balanced configurations, command frames contain the destination address, and response frames contain the transmitting station address.

The **control** field defines the exact function of LAPB and SDLC frames using the general HDLC frame format. It performs various functions, depending upon whether a frame is an *Information*, *Supervisory*, or *Unnumbered* frame (see below).

There are two important differences between LAPB and SDLC frames with regard to the control field:

- One is that each type of configuration has its own mode-setting command to establish link-level contact as either balanced or unbalanced.
- The other difference is that in the balanced configuration of LAPB, all stations can send commands and responses. In the unbalanced configuration of SDLC, a frame sent by a primary station is a command, and a frame sent by a secondary station is a response.

There are three HDLC frame types:

Information

I frames are used to convey or acknowledge end-user data between two devices. The control field of this type of frame contains both a *send sequence number* N(S) for itself and a *receive sequence number* N(R) for the next frame expected from the other station.

Supervisory

S frames are used to control the flow of data. The control field of this frame type contains an N(R) number but not an N(S) number. In addition, the control field contains the indications listed in Figure 1–14 as either a command or a response.

Command/Response	Action
RR	<i>Receive Ready.</i> I frame transmission and/or acknowledgement can proceed.
RNR	<i>Receive Not Ready.</i> I frame transmission and acknowledgement are blocked.
REJ	<i>Reject.</i> Retransmission starting with N(R) number is requested.

Figure 1–14. Commands and responses used in HDLC (SDLC and LAPB) supervisory frames.

Unnumbered

U frames are also used for control purposes. The control field of this frame type contains neither N(S) nor N(R) numbers but may contain control information or data. The commands and responses appearing in the unnumbered frame control field are listed in Figure 1–15.

Command/Response	Action
SNRM	<i>Set Normal Response Mode.</i> Configure a secondary station in a mode that precludes it from sending unsolicited frames. The primary station controls all message flow. Used in SDLC.
SNRME	<i>Set Normal Response Mode (Extended).</i> Same as SNRM except that it selects modulo 128 sequence numbering rather than modulo 8 sequence numbering.
SABM	<i>Set Asynchronous Balanced Mode.</i> Configure stations as peers with each other. Used in LAPB.
SABME	<i>Set Asynchronous Balanced Mode (Extended).</i> Same as SABM except that it selects modulo 128 sequence numbering rather than modulo 8 sequence numbering.
DISC	<i>Disconnect.</i> Command used to place a secondary station in the disconnected mode (not operational).
DM	<i>Disconnected Mode.</i> Used by secondary station to indicate that it is in the disconnected mode.
UA	<i>Unnumbered Acknowledgment.</i> Response to set mode commands and to DISC.
FRMR	<i>Frame Reject.</i> The format of a received frame was invalid. The information field contains the reason.
XID	<i>Exchange Identification.</i> Used between two stations to exchange identification and the characteristics of the two stations.
UI	<i>Unnumbered Information.</i> Used for transmission of user data in an unsequenced frame.

Figure 1–15. Commands and responses used in HDLC (SDLC and LAPB) unnumbered frames.

The only HDLC frames that are allowed to contain data after the header are I, UI, TEST, XID, and FRMR types.

Phases of Link Control

The following five phases represent the fundamental sets of activities on a synchronous line using HDLC (SDLC and LAPB):

Connect phase

Establishes a connection over a switched facility. The process includes off-hook signalling, switching, and exchange of identification.

Link establishment phase

Typical processes in this phase include initializing data transfer (over an already established physical link) and polling.

For example, when a DCE is able to proceed with a connection, it sends DISC frames with the *Poll* bit set on. When a DTE wants to reconnect to the DCE, it waits until it receives a DISC frame and then responds with a UA frame with the *Final* bit set on. The DCE is responsible for setting up the link when it receives the UA frame within the required amount of time. The specific link setup command depends upon the HDLC subset in use: if it is LAPB, the commands are SABM and SABME; if it is SDLC, the commands are SNRM and SNRME. After the DCE receives the link setup command in the required amount of time, it responds with a UA frame, and the link is now in an “up” state.

Information transfer phase

The information transfer phase follows link establishment. It includes the acknowledgment process as well as the actual data transfer between connected stations.

For example, the DTE sends an I frame. The DCE may acknowledge with either an RR frame or another I frame (if it too has data to send), and it may delay acknowledgement. When the DCE does acknowledge, it sends an N(R) number that is inclusive of all traffic transmitted and accepted.

Termination phase

The termination phase relinquishes control of the link following transmission of data. In an unbalanced configuration, the secondary station returns control to the primary station.

Clear phase

The clear phase releases the facility. For example, the DTE sends a DISC frame to tear down a link. The DCE responds by transmitting a UA frame.

StarLAN™ Network Architecture

The original 1 Mbps StarLAN is a type of LAN suitable for moderate-speed interconnection of computers and computer-controlled devices over moderate distances. The StarLAN architecture is defined by implementations from several manufacturers as a proposed variation of ANSI/IEEE standard 802.3, ISO/DIS standard 8802/3, and FIPS standard 107. StarLAN is a slower speed variant of the Ethernet network (see “Ethernet® Network Architecture” on page 1–3).

Physical Interconnection and Speed

Stations connected to a StarLAN network are connected so that every station hears what any station transmits. The delay between transmission and reception depends only on the propagation delays through the wires and attaching devices. In this way StarLAN is similar to networks like Ethernet and IBM PC Network (see “StarLAN™ Network Architecture” on page 1–25 and “IBM PC Network™ (Broadband) Architecture” on page 1–31). It differs fundamentally, however, from networks like the IEEE 802.5 token ring, where stations are wired in a logical ring so that each station hears only what its upstream neighbor transmits (see “Token Ring Network Architecture” on page 1–12).

The original StarLAN transmission speed is 1 Mbps, or 125,000 bytes per second, sent in baseband (non-modulated, non-RF) form. The ANSI/IEEE documents refer to this as the *1BASE5* (1 Mbps, baseband, and 500 meters per segment) standard. Another common standard is *10BASE5* (10 Mbps, baseband, and 500 meters per segment), for Ethernet. Because of the access control mechanism described below, the effective network throughput is significantly less than the transmission speed. StarLAN 10 is a newer version, and it operates at 10 Mbps.

Although the StarLAN network is logically a bus because every station hears what every other station transmits, it is physically wired as a hierarchical branching star. Each station is wired to a repeater unit variously called a *network extension unit* or *hub*. Each hub can accept connections from 7 to 11 stations or other hubs, and also has an output signal that is used to connect to a hub at the next higher level in the hierarchy, if any. Hubs can be layered up to five deep. The top hub is the *header* or *master* hub of the network.

The cable used to connect stations to hubs and between hubs contains two twisted pairs of voice-grade 24 AWG or heavier wire. One pair is for transmitted data and one for received data. The standard connectors are four-pair RJ-45 telephone plugs and RJ-18 jacks with two of the pairs unused. Note that these are larger than the three-pair connectors used for most phones.

The maximum distance between a station and a hub, or between hubs, is 250 meters. Thus the maximum cable distance from any station to the master hub is five times that, or 1250 meters. The maximum number of stations in any single network is 1024.

Some manufacturers specify or recommend limits that are more restrictive than these. In addition, some permit up to 10 stations to be connected together in a chain without any hubs if the maximum cable distance of the chain is less than 125 meters.

Access Control

StarLAN networks, like Ethernet and PC Network, use the CSMA/CD algorithm to control transmission. For a discussion of access control see "Ethernet® Network Architecture" on page 1-3.

Format of a Frame

The frame format used in StarLAN networks is identical to that used in Ethernet. See the discussion of frame format in "Ethernet® Network Architecture" on page 1-3.

Format of the Data

The data format used in StarLAN networks is identical to that used in Ethernet. See the discussion of data format in "Ethernet® Network Architecture" on page 1-3.

Assignment of Network Addresses

The assignment of network addresses in StarLAN is similar to that for Ethernet. See the discussion on the assignment of network addresses in "Ethernet® Network Architecture" on page 1-3.

ARCNET® Network Architecture

ARCNET is a LAN suitable for moderate-speed interconnection of computers and computer-controlled devices over moderate to large distances. The ARCNET architecture was originally defined by products available from Datapoint in 1977. It was probably the first commercially available LAN. Implementations have been simplified and standardized since 1982 by the availability of an ARCNET LAN controller and associated devices from Standard Microsystems Corporation. No standardization exists for ARCNET.

Physical Interconnection and Speed

ARCNET uses a token-bus topology and is connected so that every station hears what any station transmits. The delay between transmission and reception depends only on the propagation delays through the wires and attaching devices. In this way ARCNET differs fundamentally from a network like the IEEE 802.5 token ring where stations are wired in a logical ring so that each station only hears what its upstream neighbor transmits.

The ARCNET transmission speed is 2.5 megabits per second, sent in baseband (non-modulated, non-RF) form. Each data byte is preceded by a 3-bit starting sequence, so that it takes 11 bits for each byte. The useful data rate is, therefore, 8/11 of 2.5 Mbps, which is 1.8 Mbps or 227,000 bytes per second.

Although the ARCNET network is logically a bus because every station hears what every other station transmits, it is physically wired as interconnected stars. Each station is wired to a repeater unit called a *hub*. Each hub can accept connections from 6 to 16 stations or other hubs. The topology is unconstrained except that there can be no loops, and no two stations may be separated by more than ten hubs.

The cable used to connect stations to hubs and between hubs is RG-62, 93-ohm coaxial cable. Since each cable connects to a hub on one end and a station or hub on the other end, the terminators are typically built into the adapters. No separate terminating devices are needed.

The maximum distance between a station and a hub, or between hubs, is 610 meters. Since stations can be up to ten hubs apart, the largest “diameter” of an ARCNET network is, thus, 6,710 meters.

There are two optional techniques for reducing the number of active hub ports needed to connect devices to an ARCNET network:

- One is by the use of a *passive* mini-hub, an inexpensive resistive coupler with three or four ports. Passive hubs function logically the same as the more expensive *active* hubs, but may not be connected directly to other passive hubs and may attach to cables no longer than about 61 meters. Passive hubs are commonly used when a single cable from an active hub needs to connect to two or three computers in a single room.
- More recently, a *multi-drop* or *high impedance* version of the network interface transceiver has been developed. With this technique a series of computers can be connected to a single cable from a hub in daisy-chain fashion, and a single terminator is present at the end of the cable. The cable length is typically limited to 305 meters and the number of devices to ten.

Logical Interconnection and Access Control

Although ARCNET networks are wired in an almost arbitrary fashion, stations (not hubs) in the network are logically considered to be connected in a ring. Each of the stations on a single network is assigned an 8-bit address, typically by setting mechanical switches on the adapter card, but sometimes by the software driver. The logical ring consists of stations considered in the numerical order of their station addresses, which has nothing to do with their physical locations or the manner in which they are cabled to the network.

When there is no data traffic on the network, there is a continual transmission of the *token message* (see the first message in Figure 1–16), also called the *invitation to transmit message*. The station that has just received the token may transmit a single data frame to any other station on the network. After that data frame is complete (or immediately after it received the token, if it has no data to

transmit), the station must then send the token to the station on the network that has the next higher station address. It takes about 30 microseconds (plus cable propagation delays) for a station that has nothing to transmit to send the token to the next station.

Remember that all stations can hear what any station transmits, but the token-passing scheme ensures that only one station is transmitting at a time. Each station listens for two kinds of transmissions:

- A *data message* (see the fifth message in Figure 1–16) being sent to it from any station on the network that just got the token and wants to send it data.
- The token sent to it from the station that has the next lower station address.

The logic within the hubs depends on the fact that there is only one transmitter at a time in order to be able to allow two-way communication using a single cable. In the idle state, a hub listens for incoming signals on all its ports. As soon as it detects the start of a packet (token or data) on any port, it switches all the other ports from input to output and retransmits the incoming packet out on all of the other ports. When transmission is complete, it again enters the idle state to wait for another packet from any port.

Some newer ARCNET adapters violate the original transmission rules in order to increase performance. The most common variation, a *nodal priority scheme*, allows a station that has first transmitted to send again (perhaps by sending a token to itself) without waiting for the token to complete another round trip through the network.

Establishing and Maintaining the Token Sequence

During normal operation, the only information each station knows about the rest of the network is the address of the next higher-addressed station, to which it must send the token. But how is the token passing started (or restarted after an error), and how does each station discover its neighbor's address after a start or a restart?

The process by which the token rotation is established is called a *reconfiguration*. Whenever there is no network activity for more than 78 microseconds, all stations assume that the token doesn't exist, and each starts an internal countdown based on the value of its own station address. The first station to count down to zero is the one that generates the initial token. It is first sent to an address one higher than the station's own address in the sequence 0, 1, 2, ..., 254, 255, 0, 1, 2... If there is no activity after 74 microseconds, the conclusion is that there is no station at that address, and the next higher address is tried. This continues until a station is found, at which point the first station remembers that address as its next higher neighbor, and the neighboring station starts the search for its next higher neighbor. These searches continue until the station who originated the token receives it from its lower-addressed neighbor. At that point each station knows its higher neighbor, and normal token passing and data traffic can continue.

The entire reconfiguration process is quite fast. It can take as little as 24 milliseconds and never more than 61 milliseconds. It can occur anytime that the token has been destroyed by an error, but it normally occurs only when a new station joins the network. In that case, there is a token circulating, but it is never sent to the new station. When any station detects this situation (more than 840 milliseconds of network traffic during which it never received the token), it forces a reconfiguration by sending, without having received the token, a burst of data longer than any possible packet plus the subsequent token. That guarantees that the token will be destroyed, and then the reconfiguration algorithm will start and assign the token rotation sequence to include the new station.

Note that a full reconfiguration need not occur when a station leaves the network. In that case, the station that passes the token to the now-absent station will notice that neither token nor data is transmitted by that station within 74 microseconds after passing the token to it, so it will start sending the token to higher addresses until it finds the absent station's next neighbor and makes it its own.

When the hardware is working and is configured correctly, the algorithms for the maintenance of ARCNET token passing work extremely well, even when changes are being made to a live network. When two stations have the same address, however, or if any of several hardware faults occur (such as a station whose receiver is defective but whose transmitter and controller are working), the network will reconfigure excessively. An effective technique for isolating the problem is to break the connection between hubs; the subnetwork that contains the fault will continue to reconfigure forever, but the other will become a correctly-operating network.

Transmitting Data Frames

One advantage that ARCNET has over most other LAN designs is that it avoids sending data on the network that cannot be accepted by the destination station. In the discussion above, the operation of sending a data frame actually consists of several steps. First, the station that just received the token transmits a *free buffer inquiry message* (see the second message in Figure 1-16) to the station to which it wishes to transmit data. There are several possible outcomes:

- If the receiver is present on the network and has a free buffer for the message, it sends back an *acknowledgment message* (see the third message in Figure 1-16), and then the data is sent.
- If the receiver exists on the network but has no buffer, it sends a *negative acknowledgment message* (see the fourth message in Figure 1-16), and the transmitter postpones sending the data until the next time it gets the token.
- If the station doesn't exist at all, the transmitter will hear no activity for 74 microseconds, then abandon the transmission, and pass on the token.

A second advantage of ARCNET not shared by most other networks is that there is an immediate confirmation at the data-link level that the message was

correctly received: the destination sends another *acknowledgment message* following the data. If the data was garbled, it sends nothing. After 74 microseconds with no response from the intended receiver, the transmitter knows that the message was not received correctly.

Format of a Frame

There are five messages that may appear on the ARCNET.

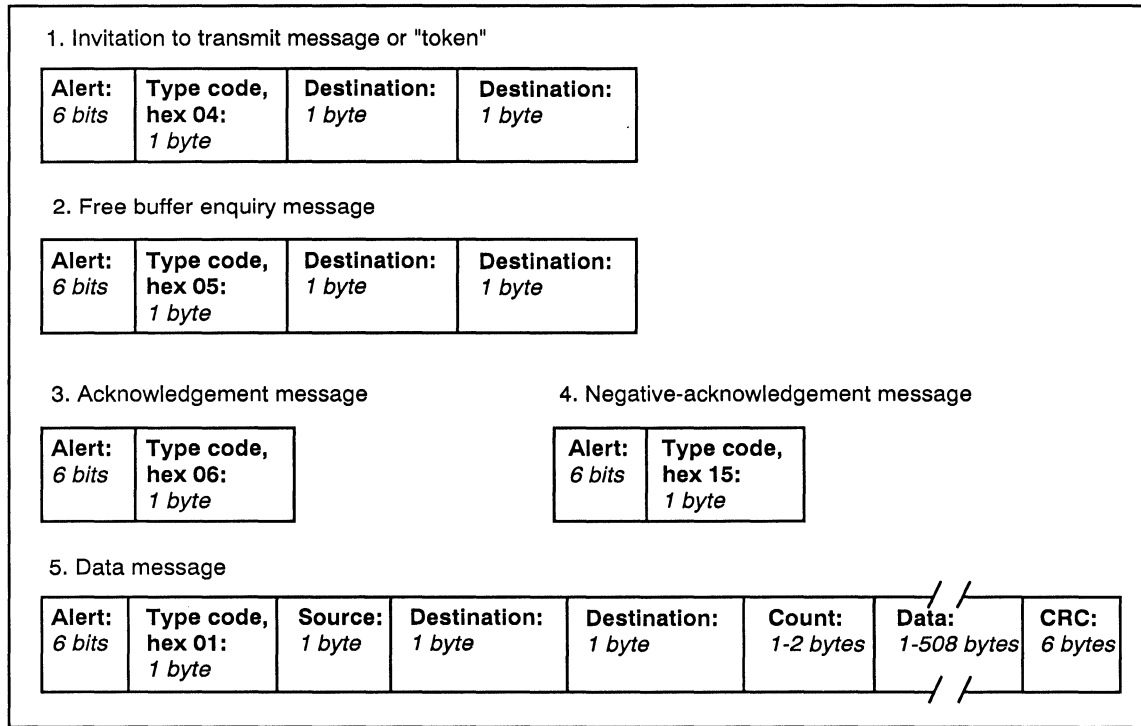


Figure 1-16. Five ARCNET messages.

Each field used in the five ARCNET messages is described below.

- **Alert.** A 6-bit announcement that a message is to follow.
- **Type code.** Indicates which one of the five types of messages it is. Each message type contains only the necessary information. Note, for example, that the token does not contain the address of the station it is *from* but only the station it is *to*.
- **Source.** The address of the station originating a data message.
- **Destination.** The address of the station slated to receive the data message. The repetition of the destination address is an artifact of the original implementation and is not strictly necessary.

- **Count.** The count field of the data message is a complicated encoding of the number of bytes in the packet that is another holdover from the details of original ARCNET implementation. For data lengths (N) from 1 to 253 bytes (i.e., “short packets”), the count field is a single byte whose value is 256-N. For data lengths from 257 to 508 bytes (i.e., “long packets”), the count field is 2 bytes. The first is zero, and the second is 512-N. Frames with data lengths from 254 to 256 cannot be sent at all.
- **Data.** Up to 508 bytes of data starting with the **System ID** (see Figure 1–17).
- **CRC.** The CRC at the end of the data frame is a two-byte *cyclic redundancy check* that the receiver uses to verify that the data has been received correctly.

Format of the Data

The Sniffer analyzer records only data messages and, thus, avoids having the buffer cluttered with the other messages not concerned with higher-level protocols.

The Sniffer analyzer records data frames in a “normalized” fashion that simplifies presentation of the data.




Figure 1–17. A data frame as recorded and “normalized” by the Sniffer Network Analyzer.

The length field indicates the number of data bytes that follow. Thus, fields that are at the same offset in the data in both long and short frames will appear in the same position in the displayed frame.

- **System ID.** A convention established by Datapoint uses the first byte of the data field to indicate what first-level protocol is being used. System IDs from hex 00 to hex 7F are reserved for Datapoint. IDs from hex 80 to hex FF are assigned to other organizations by Datapoint upon request.
- **Protocol data.** Up to 508 bytes of data starting with the **System ID**.

IBM PC Network™ (Broadband) Architecture

The IBM PC Network (Broadband) is a LAN suitable for interconnecting computers and computer-controlled devices at moderate speeds over moderate distances. The PC Network architecture was defined jointly by IBM and Sytek in 1984. It is not described by any of the existing international standards.

The original name for this network was simply PC Network, and it applied both to the hardware and to software protocols up to the application level. The full name is now IBM PC Network (Broadband), and it applies to the network

hardware and data link control protocol but not any of the higher protocols. Note that there is also now a baseband version of PC Network, but we will use “PC Network” to mean only the “IBM PC Network (Broadband).”

Speed and Transmission Technique

Every station connected to a PC Network can hear what any station transmits. The delay between transmission and reception depends only on the propagation delays through the wires and the various attaching devices. In this way, PC Network is similar to networks like Ethernet and StarLAN (see “Ethernet® Network Architecture” on page 1–3 and “StarLAN™ Network Architecture” on page 1–25). It differs fundamentally, however, from networks like the IEEE 802.5 token ring, where stations are wired in a logical ring so that each station only hears what its upstream neighbor transmits (see “Token Ring Network Architecture” on page 1–12). The transmission speed on PC Network is 2 Mbps, or 250,000 bytes per second.

The bits to be sent are not simply placed on the network cable but are used to modulate a high-frequency carrier signal by using the *frequency-shift-keying* (FSK) technique. This use of *radio frequency* (RF) transmission is what characterizes PC Network as a broadband network when compared to baseband networks such as Ethernet.

As with many bidirectional broadband systems, nodes transmit at a relatively low frequency but listen at a much higher frequency. A special device on the network called a *translator* or *headend* converts a low-frequency transmission by any station into an equivalent high-frequency signal that all stations hear. The frequency difference between the two signals is called the *offset*. Note that unlike dual-cable broadband systems, both the original low-frequency signal and the translated high-frequency signal are present on the same cable at the same time. The cable can also carry completely unrelated information on other channels. Two examples are closed-circuit TV transmissions used by plant security in some installations and standard cable TV signals.

The total bandwidth used by a single transmission is 12 MHz (one 6-MHz transmit channel and one 6-MHz receive channel). This applies to the entire PC Network since all stations typically use the same frequencies. It is the same frequency-division multiplexing scheme used by CATV systems. The entire PC Network uses the equivalent of only two TV channels on the cable. Whether the network is compatible with a particular CATV system depends on the choice of carrier frequencies.

Broadband systems differ in the amount of cable bandwidth devoted to *forward* traffic (from the headend to nodes) versus *return* traffic (from nodes to the headend). Systems designed primarily for television are typically *low-split* so that most traffic is in the forward direction. Computer-oriented systems are *mid-split* or *high-split* and allow approximately equal traffic in both directions.

There are four different versions of the PC Network adapter that differ only in their frequency assignments. The numbers in the first two columns of Figure 1–18 are the center frequencies (in megahertz) of the 6-MHz channels. The symbols in parentheses are the channel names.

Transmit	Receive	Offset
50.75 (T14)	219.00 (J)	168.25
56.75 (2')	249.00 (O)	192.25
62.74 (3')	255.00 (P)	192.25
50.75 (T14)	243.00 (N)	192.25

Figure 1–18. Frequency assignments for different versions of the PC Network adapter.

The assignments on the first line are those for the original 1984 PC Network, which used a non-standard offset. The other assignments were added later and use an offset that is also the standard for MAP (IEEE 802.4) networks. All the adapters in a single network must use the same carrier frequencies, although multiple independent networks may operate at different frequencies simultaneously on the same cable system.

Physical Interconnection

Stations are connected to an attachment point of the network using standard CATV hardware: RG-59, 75-ohm coaxial cable and F-connectors. For higher reliability, the screw-on type connectors should be used, although a push-on connector will work as long as it makes a tight connection.

Station cables typically attach to a device called a *signal splitter, directional coupler, or attenuator*. The usual topology is a tree with the headend at the root and splitters branching off to other splitters or nodes. It is also possible to use a more bus-oriented topology with the headend at a less central location.

A broadband cable system must be designed carefully so that both receive and transmit signals have the proper amplitude at every node and at the headend. For small networks there are kits available for various combinations of cable distances and number of nodes. Kits available for the PC Network limit the number of stations to 72 and the maximum radius of the network to 1000 feet.

Larger or unusual networks must be custom designed using carefully chosen splitters and attenuators and may require special RF test equipment for setup and maintenance. Among the problems to be solved is that different frequencies are attenuated by different amounts. To compensate for this, the designer must use a frequency-dependent attenuator called a *tilt compensator or equalizer*.

Access Control

PC Network LANs, like Ethernet and StarLAN, use the CSMA/CD algorithm to control transmission. See the discussion of access control in “Ethernet® Network Architecture” on page 1–3.

Format of a Frame

The frame format used in PC Network is very similar to that used in Ethernet. See the discussion of frame format in “Ethernet® Network Architecture” on page 1–3.

Format of the Data

The data format used in PC Network is very similar to that used in Ethernet. See the discussion of data format in “Ethernet® Network Architecture” on page 1–3.

Assignment of Network Addresses

The assignment of network addresses in PC Network is similar to that in Ethernet. See the discussion on the assignment of network addresses in “Ethernet® Network Architecture” on page 1–3.

LocalTalk® Network Architecture

LocalTalk is a type of LAN suitable for low-speed interconnection of computers and computer-controlled devices over short distances. The LocalTalk architecture is defined by Apple Computer, Inc.

Physical Interconnection and Speed

Stations connected to a LocalTalk network are all connected to the same wire, so that every station hears what any station transmits. The delay between transmission and reception depends only on the propagation delays through the wires and attaching devices. In this way LocalTalk is similar to networks like Ethernet, StarLAN and IBM PC Network (see “Ethernet® Network Architecture” on page 1–3, “StarLAN™ Network Architecture” on page 1–25, and “IBM PC Network™ (Broadband) Architecture” on page 1–31). It differs fundamentally, however, from networks like the IEEE 802.5 token ring, where stations are wired in a logical ring so that each station hears only what its upstream neighbor transmits (see “Token Ring Network Architecture” on page 1–12).

The LocalTalk transmission speed is 230.4 Kilobits per second (Kbps), or 28,800 bytes per second. LocalTalk’s signalling standard is EIA modified RS-422 (balanced voltage), and signal encoding uses a technique known as FM-0 (also called *biphase space*). Because of the access control mechanism described below, the effective network throughput is significantly less than the transmission speed.

LocalTalk is a *bus* network system. A *connector box* is attached to all devices on the network with either a DIN-8 minicircular plug or DB-9 plug, depending upon the device. The connector boxes are then connected to each other with LocalTalk cable equipped with DIN-8 plugs. LocalTalk cabling consists of shielded twisted pair 22 AWG stranded wire.

The maximum network cable length is 300 meters. The maximum number of connector boxes in any single network is 32.

Access Control

LocalTalk uses an access discipline to control transmission called *carrier sense multiple access with collision avoidance* (CSMA/CA). A similar architecture is *carrier sense multiple access with collision detection* (CSMA/CD), implemented in Ethernet and its variants, StarLAN and PC Network (see "Ethernet® Network Architecture" on page 1–3).

Like the CSMA/CD technique, CSMA/CA requires each node wishing to transmit data to check the transmission medium before sending (e.g., carrier sense) and permits more than one node to access the link (e.g., multiple access), but not simultaneously. Because of multiple access, both techniques run the risk of two or more nodes attempting to transmit at the same time and having their respective frames collide. The difference between how each handles this situation is significant.

With CSMA/CA, hardware does not actually detect collisions. Rather, collisions are inferred through the use of a "handshake" mechanism. The handshake involves a transmission dialog using LocalTalk Link Access Protocol (LLAP) control packets prior to sending any data. The transmission dialog begins when the transmitting node's physical layer senses that the line is idle. After sensing an idle line, the transmitting node has a wait period of 400 microseconds plus a randomly generated period. It then sends an LLAP *request-to-send* (RTS) control packet to the intended receiver. The transmitter waits up to 200 microseconds for an LLAP *clear-to-send* (CTS) control packet from the receiver.

If the transmitter receives the CTS packet, it considers the handshake successful. A successful handshake means that a collision did not occur. All other nodes defer 200 microseconds for the transmitter to send its data packet.

If the transmitter does not receive the CTS packet within the allotted 200 microseconds, the transmitter backs off and retries after a random wait period. The transmitter assumes that a collision took place and retries for up to 32 times before reporting failure to its client. If in fact the RTS packet did collide, the receiving node uses the frame check sequence to discover the corrupted packet and to discard it.

The transmission dialog starts similarly for broadcast packets. But after a transmitter sends an RTS packet with the destination address set to all nodes, it does not expect to receive a CTS packet. Rather, it continues to sense the line. If the line remains idle for 200 microseconds, the transmitter sends its data frame.

If the broadcast RTS packet forces a collision, the other transmitting node will back off. The broadcasting node delays until the link is idle again and until the wait period is over. It tries again for up to 32 times before reporting failure to its client, although collisions with other broadcasting nodes may not be noticed.

Format of a Frame

A LocalTalk network uses the LLAP at the data-link layer to allow network devices to share the communication medium. The LLAP frame **preamble** identifies the start of the frame with *flag* bytes. A **flag** byte also identifies the end of the frame. A flag is a frame delimiter that is identifiable with a special bit sequence, 01111110. Sometimes a flag-like sequence will be inserted into the user data stream by the application process and creates the possibility for an error. To protect against such errors, LLAP uses a technique known as *bit-stuffing* (see the discussion of bit-stuffing in "WAN/Synchronous Architecture" on page 1-20).

Preamble: <i>2 or more flag bytes</i>	Destination Node ID: <i>1 byte</i>	Source Node ID: <i>1 byte</i>	LLAP Type: <i>1 byte</i>	Protocol Data: <i>2 to 600 bytes</i>	FCS: <i>2 bytes</i>	Flag: <i>1 byte</i>	Abort Sequence: <i>12 to 18 1s</i>
---	--	---	------------------------------------	--	-------------------------------	-------------------------------	--

Figure 1-19. LLAP frame format.

LLAP uses a 1-byte node identifier number to identify each node on a link. For transmitting and receiving stations, these serve as **source node** and **destination node** data-link addresses.

The **LLAP type** field specifies either a control packet or a data packet. There are four types of control packets, as listed in Figure 1-20. Control packets do not have a data field. Two are involved with the dynamic assignment of addresses in LocalTalk. Unlike other network data links, LocalTalk nodes do not have fixed unique addresses. Nodes newly activated on a network choose an address for themselves. They use an **enquiry** packet to verify if the address they chose is unique. If it is already in use, the node with that address responds with an **acknowledgment** packet, and the new node tries again with another enquiry.

LLAP Type	Hex Value	Action
lapENQ	hex 81	Enquiry. Used for dynamic assignment of node address.
lapACK	hex 82	Acknowledgment. Responds to an Enquiry packet.
lapRTS	hex 84	Request-to-send. Notifies destination node that a data packet is ready.
lapCTS	hex 8	Clear-to-send. Responds to RTS indicating readiness to accept packet.

Figure 1-20. Indications used in LLAP control frames.

The other two control packets are part of LLAP's RTS-CTS handshake mechanism. As noted above in "Access Control," RTS-CTS handshake is successful if the node transmitting the **request-to-send** receives a **clear-to-send**

packet from the receiving node in the appropriate time. The handshake is not successful when the CTS packet is not received as expected.

Format of the Data

If the packet type field indicates a data packet, then the first two bytes of the data field must contain **data length** information in bytes.




Figure 1-21. LLAP data frame format.

The low-order 10 bits of the data length field contain the size in bytes (most significant bits first) of the data field, including the two data length bytes. Higher-level protocols reserve the use of the high-order 6 bits of the data length field.

Major Protocol Suites

Overview

Following a brief introductory section on the Open Systems Interconnection (OSI) layered protocol model, Chapter 2 provides some general information on each of the major protocol suites interpreted by the Sniffer analyzer. The information includes a short discussion of the suite itself, a diagram that maps each protocol to the OSI model, and a summary of the services provided by each protocol.

Introduction

The Sniffer analyzer does not simply monitor, capture, or record network traffic. It also interprets what it records. Protocol interpretation is what makes the Sniffer analyzer a valuable tool. Interpretation turns an inscrutable stream of bits and bytes into clearly labeled commands, responses, and readable text.

Interpretation routines are an integral part of the Sniffer analyzer software, built-in at the factory during a process that equips each machine for the network requested by the customer. NGC calls the interpretation routines *protocol interpreters* (PI). The interpretation they do covers the full range of the OSI seven-layer model.

<i>Application</i>	Layer 7 is concerned with the support of end-user application processes.
<i>Presentation</i>	Layer 6 provides for the representation of the data.
<i>Session</i>	Layer 5 performs administrative tasks and security.
<i>Transport</i>	Layer 4 ensures end-to-end, error-free delivery.
<i>Network</i>	Layer 3 is responsible for addressing and routing between subnetworks.
<i>Logical Link</i>	Layer 2 is responsible for the transfer of data over the channel.
<i>Physical</i>	Layer 1 handles physical signaling, including connectors, timing, voltages, and other matters.

Figure 2-1. Layers of the OSI Network Model.

At the *physical* layer, the analyzer's hardware is responsible for sending and receiving network signals. For the *logical link* layer, interpreters are included to match the networks on which the analyzer will be used.

The protocols at the higher layers—*network*, *transport*, *session*, *presentation*, and *application*—are largely independent of the physical layer.

IBM® Protocol Interpreter Suite

The Sniffer analyzer interprets frames in four families of higher-level protocols widely used on IBM networks. While IBM uses these protocols on LANs connected by token ring, they may also be found on networks connected by other media. The IBM PI suite may be installed in a Sniffer analyzer unit equipped for connection to a LAN such as token ring, Ethernet, or to a wide area network's WAN/synchronous link by way of an RS-232 or V.35 interface.

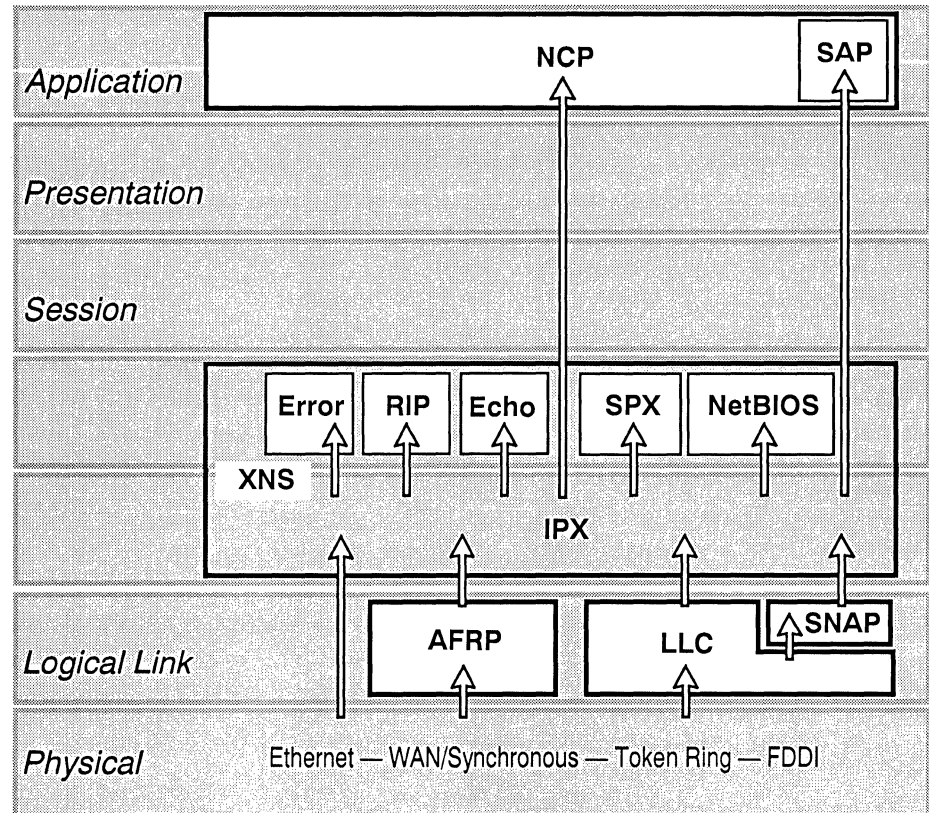


Figure 2-2. The IBM PI suite related to the OSI reference model.

Protocols Interpreted

SNA

Systems Network Architecture. SNA is IBM's name for its very extensive family of commands within a common protocol. It is widely used to connect a broad range of devices, from terminal controllers to micro- or minicomputers, to IBM mainframes. The interpreter decodes the SNA transmission header, the request/response header, and the request and response content by area, including transmission services, function management, management services, presentation services, and general data stream. Both LU2.0 and LU6.2 commands are decoded.

SMB

Server Message Block. A family of application-level commands for LAN servers developed by Microsoft® for use with the IBM PC LAN Program but frequently used in other environments as well. Many of the functions are similar to those made by an application program to DOS or to OS/2® running on a single computer. The IBM PC LAN Program sends SMBs as data within NetBIOS frames, but in other contexts, they may be sent differently. The SMB protocol for machines running under OS/2 contains extensions not present in the version for DOS machines. The PI decodes both the older DOS versions and the extended OS/2 versions.

RPL

Remote Program Load. A protocol used by IBM on the IEEE 802.5 token ring network to download initial programs into networked stations.

NETBIOS

Network Basic I/O System. A protocol implemented in the IBM PC LAN Program to support communication between symbolically named stations and the exchange of arbitrary data between symbolically named stations. (Some of the other NetBIOS implementations differ from the IBM version. The NETBIOS module of the IBM PI suite differs accordingly from the corresponding NetBIOS modules of Novell NetWare® and TCP/IP.)

IBMNM

IBM Network Management Protocols (LLC SAP F4). Used for the LAN Reporting Mechanism, Ring Error Monitor, Configuration Report Server, Ring Parameter Server, and LAN Bridge Server.

BPDU

Bridge Protocol Data Units. Used for the 802.1 spanning tree routing algorithm.

LLC

Logical Link Control (IEEE 802.2). A protocol that provides connection control and/or multiplexing to subsequent embedded protocols.

SDLC

Synchronous Data Link Control. IBM's version of the logical link layer protocol whose ISO designation is HDLC. The WAN/Synchronous Sniffer analyzer interprets the subset that provides link-level support for X.25 and SNA.

Novell® NetWare® Protocol Interpreter Suite

The Sniffer analyzer interprets the protocols used by Novell's NetWare family of products, which include an operating system for file servers as well as services in support of remote users on a variety of physical media. The interpreter suite may be installed on Sniffer systems for Ethernet, token ring, or WAN/synchronous.

Each NetWare file server runs directly under a proprietary Novell operating system. Users at DOS- or OS/2-based workstations can redirect their operating system functions to the NetWare servers. What Novell calls Network File Services (NFS) make use of NetWare Core Protocol (NCP) to transmit commands or inquiries from workstations and to receive replies from file servers. NCP in turn makes use of Novell's implementation of the XNS family of protocols developed by Xerox. These protocols are concerned with the transmission and delivery of a packet, but not its interpretation, which is left to the higher-level protocol, NCP.

At the network level, NetWare uses a datagram protocol called IPX that corresponds to Xerox's IDP (Internet Datagram Protocol). Each IPX packet identifies the network, node, and socket of its destination and of its source. A socket may be a function within a node and, hence, affects where the embedded NCP message is interpreted.

NetWare also provides a connection-oriented virtual circuit protocol called SPX (Sequential Packet Exchange) that corresponds to SPP in the XNS protocols. (However, NCP provides connection services without the use of SPX packets.) In SPX, each packet is identified in the same way as an IPX packet, but with additional fields for the source and destination connection, a sequence number within that connection, an acknowledgment number, and an allocation of the number of unacknowledged SPX packets the connection may tolerate.



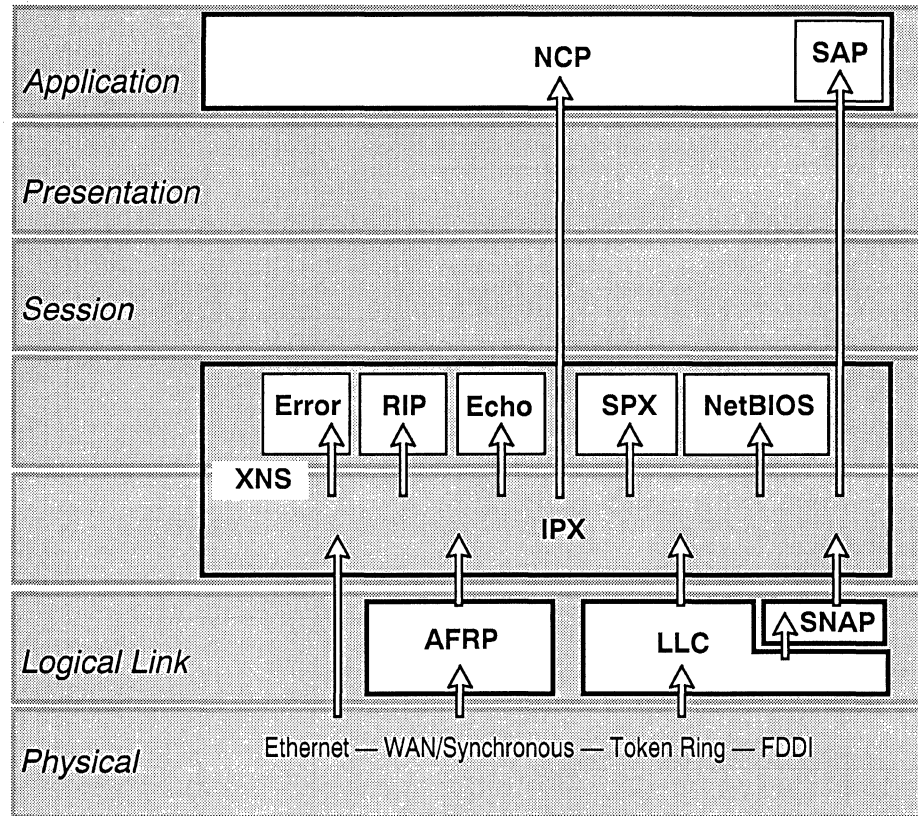


Figure 2-3. The Novell NetWare PI suite related to the OSI reference model.

Protocols Interpreted

NCP

NetWare Core Protocol. Novell's application-level protocol for the exchange of commands and data between file servers and workstations; also described as NetWare File Service Protocol (NFSP).

SAP

Service Advertising Protocol. Used by NetWare servers to broadcast the names and locations of servers and to send a specific response to any station that queries it.

NetBIOS

Network Basic I/O System. NetWare supports emulation of the protocol implemented by the IBM PC LAN Program to support communication between symbolically named stations and the exchange of arbitrary data. In the NetWare context, NetBIOS is atop IPX.

XNS

Xerox Network Systems Protocols. Within this family of protocols, the following are identified:

SPX—Sequential Packet Exchange. Novell's version of the Xerox transport-level protocol called SPP.

IPX—Internet Packet Exchange. This network-level protocol corresponds to Xerox IDP.

RIP—Routing Information Protocol. Novell's version of a protocol used to exchange routing information among gateways.

Echo—Request/response protocol used to verify the existence of a host.

Error—A protocol by which a station reports that it has received (and is discarding) a defective packet.

AFRP

ARCNET Fragmentation Protocol. Breaks up and reassembles network-layer packets so that they are acceptable to the data-link protocol and the underlying physical medium.

LLC

Logical Link Control (IEEE 802.2). A protocol that provides connection control and/or multiplexing to subsequent embedded protocols.

XNS™ Protocol Interpreter Suite

At the network, transport, and presentation layers, this PI suite handles the protocols of the Xerox Network System (XNS). After Xerox published the specifications of these protocols in 1981, several other vendors developed application-layer protocols that run on top of them. The XNS PI suite decodes SMB, a protocol used in Microsoft Networks (MS-NET™) and the IBM OS/2 LAN Manager™.

A network's upper-level protocols are largely independent of its physical layer. While Xerox developed XNS for operation with Ethernet systems, XNS protocols may also be found on networks connected by other media. The XNS PI suite may be installed in a Sniffer analyzer equipped for connection to Ethernet, token ring, or WAN/synchronous.

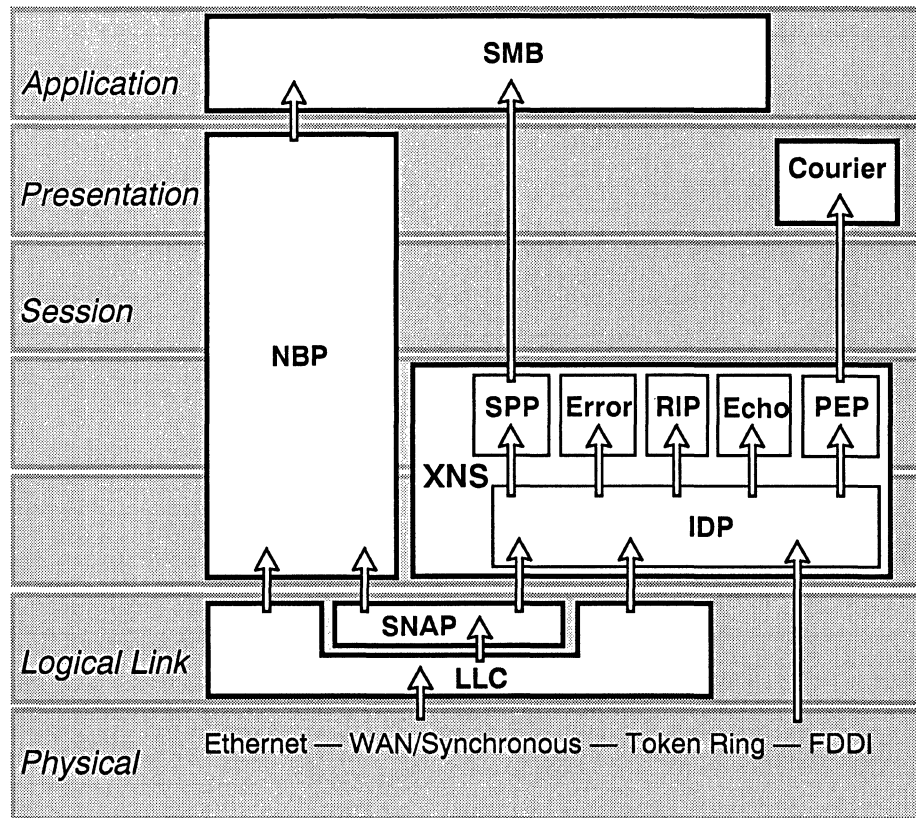


Figure 2-4. The XNS PI suite related to the OSI reference model.

Protocols Interpreted

SMB

Server Message Block. A family of application-level commands for LAN servers developed by Microsoft and IBM for use with the IBM PC LAN Program but frequently used in other environments as well. Many of the functions are similar to those made by an application program to DOS or to OS/2 running on a single computer. Although the IBM PC LAN Program sends SMBs as data within NetBIOS frames, in other contexts they may be sent differently. The XNS PI suite decodes SMB frames transported by the Xerox IDP and SPP protocols. The SMB protocol for machines running under OS/2 contains extensions not present in the version for DOS machines. The XNS PI suite decodes both the older DOS versions and the extended OS/2 versions.

NBP

NetBIOS Protocol. Used in 3Com 3+ Open software.

XNS

Xerox Network Systems Protocol. Within this family of protocols, the XNS PI suite interprets the following:

Courier—A presentation-level protocol that delivers data to such application-level protocols as XNS Printing, XNS Filing, or XNS Clearinghouse (which the XNS PI suite identifies but does not interpret).

SPP—*Sequenced Packet Protocol.* A virtual-circuit, connection-oriented protocol.

IDP—*Internet Datagram Protocol.* Delivers to an internet address a single frame as an independent entity, without regard to other packets or to the addressee's response.

PEP—*Packet Exchange Protocol.* Delivers a request and response pair; this protocol thus has a reliability greater than IDP alone, but less than achievable with SPP.

RIP—*Routing Information Protocol.* Exchanges routing information among gateways and end systems.

Echo—A request/response protocol used to verify the existence of a host.

Error—A protocol by which a station reports that it has received (and is discarding) a defective packet.

TCP/IP Protocol Interpreter Suite

The Sniffer analyzer interprets the protocols of the TCP/IP family and other related protocols. TCP/IP was developed during the 1970's by research institutions under grants from the Advanced Research Projects Agency (ARPANET), US Defense Department. Since its adoption as a standard for ARPANET in 1978, TCP/IP has become widely used in many other networks linking commercial or educational institutions. Although the U.S. Congress has mandated the eventual adoption of ISO protocols, TCP/IP is likely to remain widely used for some time.

While TCP/IP usually runs on Ethernet, its protocols may also be found on other networks. The TCP/IP PI suite may be installed in a Sniffer analyzer for Ethernet, token ring, or WAN/synchronous.

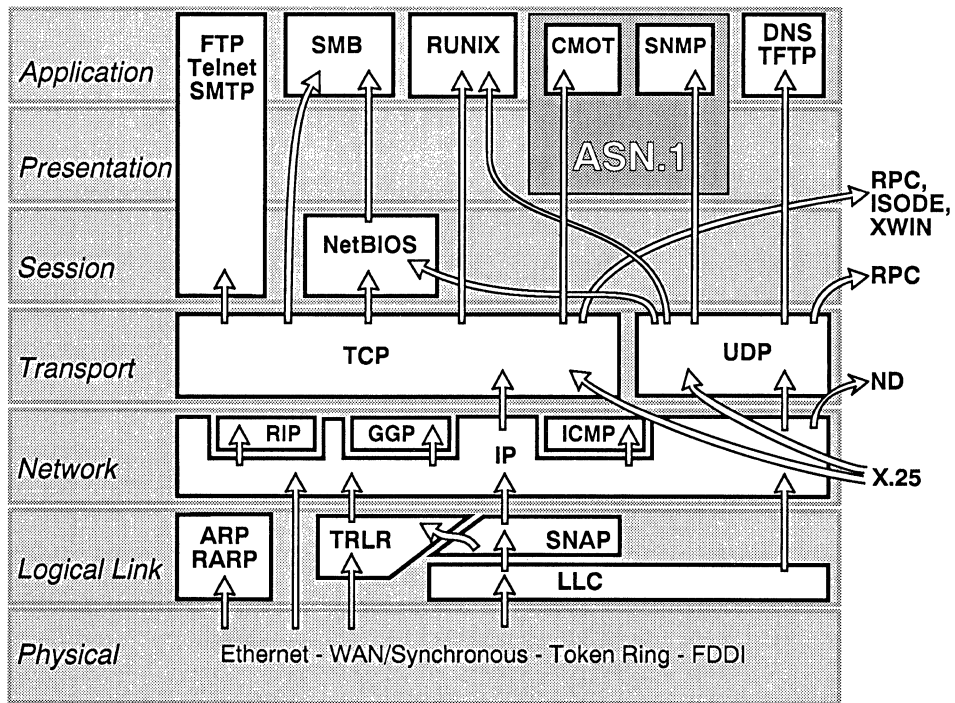


Figure 2-5. The TCP/IP PI suite related to the OSI reference model.

Protocols Interpreted

SMB

Server Message Block. A family of application-level commands for LAN servers developed by Microsoft for use with the IBM PC LAN Program and frequently used in other environments. Although the IBM PC LAN Program sends SMBs as data within NetBIOS frames, in other contexts they may be sent differently. The TCP/IP PI suite decodes SMBs transported by TCP. The SMB protocol for machines running under OS/2 contains extensions not present in the version for DOS machines. The TCP/IP PI suite decodes both the older DOS versions and the extended OS/2 versions.

NetBIOS

Network Basic I/O System. A TCP/UDP version of a protocol developed for the IBM PC LAN Program to support communication between symbolically named stations and transfer of arbitrary data. In the TCP/IP context, NetBIOS is over UDP or IP. (TCP/IP implementations of NetBIOS differ from the IBM version, and the NetBIOS module of the TCP/IP PI suite differs accordingly from the corresponding modules of the IBM PI suite and the Novell NetWare PI suite.)

FTP

File Transfer Protocol. A protocol based on TCP/IP for reliable file transfer.

TFTP

Trivial File Transfer Protocol. A simple protocol used to exchange files between networked stations with less overhead than FTP.

Telnet

A protocol for transmitting character-oriented terminal (keyboard and screen) data.

SMTP

Simple Mail Transfer Protocol. A protocol for reliable exchange of electronic mail messages.

RUNIX

Remote Unix. A protocol for handling remote requests to a UNIX™ host, including commands RLOGIN, RWHO, REXEC, RSHELL, and remote printing.

DNS

Domain Name Service. A protocol for finding information about network addresses using a database distributed among different name servers.

TCP

Transmission Control Protocol. This connection-oriented byte-stream protocol provides reliable end-to-end communication using datagrams sent over IP.

UDP

User Datagram Protocol. Transmits datagrams over IP.

IP

Internet Protocol. Handles end-to-end forwarding and long packet fragmentation control.

RIP

Routing Information Protocol. Exchanges routing information among gateways and end systems.

GGP

Gateway-to-Gateway Protocol. Exchanges routing information among IP gateways.

ICMP

Internet Control Message Protocol. Reports on difficulties in datagram transmission.

LLC

Logical Link Control (IEEE 802.2). Provides connection control and multiplexing to subsequent embedded protocols.

ARP

Address Resolution Protocol. Finds a node's DLC address from its IP address.

RARP

Reverse ARP. Finds a node's IP address from its DLC address.

SNAP

Sub-Network Access Protocol. Also called Sub-Network Access Convergence Protocol. An extension to IEEE 802.2 LLC that permits a station to have multiple network-layer protocols. The protocol specifies that DSAP and SSAP addresses must be AA hex. A field subsequent to SSAP identifies one specific protocol. (See RFC 1042 for more detailed information).

TRLR

Trailer format. A variant of IP in which the protocol headers follow rather than precede the user data.

SNMP

Simple Network Management Protocol. An application-layer protocol for managing TCP/IP based networks. SNMP runs over UDP, which in turn runs over IP. SNMP is organized in terms of network management stations (NMSs), which poll network devices (SNMP agents). At the request of an NMS, an SNMP agent accesses its management information base (MIB) and sends information to the NMS. The MIB provides a standard representation of the SNMP agent's available information and where it is stored.

CMOT

Common Management and Information Services Protocol (CMIP) over TCP. A network management protocol that uses ASN.1 encoding.

SUN[®] Protocol Interpreter Suite

The Sniffer analyzer interprets the protocols that support Sun Microsystems' Network File System (NFS). NFS allows users at workstations to mount directories of files that are located on other machines and to treat them as if they were locally available through the client's operating system. NFS provides an interface that permits a variety of machines (not necessarily under the same operating system) to play the roles of client or server. NFS is composed of a modified UNIX kernel, a set of library routines, and a collection of utilities used by machines that play the role of server.

The Sun PI suite makes use of the session and transport layers of a host network and does not include lower-level protocols of its own. Typically (but not necessarily) Sun NFS runs over TCP/IP on Ethernet. The SUN PI suite interprets frames passed to it by the TCP, IP or UDP protocols and, thus, requires the TCP/IP PI suite.

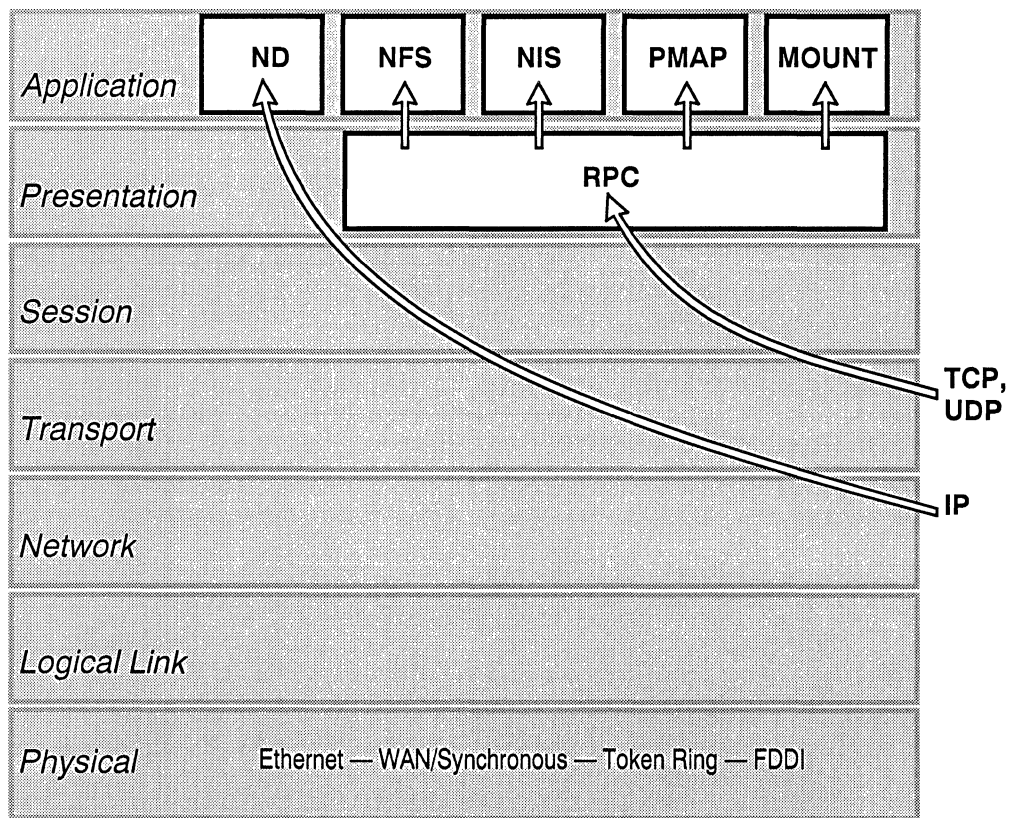


Figure 2-6. The SUN PI suite related to the OSI reference model.

Protocols Interpreted

ND

Network Disk. A protocol used to access virtual disks located remotely across the network and to boot diskless workstations.

NFS

Network File System. The high-level protocol used for communication of requests and responses between network clients and NFS servers. The Sun PI suite interprets NFS Version 2.

NIS

Network Information Services. A high-level protocol used for requests and responses regarding the availability of network hosts, services, and directories from a read-only network database.

YP

Yellow Pages. A high-level protocol used for requests and responses regarding the availability of network hosts, services and directories from a read-only network database.

PMAP

Port Mapper. A protocol for mapping RPC program numbers to TCP/IP port numbers.

MOUNT

A protocol used during initiation of a remote user's access to a network disk, including access checking and account validation.

RPC

Remote Procedure Call. A protocol for activating a function on a remote station and retrieving the result.

ISO Protocol Interpreter Suite

The Sniffer analyzer interprets the family of protocols built upon recommendations of the International Standards Organization as part of an ongoing international cooperative effort in support of Open Systems Interconnection (OSI). It decodes all layers above the physical layer, which may be any of Ethernet, token ring, or WAN/synchronous. It also decodes Microsoft SMBs.

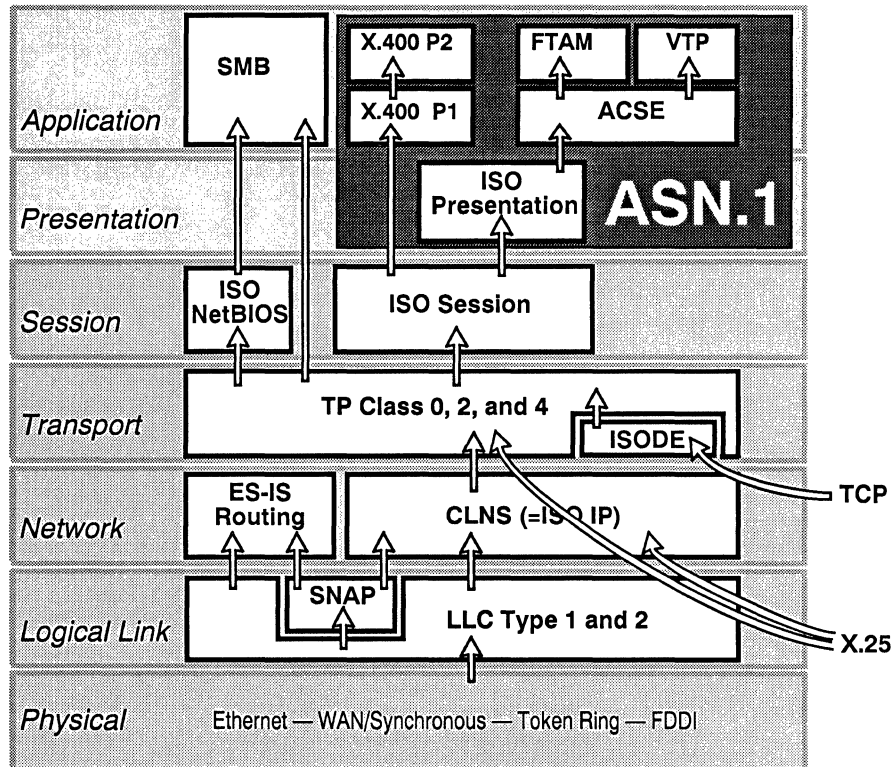


Figure 2-7. The ISO PI suite related to the OSI reference model.

Protocols Interpreted

X.400

The CCITT 1984 protocol for electronic mail. It consists of two levels: P1 for the addressing of the message's outer envelope and P2 for the inner addressing and content of a personal message.

FTAM

File Transfer, Access and Management (ISO 8571/4).

VTP

Virtual Terminal Protocol (ISO 9041).

ACSE

Association Control Service Element (ISO 8650/2). An intermediate application-level protocol used in ISO to support a number of more specific application protocols.

Presentation

(ISO 8823). Application data is encoded using the basic encoding rules (ISO 8825) of Abstract Syntax Notation One (ASN.1, ISO 8824). The user may choose to interpret part of these messages either by sharing the generic ASN.1 syntax structure or by displaying the semantics specific to the application-level protocol.

Session

(ISO 8327).

SMB

Server Message Block. A protocol developed by Microsoft for use with the IBM PC LAN Program to make requests from a user station to a server and receive replies. SMB is part of the protocol family that for DOS machines is called MS NET and for OS/2 machines is called the LAN Manager. The OS/2 version of SMB contains extensions not present in the DOS version; both versions are interpreted in the ISO PI suite.

TP

Transport Protocol (ISO 8073). The ISO PI suite interprets class 0 (for connection-oriented networks), class 4 (for connection-less networks) and the intermediate class 2.

CLNS

Connectionless Network Service Protocol (ISO 8473). Also called ISO IP, for Internetwork Protocol.

ES-IS Routing

End-System to Intermediate-System Routing (ISO 9542). A protocol within the ISO family, used to exchange routing information between gateways and hosts.

LLC

Logical Link Control (ISO 8802/2).

TCP/IP Frames

ISODE

ISO Development Environment. A protocol used to encapsulate higher-level ISO messages when they are transmitted over a network whose lower levels use TCP/IP. (ISODE serves primarily as a development technique during transition from TCP/IP to ISO protocols).

DECnet® Protocol Interpreter Suite

The Sniffer analyzer fully decodes eight protocols defined in Phase IV of Digital Equipment Corporation's *Digital Network Architecture* (DNA). It also decodes several additional protocols that, although not specified in DNA, are used in DECnet systems.

DNA was introduced in 1975 as a master plan for a family of networking hardware and software products valid across a range of machines using both wide-area and local-area networks. Implementation is in phases. Current DEC systems implement Phase IV. DEC plans that Phase V will harmonize DEC's architecture with the general OSI model adopted by the ISO.

DEC provides an implementation of DNA phase IV for each of its operating systems. On LANs, DECnet is commonly used by machines whose physical link is by Ethernet. DECnet protocols can also be used on token ring and WAN/synchronous. The DECnet PI suite can be installed with any of these.

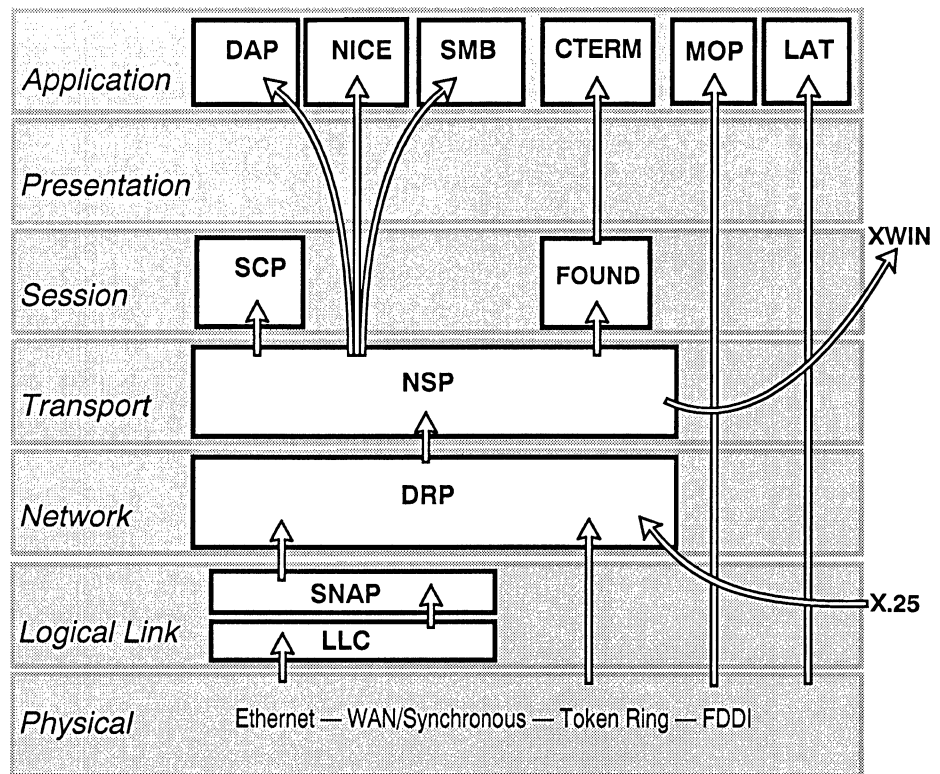


Figure 2-8. The DECnet PI suite related to the OSI reference model.

Protocols Interpreted

DAP

Data Access Protocol. A protocol that provides remote file access operations. It is a command/response protocol that allows a user process to create new files on a server, open existing files, read and write data, and so on.

NICE

Network Information and Control Exchange. A command/response protocol that provides network management information.

SMB

Server Message Block. A message type used by the IBM PC LAN Program to make requests from a user station to a server and to receive replies. Many of the functions are similar to those made by an application program to DOS running on a single computer. It is a protocol for remote file access that is very similar in function to DAP and also dwells in the application layer. It was initially developed for the IBM PC LAN Program and is supported by DEC for compatibility.

CTERM

Command Terminal. A protocol used for communicating with generic intelligent terminals, i.e., a virtual terminal protocol. It is used in conjunction with FOUND.

FOUND

Foundation Services. A protocol used for primitive terminal-handling services and to make and break logical connections between applications and terminals. It is used in conjunction with CTERM.

SCP

Session Control Protocol. A protocol that establishes virtual circuits based on NSP packets.

NSP

Network Services Protocol. A protocol that provides reliable message transmission over virtual circuits. Its functions include establishing and destroying logical links, error control, flow control, and segmentation and reassembly of messages.

DRP

DECnet Routing Protocol. The lowest-level protocol concerned with moving packets from source nodes, through routers, between and within areas, and to end nodes.

MOP

Maintenance Operations Protocol. A protocol used for network maintenance services that include downline loading, upline dumping, and remote testing and problem diagnosis.

LAT

Local Area Transport Protocol. A protocol designed to efficiently handle multiplexed terminal (keyboard and screen) traffic to and from timesharing hosts. LAT is a non-DECnet set of protocols that interfaces directly with the LAN and provides an alternative service to CTERM.

LLC

Logical Link Control (IEEE 802.2). Provides connection control and multiplexing to subsequent embedded protocols for devices on the token ring.

SNAP

Sub-Network Access Protocol. Also called Sub-Network Access Convergence Protocol. An extension to IEEE 802.2 LLC that permits a station to have multiple network-layer protocols. The protocol specifies that DSAP and SSAP addresses must be AA hex. A field subsequent to SSAP identifies one specific protocol. See RFC 1042 for more detailed information.

Banyan[®] VINES[™] Protocol Interpreter Suite

The Sniffer analyzer interprets protocols in the VINES series developed by Banyan Systems. VINES links personal computers to file servers on a LAN, perhaps with gateway links to other LANs or WANs. The user stations are PCs, typically running under DOS. Redirection permits directories on the servers to appear to the users as DOS drives, although each server in fact offers these services through processes running on the server's Unix operating system. The server role may be played by a wide range of devices from different vendors, but all appear to the user in the same way.

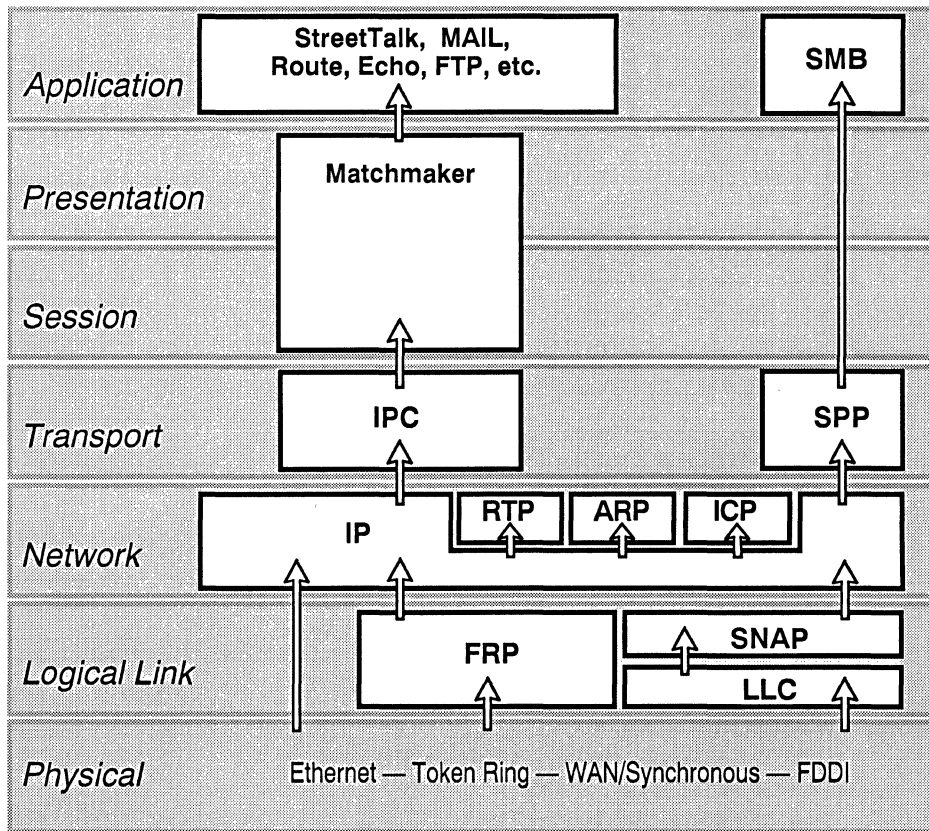


Figure 2-9. The Banyan VINES PI suite related to the OSI reference model.

Protocols Interpreted

StreetTalk

A protocol used in Banyan VINES to maintain a distributed directory of the names of network resources. Names are global across the internet and independent of the network topology.

MAIL

A protocol for the transmission of messages in the VINES distributed electronic mail system.

SMB

Server Message Block. A family of application-level commands for LAN servers developed by Microsoft for use with the IBM PC LAN Program but frequently used in other environments as well. Many of the functions are similar to those made by an application program to DOS or to OS/2 running on a single computer. The IBM PC LAN Program sends SMBs as data within NetBIOS frames. In the VINES contexts, they are transported by SPP. The SMB protocol for machines running under OS/2 contains extensions not present in the version for DOS machines. The Banyan VINES PI suite interprets both the older DOS versions and the extended OS/2 versions.

Matchmaker

A protocol used by the VINES service that provides high-level program-to-program communication and remote procedure calls. Matchmaker services include data translation as necessary to match the conventions of sender's and receiver's formats. Matchmaker is descended from the protocol that in XNS is called Courier.

In addition to MAIL and StreetTalk, the Banyan VINES PI suite identifies the following protocols that may be transmitted by a Matchmaker frame: Echo, Router, Background, File, FTP, Server, Talk, and Network Management.

IPC

Interprocess Communication Protocol. A transport-level protocol providing reliable message service and unreliable datagram service.

SPP

Sequenced Packet Protocol. A transport-level protocol providing virtual connection service, based upon the protocol of the same name in XNS.

RTP

Routing Update Protocol. Protocol used to distribute network topology information.

ARP

Address Resolution Protocol. Used for finding a node's DLC addresses from its IP address.

ICP

Internet Control Protocol. Used to broadcast notification of errors and to note changes in network topology.

IP

Internet Protocol. The protocol that moves datagrams throughout the network.

FRP

Fragmentation Protocol. Breaks up and reassembles network-layer packets so that they are acceptable to the data-link protocol and the underlying physical medium and to the IP protocol above it.

SNAP

Sub-Network Access Protocol. Also called Sub-Network Access Convergence Protocol. An extension to IEEE 802.2 LLC that permits a station to have multiple network-layer protocols. The protocol specifies that DSAP and SSAP addresses must be AA hex. A field subsequent to SSAP identifies one specific protocol. See RFC 1042 for more detailed information.

LLC

Logical Link Control (IEEE 802.2). A protocol that provides connection control and multiplexing to subsequent embedded protocols.

AppleTalk® Protocol Interpreter Suite

AppleTalk protocols link personal computers (frequently but not necessarily Apple computers) to each other and to external services such as gateways, file servers, or printers. AppleTalk is commonly used over Apple Computer's LocalTalk, Ethernet, or WAN/synchronous or may be encapsulated within packets transmitted by an unrelated protocol, for example TCP/IP.

The Sniffer analyzer interprets frames in both Phase 1 and Phase 2 of the AppleTalk family of protocols.

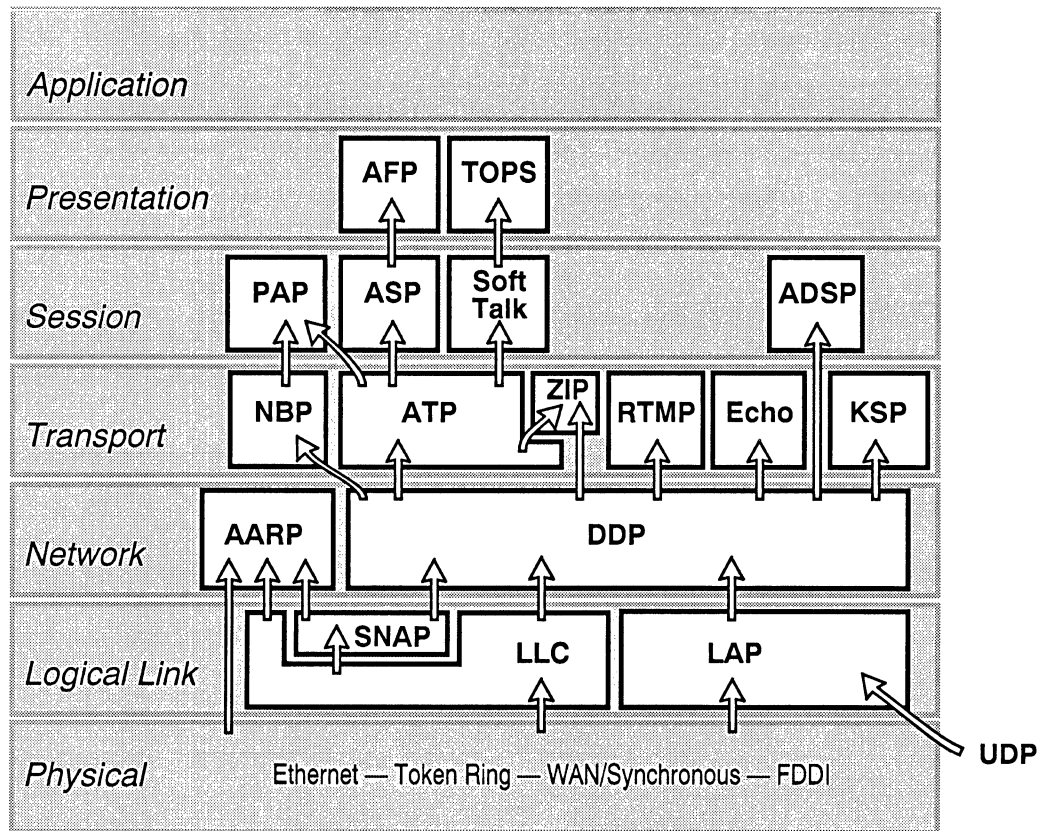


Figure 2-10. The AppleTalk PI suite related to the OSI reference model.

Protocols Interpreted

AFP

AppleTalk Filing Protocol. An application-level protocol for access to remote files.

TOPS

A presentation-level protocol used for remote access to files across different operating systems.

SoftTalk

A session-level protocol including support for remote procedure calls used to support TOPS.

PAP

Printer Access Protocol. A protocol that uses ATP XO (“exactly once”) commands to create a stream-like service for communication between user stations and the Apple LaserWriter® or similar stream-based devices.

ASP

AppleTalk Session Protocol. A general protocol, built upon ATP, providing session establishment, maintenance, and tear-down, along with request sequencing.

ADSP

AppleTalk Data Stream Protocol. A connection-oriented protocol providing a reliable, full-duplex, byte-stream service between any two sockets on an AppleTalk internet, ensuring in-sequence, duplicate-free delivery of data over its connections.

NBP

Name-Binding Protocol. Used in AppleTalk networks to permit network users to refer to network services and sockets by character names. NBP translates a character-string name within a zone into the corresponding socket address.

ATP

AppleTalk Transaction Protocol. Provides a loss-free transaction service between sockets, allowing exchanges between two socket clients in which one client requests the other to perform a particular task and report the result.

RTMP

Routing Table Maintenance Protocol. Used in AppleTalk networks to allow bridges or internet routers to dynamically discover routes to the various subnetworks of an internet. A node that is not a bridge uses a subset of RTMP (the RTMP stub) to determine the number of the network to which it is connected and the node IDs of bridges on its network.

ZIP

Zone Information Protocol. Used to maintain an internet-wide mapping of networks to zone names for the benefit of routers and as a resource for the name-binding protocol (NBP) to determine which networks belong to a given zone.

Echo

A simple protocol that allows any node to send a datagram to any other node and to receive an echoed copy of that packet in return, to verify the node's existence, or to make round trip delay measurements.

KSP

Kiewit Stream Protocol. A transport protocol resembling TCP developed at Dartmouth College for the support of terminal emulators.

AARP

AppleTalk Address Resolution Protocol. Matches the destination address corresponding to a higher-level protocol address.

DDP

Datagram Delivery Protocol. Extends the services of the underlying LAP protocol to include an internet of interconnected AppleTalk networks, with provision to address packets to sockets within a node.

SNAP

Sub-Network Access Protocol. Also called Sub-Network Access Convergence Protocol. An extension to IEEE 802.2 LLC that permits a station to have multiple network-layer protocols. The protocol specifies that DSAP and SSAP addresses must be AA hex. A field subsequent to SSAP identifies one specific protocol. (See RFC 1042 for more detailed information).

LLC

Logical Link Control (IEEE 802.2 and ISO/DIS 8802/2). A protocol that provides connection control and multiplexing to subsequent embedded protocols.

LAP

Link Access Protocol. The logical-link protocol for AppleTalk. It exists in two variants: ELAP for Ethernet and LLAP for LocalTalk.

X Windows™ Protocol Interpreter Suite

The Sniffer analyzer interprets the protocol used to transmit information between X Windows clients and servers. The protocol is independent of the lower-level frames that carry its messages. The X Windows PI suite must be installed in combination either with the TCP/IP PI suite where it interprets frames passed to it by TCP, or with the DECnet PI suite where it interprets frames passed to it by NSP. DECWindows is Digital Equipment Corporation's name for X Windows over DECnet.

X Windows is an outgrowth of Project Athena at the Massachusetts Institute of Technology in 1984. Its development was supported by contributions from Digital Equipment Corporation and IBM. Development of the X Windows system is now supported by a consortium that includes the original sponsors and more than 40 additional vendors. The X Windows PI suite interprets the protocol of the consortium's current standard, Version 11, Release 4.

The X system permits a task's graphic display to be treated independently of the task itself. An application's computations may be done anywhere— at any mainframe, minicomputer, or microcomputer that is accessible through the network. The display is handled independently for each user by a *display server*. The rest of the application's work is handled by a process that acts as a remote *client* of the end-user's display server. The client does not need to know anything about the server's hardware or software. It simply describes its output in terms of the X interface. The server must turn that description into a display. The server can maintain contact with several clients at the same time and, thus, manage multiple windows, sizing, overlaying, moving, or hiding them as the person at the server directs.

Features of the Interpreter

The X Windows protocol permits a sequence of several commands to be concatenated in a single X message. For transmission, an X message may be fragmented into several frames. The X Windows PI suite reassembles a fragmented message. The hex and detail displays show the entire X message (if captured) starting at the first of its DLC frames. The interpreter's summary window shows a separate line for each X command, regardless of the way the commands may have been packed into lower-level frames.

It may happen that the Sniffer analyzer starts recording after transmission of the initial X Windows setup message. The initial message establishes byte-ordering of the transmitted data, and synchronizes the boundaries of X commands within the transport byte stream. If the interpreter has not seen the initial message, for X messages sent over TCP, the X Windows PI suite uses a heuristic to recognize an X message and to establish the byte-order for its data.

Where a frame includes the selection of options as a sequence of bits, most Sniffer analyzer PIs show all the options, as well as an indication of which were selected. However, some X Windows options are so extensive that listing all of them would require dozens or even hundreds of lines. In such cases, the interpreter shows only the options that are selected and omits those that are not.

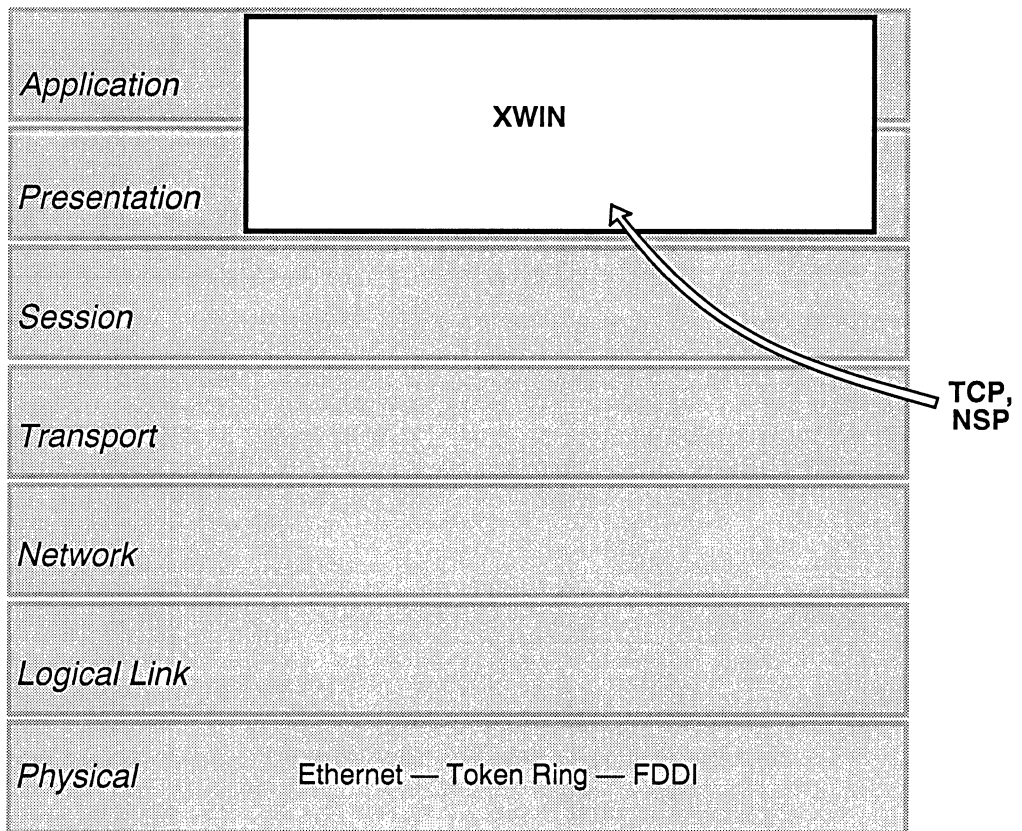


Figure 2-11. The X Windows PI suite related to the OSI reference model.

X.25 Protocol Interpreter Suite

The Sniffer analyzer X.25 PI suite fully decodes six protocols used in the communication links of WANs. It decodes the network layer 3 of the standard usually known as Recommendation X.25 of the CCITT. Also, it decodes certain protocols commonly used above X.25, identifies several other higher-level protocols that may be transmitted over X.25, and passes packets to the appropriate PI suites for display.

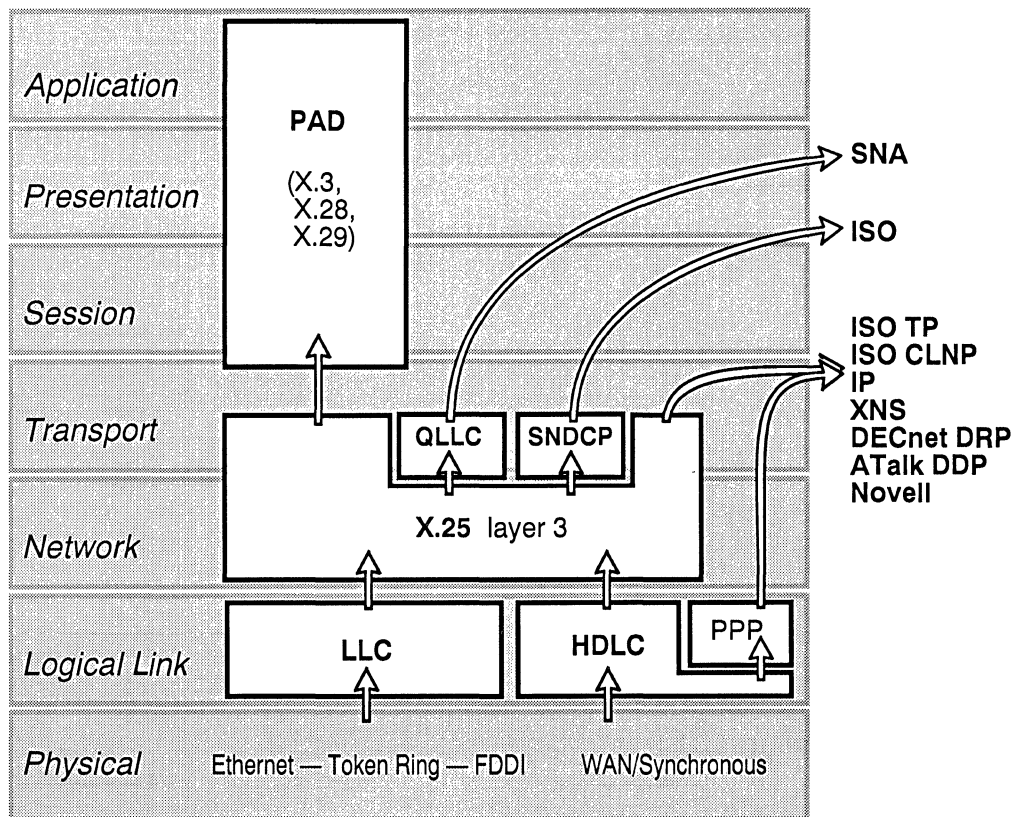


Figure 2-12. The X.25 PI suite related to the OSI reference model.

Protocols Interpreted

PAD

Packet Assembler/Disassembler Protocol. This protocol family provides buffering between traffic at a terminal or similar character-oriented device and the block-oriented communications of a packet-switched network. CCITT recommendation X.3 describes a virtual device that acts as intermediary between the terminal and the X.25 network. The protocol between the terminal and PAD device is described in X.28 and between the PAD device and the X.25 link in recommendation X.29.

X.25

The Sniffer analyzer X.25 PI decodes layer 3 of the 1980 and 1984 versions of CCITT recommendation X.25, including the 1984 extensions for OSI addressing and the ISO and DDN facility and diagnostic fields.

The interpreter recognizes numerous higher-level embedded protocols and (when installed) passes frames to the appropriate PI suite. Protocols thus interpreted include:

ISO TP and ISO CLNP (with the ISO PI suite); IP (with the TCP/IP PI suite); DRP (with the DECnet PI suite); XNS (with the XNS PI suite); DDP (with the AppleTalk PI suite); and NCP (with the Novell NetWare PI suite).

SNDCP

Subnetwork Dependent Convergence Protocol. This intermediate protocol provides an interface between X.25 and the transport layer of an ISO protocol. (The enclosed ISO protocols are interpreted when the ISO PI suite is also installed.)

QLLC

Qualified Logical Link Control Protocol. This intermediate protocol provides an interface between X.25 and the SNA family of protocols. (The enclosed SNA protocols are interpreted when the Sniffer analyzer IBM protocol interpreter is also installed.)

PPP

Point-to-Point Protocol (RFC 1134). This link-level protocol bypasses X.25 for communication between systems that are directly connected, running any of a variety of protocols directly over HDLC.

HDLC

High-level Data Link Control Protocol. This ISO standard protocol is widely implemented as the logical link layer for an X.25 network. (On IBM networks, the corresponding protocol is called SDLC.) The WAN/Synchronous Sniffer analyzer interprets LAPB, the subset of HDLC used to provide link-level support for X.25.

Frame Relay Protocol Interpreter Suite

The Sniffer Internetwork Analyzer interprets the Frame Relay protocol. The primary use for Frame Relay is in internetworking — connecting LANs together by means of WAN links. Frame Relay is used to encapsulate the LAN frames so that they can be forwarded properly by the WAN equipment and software.

Frame Relay operates on the Network and Logical Link layers of the OSI reference model. The frame format is optimized for use over networks that have a relatively low rate of physical errors. The number of bytes used for frame overhead is kept as low as possible relative to the number of bytes of LAN data being forwarded. This makes Frame Relay a particularly efficient protocol compared to others in current use.

The logical path from a sending frame relay link through a WAN to a receiving frame relay link is called a virtual circuit. Such a circuit is identified at its source by a Data Link Connection Identifier (DLCI). A Network Layer protocol called the Local Management Interface (LMI) is used to keep track of the virtual circuits over which Frame Relay can forward data.

Figure 2–13 shows the relationship among LMI, Frame Relay, and the physical WAN equipment.

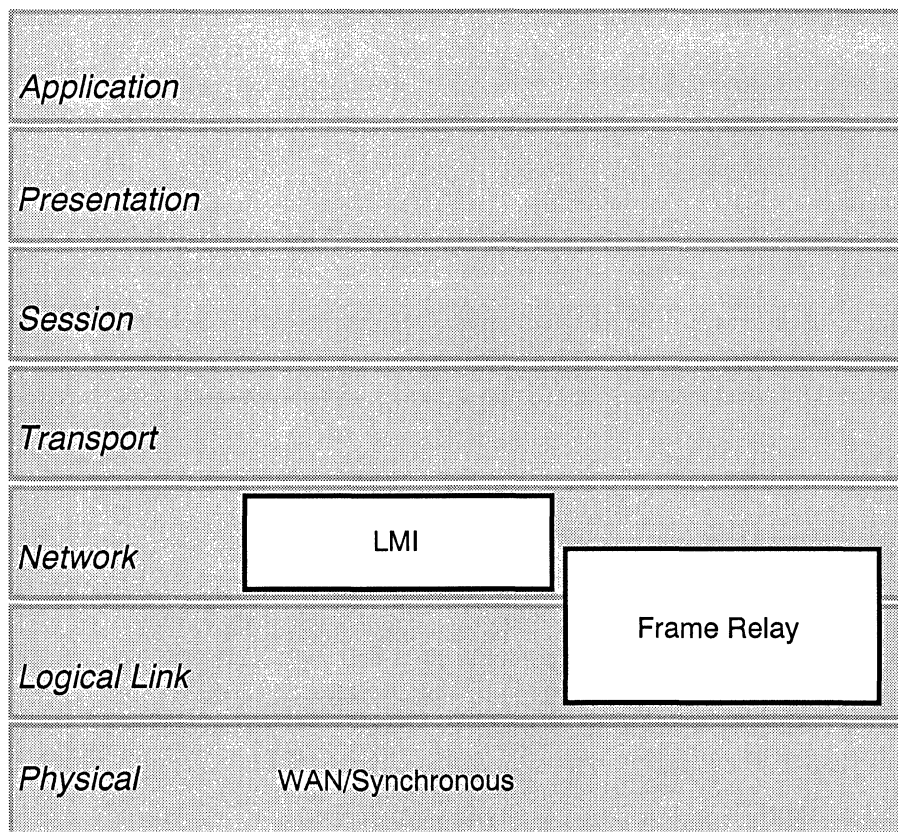


Figure 2–13. The Frame Relay PI suite related to the OSI reference model.

Protocols Interpreted

Frame Relay

The basic frame format for Frame Relay is as follows:

- A one-byte opening flag
- A short header (typically two bytes, containing ten DLCI address bits and six control bits)
- User (subscriber) data (a variable length field)
- A two-byte frame check sequence
- A one-byte closing flag

If a frame is defective in any way, or if WAN congestion precludes its delivery, the Frame Relay protocol simply discards the frame. It is the responsibility of the higher protocol layers to perform the necessary error recovery operations.

The Frame Relay protocol interpreter decodes the Frame Relay header and identifies the DLCI. It also determines from this header whether there is any congestion on the WAN, by decoding the Backward Explicit Congestion Notification (BECN) bit and the Forward Explicit Congestion Notification (FECN) bit. These bits, respectively, notify a data source or a data destination of congestion on the WAN. The interpreter also decodes the Discard Eligible (DE) bit, which can be used to set priorities, in case WAN congestion makes it necessary to discard some frames.

LMI

The LMI protocol interpreter decodes the Local Management Interface information, which is maintained by the WAN to which the Frame Relay interface is attached. The LMI frames are special Frame Relay frames, identified by DLCI = 1023. The LMI information is contained in the user data area of the Frame Relay frame. LMI frames are used to inform subscribers of the status of all the virtual circuits that are involved with the local Frame Relay link. Each subscriber periodically sends a Status Enquiry LMI frame. This can be a simple “keep alive” type of enquiry, in which the subscriber tells the WAN that it is still active, and receives a confirmation from the WAN. Alternatively, it can be a “full status” request, to which the WAN responds by giving the status of all the virtual connections attached to the local Frame Relay link.

SMT Protocol Interpreter Suite

The Sniffer analyzer interprets the SMT protocol, which is used to manage stations on an FDDI network. SMT contains these components:

- Connection management (CMT) connects nodes to the FDDI network, manages physical connection between adjacent PHYs, configures PHY and MAC entities within a node, and coordinates trace functions.
- Ring management (RMT) receives status information from MAC and CMT and reports this status to SMT and higher-level processes.
- SMT frame services provide the means to control and observe the FDDI network.

Figure 2–14 shows the relationship among SMT, MAC, and LLC. The FDDI standards deal specifically with MAC, SMT, and the physical aspects of FDDI. Although Logical Link Control (LLC) is not part of the FDDI standards, FDDI requires LLC for proper operation and data transmission.

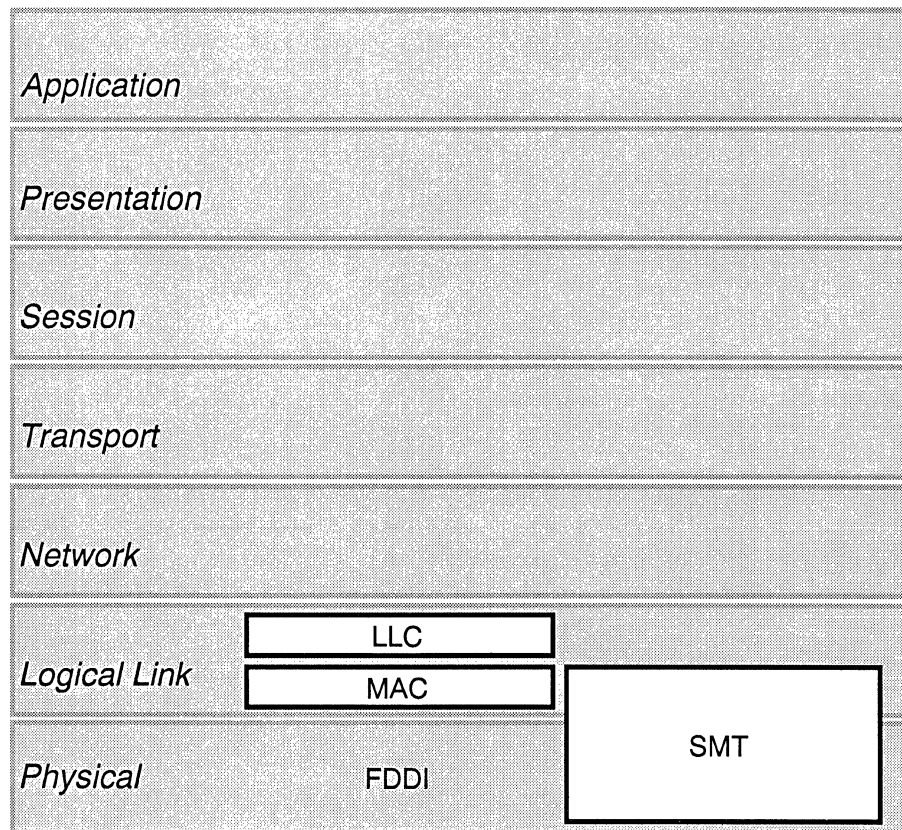


Figure 2–14. The SMT PI suite related to the OSI reference model.

Features of the Interpreter

The SMT protocol interpreter decodes the following SMT frames:

NIF

Neighbor Information Frame. Each station sends out neighbor information frames (NIFs) to its downstream neighbor at intervals of 30 seconds or less. This keeps each station informed as to the address of its upstream neighbor. In addition, a monitoring station can use these frames to construct a logical ring map showing the order in which the stations on the ring appear in the token path.

SIF

Status Information Frame. A station can use a status information frame (SIF) to request or to provide information about station configuration and operating parameters. SIF frames can be sent to all stations, groups of stations, or to individual stations. The SIF responses can be used to create a logical ring map showing detailed configuration information. They can also be used to create a physical map showing the position and connection status of each station.

ECF

Echo Frame. An SMT echo frame (ECF) is used for SMT-to-SMT loopback testing on an FDDI ring. An echo frame may contain any amount of data up to the maximum frame size supported by FDDI.

RAF

Resource Allocation Frame. Resource allocation frames (RAFs) are used to coordinate the allocation of resources among FDDI stations.

RDF

Request Denied Frame. A request denied frame (RDF) is sent in response to a request for an unsupported frame class, type, or version.

ESF

Extended Service Frame. Extended service frames (ESFs) are used for extending and exercising new SMT services.

SRF

Status Report Frame. Status report frame (SRFs) are used to periodically announce station status that may be of interest to the manager of an FDDI ring.

PMF

Parameter Management Frame. Parameter management frames (PMFs) are used for remote management of station attributes.

SNIFFER® NETWORK ANALYZER

APPENDIX A: GLOSSARY

A

Glossary

1BASE5	The implementation of the IEEE 802.3 (StarLAN) standard using 1 megabit per second transmission on a baseband medium whose maximum segment length is 500 meters.
10BASE2	The implementation of the IEEE 802.3 (Ethernet) standard using 10 megabit per second transmission on a baseband medium whose maximum segment length is 185 meters.
10BASE5	The implementation of the IEEE 802.3 (Ethernet) standard using 10 megabit per second transmission on a baseband medium whose maximum segment length is 500 meters.
10BASE-T	The implementation of the IEEE 802.3 (Ethernet) standard using 10 megabit per second transmission on a baseband medium. The standard provides a means for attaching AUI-compatible devices to 24 gauge, unshielded twisted pair cable, instead of the usual coaxial media.
3Com 3+	A networking system from 3Com Corporation using parts of the XNS and Microsoft/IBM PC LAN program protocols.
3Plus	3Com's implementation of XNS. Interpreted by the XNS PI suite.
802.2	The IEEE standards designation for the LLC sublayer protocol that provides both datagram and reliable connection transmission.
802.3	The IEEE standards designation for the CSMA/CD network access method. Similar to (and often used interchangeably with) Ethernet.
802.5	The IEEE standards designation for the token ring network access method.
AARP	AppleTalk Address Resolution Protocol. For outgoing packets, supplies the hardware destination address corresponding to a higher-level protocol address, and filters incoming packets to pass only those that are broadcast or specifically addressed to it. Interpreted in the AppleTalk PI suite.
AC	Access control. A DLC byte on IEEE 802.5 token ring networks that contains the token indicator and frame priority information.
ACSE	Association Control Service Element. An ISO application-level protocol interpreted in the ISO PI suite.
ACTPU	Activate Physical Unit. An SNA message sent to start a session.

ADSP	AppleTalk Data Stream Protocol. A connection-oriented protocol providing a reliable, full-duplex, byte-stream service between any two sockets on an AppleTalk internet, ensuring in-sequence, duplicate-free delivery of data over its connections. Interpreted in the AppleTalk PI suite.
AEP	AppleTalk Echo Protocol. See Echo.
AFP	AppleTalk Filing Protocol. A presentation-level protocol for access to remote files. Interpreted in the AppleTalk PI suite.
ALAP	AppleTalk Link Access Protocol. See LAP.
API	Application Program Interface. The specification of functions and data used by one program module to access another; the programming interface that corresponds to the boundary between protocol layers.
APPC	Advanced Program-to-Program Communications. A communications system used to communicate between transaction programs on IBM computers; APPC uses the LU 6.2 subset of SNA.
ARCNET	A baseband token-passing network originally designed by the Datapoint Corporation that communicates among up to 255 stations at 2.5 Mbps.
ARP	Address Resolution Protocol. (1) A protocol within TCP/IP for finding a node's DLC addresses from its IP address. Interpreted in the TCP/IP PI suite. (2) Interpreted in the Banyan VINES PI suite.
ASCII	American Standard Code for Information Interchange. A mapping between numeric codes and graphical characters used almost universally for all personal computer and non-IBM mainframe applications.
ASN.1	Abstract Syntax Notation One. A set of conventions governing the ISO presentation layer. Interpreted in the ISO PI suite.
ASP	AppleTalk Session Protocol. A general protocol, built upon ATP, providing session establishment, maintenance, and tear-down, along with request sequencing. Interpreted in the AppleTalk PI suite.
ATP	AppleTalk Transaction Protocol. Provides a loss-free transaction service between sockets, allowing exchanges between two socket clients in which one client requests the other to perform a particular task and report the result. Interpreted in the AppleTalk PI suite.
AUI	Attachment Unit Interface. Drop cable for Ethernet between station and transceiver.

Background services	A protocol transmitted by a Matchmaker frame in Banyan VINES.
baseband	A transmission technique that sends data bits without using a much higher carrier frequency (contrast with broadband). The entire bandwidth of the transmission medium is used by one signal.
BIND	An SNA message sent to activate a session between LUs.
BIS	Bracket Initiation Stopped. An SNA message sent to indicate that the sending station will not attempt to initiate any more brackets.
BNC	A standardized coaxial cable connector; used for Thin Ethernet ("Cheapernet") cables and ARCNET networks.
BOOTP	Boot Protocol. A protocol within TCP/IP that is used for downloading initial programs into networked stations. Interpreted in the TCP/IP PI suite.
broadband	A transmission technique that sends data bits encoded within a much higher radio-frequency carrier signal. The transmission medium may be shared by many simultaneous signals since each one only uses part of the available bandwidth.
broadcast	(1) A message directed to all stations on a network or collection of networks. (2) A destination address that designates all stations.
CCITT	International Consultative Committee for Telephony and Telegraphy. CCITT is a member of the International Telecommunications Union (ITU) that is, in turn, a specialized body within the United Nations. It sponsors a number of standards dealing with data communications networks, telephone switching standards, digital systems, and terminals.
CGA	Color Graphics Adapter. The interface between a personal computer and a medium-resolution color monitor.
CLNS	Connectionless Network Service Protocol (also called ISO IP). Interpreted in the ISO PI suite.
CMIP	Common Management Information and Services Protocol. When used with TCP/IP, it is also known as CMOT.
CMOT	Common Management Information and Services Protocol Over TCP. A management protocol for networks; it uses ASN.1 encoding. Interpreted in the TCP/IP and ISO PIs.

Courier	A presentation-level protocol in XNS (similar to RPC in the Sun protocol family); it delivers data to such application-level protocols as XNS Printing, XNS Filing, or XNS Clearinghouse.
CRC	Cyclic Redundancy Check. A check-word, typically two or four bytes at the end of a frame, used to detect errors in the data portion of the frame.
CSMA/CA	Carrier Sense Multiple Access with Collision Avoidance. A <i>random access</i> or <i>contention-based</i> control technique; the algorithm used in LocalTalk networks to control transmission.
CSMA/CD	Carrier Sense Multiple Access with Collision Detection. A <i>random access</i> or <i>contention-based</i> control technique; the algorithm used by IEEE 802.3 and Ethernet networks to control transmission.
CTERM	Command Terminal. A protocol within DECnet for communicating with generic intelligent terminals, that is, a virtual terminal protocol. Interpreted in the DECnet PI suite.
DAP	Data Access Protocol. The DECnet protocol that provides remote file access. Interpreted in the DECnet PI suite.
DB-9	A 9-pin standardized connector used in personal computers for a token ring network connection (female), serial I/O port (male), and RGBI output. Also used for LocalTalk.
DB-15	A 15-pin standardized connector used at the transceiver, the drop cable, and the station of IEEE 802.3 or Ethernet network components.
DB-25	A 25-pin standardized connector used in personal computers for parallel output ports (female connector on IBM PC chassis) or for serial I/O ports (male connector on IBM PC chassis).
DCE	Data Circuit-terminating Equipment (also called Data Communications Equipment). On a serial communications link, the device that connects the DTEs into the communication line or channel.
DDP	Datagram Delivery Protocol. Extends the services of the underlying LAP protocol to include an internet of interconnected AppleTalk networks, with provision to address packets to sockets within a node. Interpreted in the AppleTalk PI suite.
DFC	Data Flow Control. An SNA subprocess for reliable message transfer.
DIS	Draft International Standard.

DISC	Disconnect. An LLC non-data frame indicating that the connection established by an earlier SABM or SABME is to be broken.
DIX	DEC/Intel/Xerox. Used to refer to an early version of Ethernet.
DLC	Data Link Control. The lowest protocol level within the transmitted network frame; fields typically include the Destination address, and Source address, and perhaps other control information.
DLL	Downline load. A protocol within the Datapoint RMS family used for downloading initial programs into networked stations.
DM	Disconnected Mode;. An LLC message acknowledging that a previously established connection has been broken.
DNS	Domain Name Service. A protocol within TCP/IP for finding out information about resources using a database distributed among different name servers. Interpreted in the TCP/IP PI suite.
DOS	Disk Operating System. The most common operating system for IBM-compatible personal computers.
DRP	DECnet Routing Protocol. The lowest-level DECnet protocol, concerned with moving packets from endnodes through routers to other endnodes. ("Routing" in DNA terminology corresponds to the ISO model's "Network" layer).
DSAP	Destination Service Access Point. The LLC SAP for the protocol expected to be used by the destination station in decoding the frame data.
DTE	Data Terminal Equipment. On a serial communications link, a generic term used to describe the host or end-user machine.
EBCDIC	Extended Binary-Coded-Decimal Interchange Code. A mapping between numeric codes and graphical characters used for IBM mainframe computers and communications protocols defined by IBM.
Echo	(1) A request/response protocol within XNS used to verify the existence of a host. (2) A protocol within AppleTalk that allows any node to send a datagram to any other node and to receive an echoed copy of that packet in return to verify the existence of that node or to make round trip delay measurements. Interpreted in the AppleTalk PI suite. (3) A protocol transmitted by a Matchmaker frame in Banyan VINES.

EGP	Exterior gateway protocol. A protocol within TCP/IP used to exchange routing information among gateways belonging to the same or different systems. A generalization of GGP.
ELAP	See LAP.
Error	A protocol within XNS by which a station reports that it has received (and is discarding) a defective packet. Interpreted in the XNS PI suite.
ES-IS Routing	End-System to Intermediate-System Routing. A protocol within the ISO family used to exchange routing information between gateways and hosts. Interpreted in the ISO PI suite.
Ethernet	A CSMA/CD network standard originally developed by Xerox; similar to (and often used interchangeably with) the IEEE 802.3 standard.
Ethertype	A 2-byte protocol-type code in Ethernet frames used by several manufacturers but independent of the IEEE 802.3 standard.
FC	Frame control. On a token ring network, the DLC byte that contains the frame's type.
FCS	Frame check sequence. A redundant check field used to increase the probability of error-free transmission on the network.
FDDI	Fiber Distributed Data Interface. ANSI/ISO standards that defines a 100Mb/s LAN over a fiber-optic media using a timed token over a dual ring of trees.
FID	Format Identification. A field in the SNA Transmission header indicating the type of nodes participating in the conversation. LU 6.2 nodes are type 2.
FMD	Function Management Data. A class of data embedded at the start of SNA RUs.
FMH	Function Management Header. The header part of SNA FMD containing addressing and transmission control information.
FOUND	Foundation Services. A protocol within DECnet used for primitive terminal-handling services. Interpreted in the DECnet PI suite.
frame	The multi-byte unit of data transmitted at one time by a station on the network; synonymous with Packet.
FRMR	Frame Reject. An LLC command or response indicating that a previous frame had a bad format and is being rejected. The REJ frame contains five bytes of data explaining why and how the previous frame was bad.

FRP	Fragmentation Protocol. Breaks up and reassembles network-layer packets so that they are acceptable to the data-link protocol and the underlying physical medium; used on networks whose physical medium is ARCNET. Interpreted in the Banyan VINES PI suites.
FS	Frame status. A byte appended to a token ring network frame following the CRC. It contains the Address Recognized and Frame Copied bits.
FTAM	File Transfer, Access and Management. An application-level protocol within the ISO suite, on top of ACSE.
FTP	File Transfer Protocol. (1) A protocol based on TCP/IP for reliable file transfer. Interpreted in the TCP/IP PI suite. (2) A protocol transmitted by a Matchmaker frame in Banyan VINES.
Functional address	A limited broadcast destination address for IEEE 802.5 token ring networks. Individual bits in the address specify attributes that stations eligible to receive the frame should have. Similar to "multicast address."
GGP	Gateway-to-gateway protocol. A protocol within TCP/IP used to exchange routing information between IP gateways and hosts. Interpreted in the TCP/IP PI suite. See also EGP.
hub	A concentrator and repeater for the StarLAN or the ARCNET network. For StarLAN, it is more properly known as a Network Hub Unit or as a Network Extension Unit.
I	Information. An LLC, HDLC, or SDLC frame type used to send sequenced data that must be acknowledged.
ICMP	Internet Control Message Protocol. A protocol within TCP/IP used principally to report errors in datagram transmission. Interpreted in the TCP/IP PI suite.
ICP	Internet Control Protocol. Used to broadcast notification of errors and to note changes in network topology in Banyan VINES. Interpreted in XNS PI suite.
IDP	Internet Datagram Protocol. Delivers to an internet address a single frame as an independent entity, without regard to other packets or to the addressee's response.
IEEE	Institute of Electrical and Electronics Engineers, Inc. Standards documents are available from them at 345 East 47th Street, New York, NY 10017.

IONET	Input/Output Network. A device message protocol used by Datapoint.
IP	Internet Protocol. The lowest-level protocol under TCP/IP that is responsible for end-to-end forwarding and long packet fragmentation control. Interpreted in the TCP/IP PI suite. A similar protocol is interpreted in the Banyan VINES PI. See also the IPX and ISO IP protocols.
IPC	Interprocess Communication Protocol. A transport-level protocol in Banyan VINES, providing reliable message service and unreliable datagram service. Interpreted in the Banyan VINES PI suite.
IPX	Internet Protocol. Novell's implementation of Xerox Internet Datagram Protocol. Interpreted in the Novell NetWare PI suite.
ISO	International Organization for Standardization (or International Standards Organization). (1) A consortium that is establishing a suite of networking protocols; (2) The protocols standardized by that group.
ISODE	ISO Development Environment. Protocol for transmitting higher-level ISO protocols over a network whose lower levels are handled by TCP/IP. Interpreted in the TCP/IP and ISO PI suites.
ISO IP	The ISO standard Internet Protocol. Interpreted in the ISO PI suite.
KSP	Kiewit Stream Protocol. A transport protocol resembling TCP developed at Dartmouth College for the support of terminal emulators connected to AppleTalk networks; interpreted in the AppleTalk PI suite.
LAN	Local Area Network. The hardware and software used to connect computers together in a limited geographical area.
LAP	Link Access Protocol. The logical level protocol for AppleTalk. It exists in two variants: ELAP (for Ethernet) and LLAP (for LocalTalk networks). Interpreted in the AppleTalk PI.
LAPB	Link Access Procedure, Balanced. A subset of HDLC.
LLAP	See LAP.
LAT	Local Area Transport. The DECnet protocol that handles multiplexed terminal (keyboard and screen) traffic to and from timesharing hosts. Interpreted in the DECnet PI suite.

LLC	Logical Link Control. A protocol that provides connection control and multiplexing to subsequent embedded protocols; standardized as IEEE 802.2 and ISO/DIS 8802/2.
LOOP	Loopback protocol. A protocol under Ethernet for sending diagnostic probe messages.
LSA	Lost Subarea. An SNA error condition.
LU 6.2	Logical Unit 6.2. A subset of the SNA protocols used for peer-to-peer communications between computers.
LUSTAT	Logical Unit Status. An SNA message used to send status information.
MAC	Medium Access Control. The protocol level that describes network management frames sent on the 802.5 token ring. Most MAC frames are handled transparently by the network adapter.
Mail Service	Protocol used (in conjunction with StreetTalk) for the transmission of messages in the VINES distributed electronic mail system. Interpreted in the Banyan VINES PI suite.
Manchester encoding	A data encoding technique that uses a transition at the middle of each bit period that serves as a clock and also as data.
MAP	Manufacturing Automation Protocol. A multilayer networking protocol developed primarily by General Motors for manufacturing control applications.
Matchmaker	Protocol used by the VINES service that provides high-level program-to-program communication, including translation as necessary to match the conventions of sender's and receiver's formats. Matchmaker is descended from XNS Courier. Interpreted in the Banyan VINES PI suite.
MAU	Multiple Access Unit (also Medium Attachment Unit). The wiring concentrator or transceiver used for attaching stations connected to the network.
MIB	Management Information Data Base. The structured database of network statistical information used by the SNMP and CMIP protocols.
MOP	Maintenance Operations Protocol. A protocol under DECnet for remote testing and problem diagnosis. Interpreted in the DECnet PI suite.
MOUNT	A protocol developed by Sun Microsystems that provides request access checking and user validation. It is used in conjunction with NFS. Interpreted in the Sun PI suite.

multicast	(1) A message directed to a group of stations on a network or collection of networks (contrast with broadcast). (2) A destination address that designates such a subset.
N(R)	Receive sequence number. An LLC or HDLC field for I frames that indicates the sequence number of the next frame expected; all frames before N(R) are thus implicitly acknowledged.
N(S)	Send sequence number. An LLC or HDLC field for I frames that indicates the sequence number of the current frame within the connection.
NBP	(1) Name-Binding Protocol. Used in AppleTalk networks to permit network users to use character names for network services and sockets. NBP translates a character-string name within a zone into the corresponding socket address. Interpreted in the AppleTalk PI suite. (2) NetBios Protocol. Used in 3Com 3+ Open software. Interpreted in the XNS PI suite.
NC	Network Control. An SNA subprocess.
NCP	NetWare Core Protocol. Novell's application-level protocol for the exchange of commands and data between file servers and workstations. Interpreted in the Novell NetWare PI suite.
ND	Network Disk. A protocol within the Sun NFS family used to access virtual disks located remotely across the network. Interpreted in the TCP/IP PI suite.
NetBIOS	Network Basic I/O System. (1) A protocol implemented by the PC LAN Program to support symbolically named stations and the exchange of arbitrary data. (2) The programming interface (API) used to send and receive NetBIOS messages. There exist several different and incompatible implementations of NetBIOS, and separate PIs for them, including the IBM and the TCP/IP PI suites.
NETBLT	Network Block Transfer. A protocol within an earlier version of TCP/IP. Not interpreted in the TCP/IP PI suite.
NetWare	The networking system designed by Novell Inc. and the protocols used therein.
Network Management	A protocol transmitted by a Matchmaker frame in Banyan VINES.
NEU	Network Extension Unit. A concentrator and repeater for StarLAN networks.

NFS	Network File System. A protocol developed by Sun Microsystems for requests and responses to a networked file server. Interpreted in the Sun PI suite.
NHU	Network Hub Unit. A concentrator and repeater for StarLAN networks.
NICE	Network Information and Control Exchange. The DECnet protocol for network management. Interpreted in the DECnet PI suite.
NIF	Neighbor Information Frame. Used by stations on an FDDI ring to announce their addresses to downstream neighbors.
NSP	Network Services Protocol. The DECnet protocol that provides reliable message transmission over virtual circuits. Interpreted in the DECnet PI suite.
OpenNET	A networking system from the Intel Corporation that uses parts of the OSI standards and components of the Microsoft/IBM PC LAN program. Interpreted in the ISO PI suite.
OSI	Open Systems Interconnection. A generalized model of a layered architecture for the interconnection of systems.
packet	The multi-byte unit of data transmitted at one time by a station on the network. Synonymous with Frame.
PAP	Printer Access Protocol. A protocol within AppleTalk that uses ATP XO commands to create a stream-like service for communication between user stations and the Apple LaserWriter or similar stream-based devices. Interpreted in the AppleTalk PI suite.
PC*I	Personal Computer Integration. Data General's nomenclature for their networking system. Protocols used include the ISO IP and TP4 levels and the Microsoft/IBM PC LAN program SMB protocols. Interpreted in the ISO PI suite.
PCF	Physical Control Fields. The part of the token ring DLC header that includes the AC and FC fields.
PDU	Protocol Data Unit. The data delivered as a single unit between peer processes on different computers.
PEP	Packet Exchange Protocol. A protocol within the XNS family used to exchange datagrams. Interpreted in the XNS/MS-Net PI suite.
PI	Protocol Interpreter. A program that knows the frame format and transaction rules of a communications protocol and can decode and display frame data.

PMAP	Port Mapper. A protocol developed by Sun Microsystems for mapping RPC program numbers to TCP/IP port numbers. Interpreted in the Sun PI suite.
PUP	PARC Universal Packet. A type of Ethernet packet formerly used at the Xerox Corporation's Palo Alto Research Center. Interpreted in the XNS/MS-Net and the TCP/IP PIs but not included in their protocol diagrams since no longer in regular use.
RARP	Reverse Address Resolution Protocol. A protocol within TCP/IP for finding a node's IP address given its DLC address. Interpreted in the TCP/IP PI suite.
RDP	Reliable datagram protocol. A protocol within an earlier version of TCP/IP. Not interpreted in the TCP/IP PI suite.
REJ	Reject. An LLC frame type that requests retransmission of previously sent frames.
REM	Ring Error Monitor. A station on the 802.5 token ring network that collects MAC-level error messages from the other stations.
RFC	Request For Comment. Designation used in DoD/TCP protocol research and development.
RG-58	The designation for 50-ohm coaxial cables used by Cheapernet (thin Ethernet).
RG-59	The designation for 75-ohm coaxial cables used by PC Network (broadband).
RG-62	The designation for 93-ohm coaxial cables used by ARCNET.
RGBI	Red-Green-Blue-Intensity. An interface used for attaching a color monitor to a personal computer; DB-9 connectors are typically used.
RH	Request/response header. An SNA control field prior to a Request Unit or Response unit.
RI	Routing Information. A protocol at the logical link level for devices operating on the token ring. Interpreted by the token ring and Ethernet Distributed Sniffer™ System independent of other PIs.
RII	Routing Information Indicator. If the first bit in the source address field of a token ring frame is 1, then the data field begins with Routing Information. Interpreted by the token ring and Ethernet Distributed Sniffer™ System independent of other PIs.

RIP	Routing Information Protocol. A protocol within the XNS and TCP/IP families used to exchange routing information among gateways. Interpreted in the XNS PI suite and in the TCP/IP PI suite.
RJ-45	The designation for the 8-wire modular connectors used for StarLAN and 10BASE-T networks. It is similar to, but wider than, the standard (RJ-11) telephone modular connectors.
RMS	Resource Management System. A set of protocols used by Datapoint to communicate from client stations to servers.
RNR	Receive Not Ready. An LLC and HDLC command or response indicating that transmission is blocked.
Router	(1) An internet linking device operating at network layer 3. (2) A protocol transmitted by a Matchmaker frame in Banyan VINES.
RPC	Remote Procedure Call. A protocol for activating functions on a remote station and retrieving the result. Interpreted in the Sun PI suite. A similar protocol exists in Xerox XNS.
RPL	Remote Program Load. A protocol used by IBM on the IEEE 802.5 token ring network to download initial programs into networked stations. Interpreted in the IBM PI suite.
RPS	Ring Parameter Server. A station on a token ring network that maintains MAC-level information about the LAN configuration such as ring numbers and physical location identifiers.
RR	Receive ready. An LLC non-data frame indicating readiness to receive data from the other station.
RS-232C	Recommended Standard 232. EIA standard defining electrical characteristics of the signals in the cables that connect a DTE and a DCE.
RSTAT	Remote status. A protocol with the Sun NFS family used to exchange statistics on network activity. Interpreted in the Sun PI suite.
RTMP	Routing Maintenance Protocol. Used in AppleTalk networks to allow bridges or internet routers dynamically to discover routes to the various networks of an internet. A node that is not a bridge uses a subset of RTMP (the RTMP stub) to determine the number of the network to which it is connected and the node IDs of bridges on its network. Interpreted in the AppleTalk protocol interpreter.

RTP	Routing Update Protocol. Used to distributed network topology information. Interpreted in the Banyan VINES PI suite.
RU	Request Unit/Response unit. The part of an SNA frame after the RH that contains the details of a request or its response.
RUnix	Remote Unix. A protocol atop TCP/IP for issuing remote requests over the network to a UNIX host.
S	Supervisory. An LLC, HDLC, or SDLC frame type used for control functions.
SABM	Set Asynchronous Balanced Mode. An LLC non-data frame requesting the establishment of a connection over which numbered I frames may be sent.
SABME	Set Asynchronous Balanced Mode (Extended). SABM with two more bytes in the control field. Used in LAPB.
SAP	Service Access Point. (1) A small number used by convention or established by a standards group, that defines the format of subsequent LLC data; a means of demultiplexing alternative protocols supported by LLC. (2) Service Advertising Protocol; Used by NetWare servers to broadcast the names and locations of servers and to send a specific response to any station that queries it.
SBI	Stop Bracket Initiation. An SNA message sent to request that the other station not initiate any more brackets.
SC	Session Control. An SNA subprocess for establishing and maintaining connections.
SCP	Session Control Protocol. The DECnet protocol concerned with the establishment of virtual circuits over which NSP transfers data; interpreted in the DECnet PI suite.
SDLC	Synchronous Data Link Control. An older serial communications protocol that was the model for LLC and with which it shares many features.
SESSION	Name for the session-level protocol in the ISO series, interpreted in the ISO PI suite.
Sever	A protocol transmitted by a Matchmaker frame in Banyan VINES.
SIF	Status Information Frame. Used by stations on an FDDI ring to exchange information about station configuration and operating parameters.
SIG	Signal. A high-priority SNA message used to request permission to send.

SMB	<p>Server Message Block. A message type used by the IBM PC LAN Program to make requests from a user station to a server and receive replies. Many of the functions are similar to those made by an application program to DOS or to OS/2 running on a single computer.</p> <p>SMB is part of the protocol family that for DOS machines is called MS-NET and for OS/2 machines is called The LAN Manager. Under the IBM PC LAN Program, SMBs are sent as data within NetBIOS frames, but in other context may be transported differently. The OS/2 version of SMB contains extensions not present in the DOS version. Both versions are interpreted in the IBM , XNS, TCP/IP, ISO, DECnet, and Banyan VINES PI suites.</p>
SMT	<p>Station Management. Provides ring management, connection management, and SMT frame services for an FDDI ring.</p>
SMTP	<p>Simple Mail Transfer Protocol. A protocol within TCP/IP for reliable exchange of electronic mail messages. Interpreted in the TCP/IP PI suite.</p>
SNA	<p>Systems Network Architecture. A complex set of protocols used by IBM for network communications, particularly with mainframe computers. Interpreted in the IBM PI suite.</p>
SNAP	<p>Sub-Network Access Protocol (also sometimes called Sub-Network Access Convergence Protocol). An extension to IEEE 802.2 LLC that permits a station to have multiple network-layer protocols. The protocol specifies that DSAP and SSAP addresses must be AA hex. A field subsequent to SSAP identifies one specific protocol. Interpreted in the TCP/IP PI suite and the AppleTalk PI suite. (See RFC 1042 for further information on SNAP.)</p>
SNMP	<p>Simple Network Management Protocol. Interpreted in the TCP/IP PI suite .</p>
SNRM	<p>Set Normal Response Mode. Place a secondary station in a mode that precludes it from sending unsolicited frames. The primary station controls all message flow. Used in SDLC.</p>
SNRME	<p>Set Normal Response Mode (Extended). SNRM with two more bytes in the control field. Used in SDLC.</p>
socket	<p>A logically addressable entity or service within a node, serving as a more precise identification of sender or recipient.</p>
SPP	<p>Sequenced Packet Protocol. A virtual-circuit connection-oriented protocol in XNS.</p>

SPP	Sequenced Packet Protocol. (1) The XNS protocol that supports reliable connections using sequenced data; interpreted in the XNS PI suite. A variant called SPX is used in Novell NetWare. (2) The transport-level protocol that provides virtual connection service in Banyan VINES, based upon the protocol of the same name in XNS. Interpreted in the Banyan VINES PI suite.
SPX	Sequential Packet Exchange. Novell's version of the Xerox protocol called SPP. Interpreted in the Novell NetWare PI suite.
SQE	Signal Quality Error. The 802.3/Ethernet collision signal from the transceiver.
SQE TEST	The SQE signal generated by the transceiver at the end of a transmitted frame to check the SQE circuitry. Also known as <i>Heart Beat</i> in Ethernet.
SSAP	Source Service Access Point. The LLC SAP for the protocol used by the originating station.
SSCP	System Services Control Point. An SNA identification of communications management functions.
StarLAN	A network developed by AT&T Bell Labs and based upon a derivative of the CSMA/CD (Ethernet) network standard originally developed by Xerox; similar to (and often used interchangeably with) the IEEE 802.3 standard.
StreetTalk	Protocol used in Banyan VINES to maintain a distributed directory of the names of network resources. In VINES names are global across the internet and independent of the network topology. Interpreted in the Banyan VINES PI suite.
SUA	Stored Upstream Address. The network address of a token ring station's nearest upstream neighbor. Texas Instruments calls this the UNA (see Upstream Neighbor Address).
Talk	A protocol transmitted by a Matchmaker frame in Banyan VINES.
TC	Transmission Control. An SNA subprocess.
TCP	Transmission Control Protocol. The connection-oriented byte-stream protocol within TCP/IP that provides reliable end-to-end communication by using sequenced data sent by IP. Interpreted in the TCP/IP P.
TCP/IP	Transmission Control Protocol/Internet Protocol. A suite of networking protocols developed originally by the US Government for Arpanet and now used by several LAN manufacturers. The individual TCP/IP protocols are listed separately in this Glossary.

Telnet	Protocol for transmitting character-oriented terminal (keyboard and screen) data. Interpreted in the TCP/IP PI suite.
TFTP	Trivial File Transfer Protocol. A protocol within TCP/IP used to exchange files between networked stations. Interpreted in the TCP/IP PI suite.
TH	Transmission header. The initial part of an SNA frame immediately following the LLC header.
token	A small message used in some networks to represent the permission to transmit; it is passed from station to station in a predefined sequence.
token bus	A type of LAN where all stations can hear what any station transmits and where permission to transmit is represented by a token sent from station to station.
token ring	A type of LAN where stations are wired in a ring and each can directly hear transmissions only from its immediate neighbor. Permission to transmit is granted by a token that circulates around the ring.
TP	Transport-level Protocol. It exists in alternate forms, depending on the services it assumes are provided to it by the network level below it. TP 0 assumes that the connection is maintained at the lower level, while TP 4 assumes a connectionless network protocol, so that functionality for the establishment and maintenance of a connection are included in the transport protocol. Levels 0, 2, and 4 are interpreted in the ISO PI suite.
TRLR	Trailer format. Variant of IP in which the protocol headers follow rather than precede the user data.
TS	Transmission Services. An SNA subprocess.
UA	Unnumbered Acknowledgment. An LLC frame that acknowledges a previous SABME or DISC request.
UDP	User Datagram Protocol. A protocol within TCP/IP for sending unsequenced data frames not otherwise interpreted by TCP/IP.
UI	Unnumbered Information. An LLC, HDLC, or SDLC frame type used to send data without sequence numbers.
UNA	Upstream Neighbor Address. The network address of a token ring station's nearest upstream neighbor. IBM calls this the SUA (see Stored Upstream Address).
UNIX	A popular portable operating system written by AT&T.
VINES	Virtual Network Software. The networking operating system developed by Banyan Systems Inc., and the protocols used therein. Notable components are StreetTalk and MatchMaker.

VMTP	Versatile Message Transaction Protocol (proposed).
VTP	Virtual Terminal Protocol.
V.35	A CCITT wideband interface recommendation for WANs.
WAN	Wide Area Network. A collection of LANs, or stations and hosts, extending over a wide area that can be connected via common carrier or private lines. Typically, transmission speeds are lower on a WAN than on a LAN.
X.25	A CCITT recommendation that defines the standard communications protocol for access to packet-switched networks.
X.400	ISO standard protocol for electronic mail. Interpreted in the ISO PI suite.
XID	Exchange Identification. An LLC unnumbered frame type used to negotiate what LLC services will be used during a connection.
XNS	Xerox Network Systems. A family of protocols standardized by Xerox; in particular the Internet Transport Protocols.
X Windows	Protocol for the management of high-resolution color windows at workstations, originated by MIT, DEC, and IBM and subsequently transferred to a consortium of vendors and developers.
YP	Yellow Pages. A protocol developed by Sun Microsystems for implementing a distributed resource look-up database; similar in function to DNS. Interpreted in the Sun PI suite.
ZIP	Zone Information Protocol. Used in AppleTalk to maintain an internet-wide mapping of networks to zone names. ZIP is used by the Name-Binding Protocol (NBP) to determine which networks belong to a given zone. Interpreted in the AppleTalk PI suite.
Zone	In AppleTalk networks, a set of one or more networks within an internet, such that no network is a member of more than one zone.

Bibliography

General References

Martin, James. *Local Area Networks*. The Arben Group, 1989.

Meijer, Anton and Paul Peeters. *Computer Network Architectures*. Rockville, Maryland: Computer Science Press, 1982.

Stallings, William. *Local Networks*. Macmillan, 1990.

Tannenbaum, Andrew S. *Computer Networks*. 2d ed. Englewood Cliffs, New Jersey: Prentice Hall, 1989.

Network References

Ethernet and StarLAN

IEEE Standards for Local Area Networks: *Logical Link Control*. ANSI/IEEE Std 802.2-1985 (ISO/DIS 8802/2). IEEE publication number SH09712. Institute of Electrical and Electronics Engineers, 345 East 47th Street, New York, NY 10017.

IEEE Standards for Local Area Networks: *Carrier Sense Multiple Access with Collision Detection (CSMA/CD). Access Method and Physical Layer Specifications*. ANSI/IEEE Std 802.3-1985 (ISO/DIS 8802/3). IEEE publication number SH09738. Institute of Electrical and Electronics Engineers, 345 East 47th Street, New York, NY 10017.

The Ethernet: A Local Area Network. Data Link Layer and Physical Layer Specifications. ("The blue book.") Issued jointly by Digital Equipment Corporation, Maynard, MA, Intel Corporation, Santa Clara, CA, and Xerox Corporation, Stamford, CT. Version 2.0, November 1982. Available from: Hillary Cornell, Xerox Systems Institute, 475 Oakmead Parkway, Sunnyvale, CA 94086. (408) 737-4652.

Token Ring

Haugdahl, J. Scott. *Inside the Token Ring*. Architecture Technology Corporation, 1986. P.O. Box 24344, Minneapolis MN 55424.

IEEE Standards for Local Area Networks: *Logical Link Control*. ANSI/IEEE Std 802.2-1985 (ISO/DIS 8802/2). IEEE publication number SH09712. Institute of Electrical and Electronics Engineers, 345 East 47th Street, New York, NY 10017.

IEEE Standards for Local Area Networks: *Token Ring Access Method*. ANSI/IEEE Std 802.5-1985 (ISO/DP 8802/5), IEEE publication number SH09944.

TMS380 Adapter Chipset User's Guide. Texas Instruments Incorporated, publication number SPWU001.

Token Ring Network Architecture Reference. IBM Corporation, publication number 6165877.

Token Ring Network PC Adapter Technical Reference. IBM Corporation, publication number 69X7713.

FDDI

American National Standard for Information Systems. *Fiber Distributed Data Interface (FDDI) - Token Ring Media Access Control (MAC)*. ANSI X3.139-1987. American National Standards Institute, 1987.

American National Standards for Information Systems. *Fibre Data Distributed Data Interface (FDDI) - Token Ring Physical Layer Medium Dependent (PMD)*. ANSI X3.166-1990. American National Standards Institute, 1990

A Primer to FDDI: Fiber Distributed Data Interface. Digital Equipment Corporation, 1991.

ARCNET

ARCNET Designer's Handbook. Datapoint Corporation, publication number 61610-01.

Attached Resource Computer: Simplified User's Guide. Datapoint Corporation, document number 50298, revision number 1.

Concepts of ARC Local Networking. Datapoint Corporation, document number 50694.

Herman, Mort. "LAN controller regulates token-passing traffic." *Electronic Design* (December 22, 1983): 139-144.

Local Area Network Controller LANC. Standard Microsystems Corporation, document number 4/83-2.5M.

Murphy, John A. "Token-passing protocol boosts throughput in local networks," *Electronics* (September 8, 1982).

Shustek, Leonard J. "A Client-Server Protocol for Local Area Networks." *Systems and Software* (March 1984): 127-131.

IBM PC Network

Berry, Paul. *Operating the IBM PC Network*. Berkeley, CA: SYBEX, Inc., 1986.

Cooper, Edward. *Broadband Network Technology: An Overview for the Data and Telecommunications Industry*. Englewood Cliffs, New Jersey: Prentice-Hall, 1984.

Cunningham, John E. *Cable Television*, Second Edition. Indianapolis: Howard W. Sams & Co., 1980.

Hewlett Packard. *Cable Television System Measurements Handbook*. Santa Rosa, CA: Hewlett Packard Co., 1977.

National Cable Television Association. *Standards of Good Engineering Practices for Measurement on Cable Television Systems, Distribution System*. Washington, D.C.: NCTA, 1977.

PC Network Technical Reference. IBM Corporation, publication number 6322916.

Shrock, Clifford B. *No Loose Ends: The Tektronix Proof-of-Purchase Program for CATV*. Tektronix Application Note, 1973.

Simons, Ken. *Technical Handbook for CATV Systems*, 3rd edition. Hatboro, PA: General Instrument Corporation, Jerrold Division, 1968.

LocalTalk

Sidhu, Gursharan S., Richard F. Andrews and Alan B. Oppenheimer. *Inside AppleTalk*. Addison-Wesley Publishing Company, 1989.

Synchronous

Data Communication Networks Interfaces: Recommendations X.20—X.32, Red Book, Volume VIII-Fascicle VIII.3. Geneva: International Telecommunications Union-CCITT, 1985.

Protocol References

IBM

Advanced Program-to-Program Communication for the IBM Personal Computer. Programming Guide. IBM Corporation, publication number 61X3842.

Haugdahl, J. Scott. *Inside NetBIOS*. Architecture Technology Corporation, 1986. P.O. Box 24344, Minneapolis MN 55424.

Systems Network Architecture Reference Summary. IBM Corporation, publication number GA27-3136.

Novell NetWare

Sheldon, Tom. *Novell NetWare: The Complete Reference*. Berkeley, California: McGraw-Hill, 1989.

XNS

Internet Transport Protocols. Xerox Systems Integration Standard X.S.I.S. 028112, December 1981.

TCP/IP

Cerf, V.G. and R.E. Kahn. "A Protocol for Packet Network Interconnection." *IEEE Trans. Commun.* COM-22:637—648 (May 1974).

Comer, Douglas E. *Internetworking With TCP/IP: Principles, protocols, and Architecture*. Englewood Cliffs, New Jersey: Prentice-Hall, 1988.

DDN Protocol Handbook.

Vol. 1: DOD Military Standard Protocols. NIC-5004.

Vol. 2: DARPA Internet Protocols. NIC-5005.

Vol. 3: Supplement. NIC-5006.

US Defense Communications Agency, December 1985. Available from: DDN Network Information Center, DDN Network Information Center, SRI International, Room EJ291, 333 Ravenswood Avenue, Menlo Park, CA 94025 (800) 235-3155; (415) 859-3695. NIC@SRI-NIC.ARPA; or from Defense Technical Information Center, Cameron Station, Alexandria, VA 22314 (202) 274-7633.

OSI

Day, J.D. and H. Zimmerman. "The OSI Reference Model." *Proceedings of the IEEE* 71 (1983): 1334-1340.

Henshall, J. and A. Shaw. *OSI Explained. End to End Computer Communication Standards*. Chichester, England: Ellis Horwood, 1988.

Linnington, P.F. "Fundamentals of the Layer Service Definitions and Protocol Specifications." *Proceedings of the IEEE* 71 (1983): 1341-1345.

Rose, Marshall T. *The Open Book: A Practical Perspective on OSI*. Englewood Cliffs, New Jersey, 1990.

Zimmerman, H. "OSI Reference Model—The ISO Model of Architecture for Open Systems Interconnection." *IEEE Trans. Commun.* COM-28:425—432 (April 1980).

DECnet

Malamud, Carl. *DEC Networks and Architectures*. New York: McGraw-Hill Book Company, 1989.

AppleTalk

Sidhu, Gursharan S., Richard F. Andrews and Alan B. Oppenheimer. *Inside AppleTalk*. Addison-Wesley Publishing Company, 1989.

X Windows

Scheifler, Robert, James Gettys, and Ron Newman. *X Window System, Library, and Protocol Reference*. Digital Press, 1988.

X.25

Deasington, R.J. *X.25 Explained: Protocols for Packet Switched Networks*. 2d ed. Chichester, England: Ellis Horwood, 1988.

Dhas, C.R. and V.K. Konangi. "X.25: An Interface to Public Packet Networks." *IEEE Commun. Magazine* IT-22 (1976): 118-125.

The X.25 Protocol and Seven Other Key Protocols. Belmont, California: Lifetime Learning Publications.

SNA

IBM Systems Network Architecture Formats. IBM Corporation, publication number GA27-3136-10

SMT

A Primer to FDDI: Fiber Distributed Data Interface. Digital Equipment Corporation, 1991.

SNIFFER® NETWORK ANALYZER

INDEX



Index

Numerics

- 10BASE2 1-3
- 10BASE5 1-3, 1-25
 - address format 1-7
- 10BASE-T 1-3, 1-5
 - topology 1-5
- 1BASE5 1-3, 1-25
 - address format 1-7
- 22 AWG wire 1-34
- 24 AWG wire 1-25
- 802.3
 - address format 1-7, 1-11
 - IEEE standard 1-3, 1-25
- 802.5
 - address format 1-11
 - IEEE standard 1-12

A

- abort sequence
 - frame field 1-36
- AC
 - frame field 1-15
- access control
 - ARCNET 1-27
 - CSMA/CA 1-35
 - CSMA/CD 1-6, 1-35
 - Ethernet 1-6, 1-35
 - LLAP 1-6, 1-35
 - LocalTalk 1-6, 1-35
 - PC Network 1-6, 1-33, 1-35
 - StarLAN 1-6, 1-26, 1-35
 - token ring 1-13
- access control byte 1-14
- access unit 1-12, 1-13
- acknowledged connectionless service 1-11
- acknowledgement packet 1-36
- ACSE 2-19
- active monitor 1-14, 1-16

- adapter
 - ARCNET 1-27, 1-28
 - Ethernet 1-5
 - LocalTalk 1-34
 - PC Network 1-32
 - StarLAN 1-25
 - synchronous 1-20
 - token ring 1-13
- address
 - 802.3 format 1-11
 - 802.5 format 1-11
 - ARCNET network 1-27
 - Ethernet format 1-11
 - format 1-7
 - frame field 1-21, 1-36
 - order in which transmitted 1-11
 - PC Network format 1-11, 1-34
 - PC Network original format 1-11
 - recognized 1-15
 - StarLAN format 1-11, 1-26
 - station 1-7
 - token ring format 1-11, 1-16
 - two stations with same address 1-29
 - WAN/synchronous 1-21
 - Address Resolution Protocol 2-14, 2-26
 - ADSP 2-28
 - Advanced Research Projects Agency 2-12
 - AFP 2-27
 - AFRP 2-9
 - agency code
 - frame field 1-16
 - alert
 - frame field 1-30
 - amplitude modulation 1-20
 - ANSI/IEEE 802.3
 - Ethernet standard 1-3
 - StarLAN standard 1-25
 - ANSI/IEEE 802.5
 - token ring standard 1-12
 - AppleTalk
 - Phase 1 2-27
 - Phase 2 2-27
 - protocol interpreter suite 2-27
 - AppleTalk Address Resolution Protocol 2-29

AppleTalk Data Stream Protocol 2-28

AppleTalk Filing Protocol 2-27

AppleTalk Session Protocol 2-28

AppleTalk Transaction Protocol 2-28

application layer 2-4

architecture

 Ethernet 1-3

 LocalTalk 1-34

 PC Network 1-31

 StarLAN 1-25

 synchronous 1-20

 token ring 1-12

ARCNET

 advantages 1-29

 architecture 1-26

 messages 1-29, 1-30

 nodal priority scheme 1-28

 speed 1-27

 standard 1-26

 token bus LAN 1-26

 topology 1-26, 1-27

ARCNET Fragmentation Protocol 2-9

ARP 2-14, 2-26

ARPANET

 Advanced Research Projects Agency 2-12

ASP 2-28

Association Control Service Element 2-19

ATP 2-28

ATP XO 2-28

Attachment Unit Interface cable 1-5

Attachment Unit Interface, see AUI 1-4

attenuator 1-33

AUI 1-5

 cable 1-4, 1-5

 connector 1-5

B

backoff algorithm 1-6

bandwidth

 PC Network 1-32

Banyan VINES

 protocol interpreter suite 2-24

baseband 1-3, 1-20, 1-25, 1-27

bits per second

 propagation speed 1-3, 1-12, 1-25, 1-27, 1-34

bit-stuffing 1-36

BNC connector

 network transceiver 1-5

broadband 1-5, 1-32

bus

 network topology 1-3, 1-25, 1-33, 1-34

 token bus network topology 1-27

bytes per second

 propagation speed 1-3, 1-12, 1-25, 1-34

C

cable

 ARCNET 1-27

 AUI 1-4, 1-5

 cheapernet 1-4

 coaxial 1-4

 drop 1-4

 Ethernet 1-3, 1-4

 FDDI 1-18

 LocalTalk 1-34

 PC Network 1-32

 RG-58/U 1-4

 RG-59 1-33

 RG-62 1-27

 StarLAN 1-25

 thin Ethernet 1-4

 token ring 1-12, 1-13

 twisted-pair 1-3, 1-5, 1-34

carrier sense multiple access 1-6, 1-35

CATV 1-32

CCITT 2-32

 V-Series standards 1-20

 X-Series standards 1-20

channels

 PC Network 1-32

Cheapernet

 jumper on Sniffer analyzer's adapter card 1-5

cheapernet

 cable 1-4

class 1

 connection-oriented networks 2-19

class 2

 intermediate networks 2-19

class 4

 connection-less networks 2-19

clear-to-send packet 1-35, 1-36

CLNS 2-19

CMOT 2-15

CMT 1-19, 2-36

- coaxial
 - cable 1-4
 - N-series connector 1-4
 - collision
 - detection 1-6
 - fragment 1-6
 - transceiver's role in detection 1-6
 - collision detect signal 1-6
 - Collisions 1-6
 - Command Terminal 2-22
 - Common Management and Information Services Protocol (CMIP) over TCP 2-15
 - Connectionless Network Service Protocol 2-19
 - connection-oriented service 1-11
 - connector
 - AUI 1-5
 - BNC 1-5
 - box 1-34
 - DB-15 1-4, 1-5
 - DB-9 1-12, 1-34
 - DIN-8 1-34
 - F 1-33
 - IBM data 1-12
 - N-series coaxial 1-4
 - RJ-11 1-12
 - RJ-18 1-25
 - RJ-45 1-5, 1-25
 - RS-232 1-20
 - T 1-4
 - connector box 1-35
 - contention 1-14, 1-16
 - control
 - frame field 1-9, 1-16, 1-22
 - count
 - frame field 1-31
 - Courier 2-11, 2-25
 - CRC
 - frame field 1-31
 - CSMA/CA
 - access control 1-6, 1-35
 - CSMA/CD
 - access control 1-6, 1-35
 - CTERM 2-22, 2-23
 - cyclic redundancy check 1-31
- D**
- DAP 2-22
 - data
 - frame field 1-7, 1-9, 1-15, 1-31, 1-37
 - Data Access Protocol 2-22
 - data circuit-terminating equipment 1-20
 - data communications equipment 1-20
 - data format
 - ARCNET 1-31
 - Ethernet vs. IEEE 802.3 1-8
 - LocalTalk 1-37
 - token ring 1-15
 - data frame 1-8, 1-31
 - CRC field 1-31
 - data field 1-7, 1-15, 1-31, 1-37
 - format "normalized" by 1-31
 - length 1-31
 - data terminal equipment 1-20
 - Datagram Delivery Protocol 2-29
 - Datapoint Corporation 1-26
 - system codes 1-31
 - DB-15 connector 1-4
 - network transceiver 1-5
 - DB-9 connector 1-12, 1-34
 - DCE 1-20
 - DD 2-29
 - DECnet
 - OSI 2-21
 - Phase IV 2-21
 - Phase V 2-21
 - protocol interpreter suite 2-21
 - DECnet Routing Protocol 2-23
 - DECWindows 2-30, 2-36
 - delimiter
 - token ring 1-14
 - destination address
 - frame field 1-7, 1-15, 1-30, 1-36
 - detail view
 - display FS byte 1-15
 - Digital Network Architecture 2-21
 - DIN-8 connector 1-34
 - directional coupler 1-33
 - DISC
 - HDLC protocol 1-23
 - LLC protocol 1-10
 - disconnect
 - HDLC unnumbered frame 1-23
 - LLC unnumbered frame 1-10

disconnected mode
 HDLC unnumbered frame 1-23
 LLC unnumbered frame 1-10

DM
 HDLC protocol 1-23
 LLC protocol 1-10

DNA 2-21

DNS 2-13

Domain Name Service 2-13

drop cable 1-4, 1-5, 1-6

DRP 2-23

DSAP
 frame field 1-9, 1-16
 LLC protocol 1-9

DTE 1-20

E

early token release algorithm 1-14

ECF 2-37

Echo 2-9, 2-11, 2-29

EDEL
 frame field 1-14

ELAP
 Ethernet 2-29

ending delimiter
 frame field 1-14

End-System to Intermediate-System Routing 2-19

enquiry packet 1-36

equalizer 1-33

Error 2-9, 2-11

error
 control 2-22
 handling 1-28
 recovery 1-11, 1-16
 reporting 1-16
 transmission 1-14

ESF 2-37

ES-IS Routing 2-19

Ethernet 1-3, 1-25
 access control 1-6, 1-35
 address format 1-11
 architecture 1-3
 ELAP 2-29
 LocalTalk over 2-27
 other systems 1-5
 RI field 1-7
 speed 1-3
 standard 1-3

 topology 1-3, 1-25, 1-32, 1-34

EtherTalk Link Access Protocol 2-29

Ethertype
 frame field 1-8
 vs. IEEE 802.3 1-8

exchange identification
 HDLC unnumbered frame 1-23
 LLC unnumbered frame 1-10

F

F connector 1-33

fault, hardware 1-29

FC
 frame field 1-15

FCS
 frame field 1-7, 1-15, 1-21, 1-36

FDDI 1-16, 1-17

 LLC 1-17

 MAC 1-17

 PHY 1-17

 PMD 1-17

 SMT 1-17

fiber optics 1-5

File Transfer Protocol 2-13

File Transfer, Access and Management 2-18

FIPS
 Ethernet standard 1-3
 StarLAN standard 1-25

flag
 frame field 1-36

forward traffic 1-32

FOUND 2-22

Foundation Services 2-22

Fragmentation Protocol 2-26

frame 1-13
 ARCNET format 1-30
 check sequence 1-7
 Ethernet format 1-7
 LocalTalk format 1-36
 PC Network format 1-7, 1-34
 StarLAN format 1-7, 1-26
 synchronous format 1-21
 token ring format 1-14

frame copied 1-15

frame reject
 HDLC unnumbered frame 1-23
 LLC unnumbered frame 1-10

-
- Frame Relay 2-34, 2-35
 - LMI 2-35
 - Frame Status byte, see FS 1-15
 - free token 1-13
 - frequency assignments
 - PC Network 1-32
 - frequency modulation 1-20
 - frequency-division multiplexing 1-32
 - frequency-shift-keying 1-32
 - FRMR
 - HDLC protocol 1-23
 - LLC protocol 1-10
 - FRP 2-26
 - FS
 - display in detail view 1-15
 - frame field 1-15
 - FSK 1-32
 - FTAM 2-18
 - FTP 2-13

 - G**
 - Gateway-to-Gateway Protocol 2-14
 - GGP 2-14
 - ground
 - network 1-4, 1-33

 - H**
 - handshake mechanism
 - LLAP 1-35
 - hardware
 - isolating fault in LAN equipment 1-29
 - HDLC 1-20, 2-33
 - address field 1-21
 - control field 1-22
 - frame format 1-21
 - frame types 1-22
 - information frame type 1-22
 - protocol 1-9
 - relation to LLC 1-9
 - supervisory frame type 1-22
 - headend 1-32
 - High-level Data Link Control 1-20, 2-33
 - High-level Data Link Control, see HDLC 1-9, 1-21
 - hub
 - StarLAN 1-25
 - hub unit 1-27
 - techniques for reducing the number needed 1-27

 - I**
 - IBM
 - protocol interpreter suite 2-5
 - IBM Data Connector 1-12
 - IBM LAN Support Program
 - NetBIOS 1-8
 - IBM Network Management Protocols 2-6
 - IBM PC LAN 2-19, 2-22, 2-25
 - IBM Token-Ring Network 1-12
 - IBMNM 2-6
 - ICMP 2-14
 - ICP 2-26
 - ID
 - frame field 1-36
 - IDP 2-7, 2-11
 - IEEE 802.2 2-6, 2-9, 2-14, 2-23, 2-26, 2-29
 - IEEE 802.3
 - address format 1-7, 1-11
 - data format 1-8
 - Ethernet standard 1-3
 - StarLAN standard 1-25
 - vs. Ethernet 1-8
 - IEEE 802.5
 - address format 1-11
 - token ring standard 1-12
 - impedance
 - network cable 1-4, 1-33
 - information
 - HDLC frame type 1-22
 - LAPB frame type 1-22
 - LLC frame type 1-9
 - SDLC frame type 1-22
 - insertion in ring 1-13
 - International Consultative Committee for Telephony and Telegraphy, see CCITT 1-20
 - Internet Control Message Protocol 2-14
 - Internet Control Protocol 2-26
 - Internet Datagram Protocol 2-7, 2-11
 - Internet Packet Exchange 2-7, 2-9
 - Internet Protocol 2-14, 2-26
 - Internetwork Protocol 2-19
 - interpretation
 - protocols 2-3
 - Interprocess Communication Protocol 2-25
 - IP 2-14, 2-26
 - IPC 2-25
 - IPX 2-7, 2-9
-

- ISO 1-20
 - DECnet 2-21
 - International Standards Organization 2-18
 - LLC standard 1-8
 - Open Systems Interconnection 2-18
 - protocol interpreter suite 2-18
 - StarLAN standard 1-25
 - token ring standard 1-12
 - X.25 standard 1-20

- ISO 8327 2-19

- ISO 8473 2-19

- ISO 8571/4 2-18

- ISO 8650/2 2-19

- ISO 8802/2 2-19

- ISO 8823 2-19

- ISO 9041 2-19

- ISO Development Environment 2-20

- ISO IP 2-19

- ISO/DIS 8802/2 2-29

- ISODE 2-20

J

- jabber detection
 - transceiver 1-7

- jamming signal 1-6

- jumper
 - thick vs. thin Ethernet 1-5

K

- Kiewit Stream Protocol 2-29

- KSP 2-29

L

- LAN Manager 2-19

- LAP 2-29
 - on AppleTalk networks 2-29

- LAPB
 - address field 1-21
 - comparison to SDLC 1-21
 - control field 1-22
 - frame format 1-21
 - frame types 1-22
 - information frame type 1-22

- LAT 2-23

- LattisNet 1-5

- length
 - ARCNET cable segment 1-27

- Ethernet cable segment 1-4
- field in data frame 1-31
- frame field 1-8, 1-37
- LocalTalk network cable 1-35
- PC Network cable segment 1-33
- StarLAN cable segment 1-25
- synchronous cable segment 1-20
- token ring cable segment 1-12

- Link Access Procedure, Balanced, see LAPB 1-21

- Link Access Protocol 2-29

- link control
 - synchronous 1-24

- LLAP
 - access control 1-35
 - frame format 1-36
 - handshake mechanism 1-35
 - LocalTalk 2-29

- LLAP type
 - frame field 1-36

- LLC 1-20, 2-6, 2-9, 2-14, 2-19, 2-23, 2-26, 2-29
 - acknowledged connectionless service 1-11
 - connection-oriented service 1-11
 - FDDI 1-17
 - frame 1-16
 - frame format 1-21
 - frame types 1-9
 - header format 1-8
 - protocol 1-9
 - relation to HDLC 1-9
 - types of operation 1-11
 - unacknowledged connectionless service 1-11

- LMI 2-35

- Local Area Transport Protocol 2-23

- local code
 - frame field 1-16

- LocalTalk 2-27
 - access control 1-6, 1-35
 - architecture 1-34
 - LLAP 2-29
 - LLAP frame format 1-36
 - speed 1-34
 - standard 1-34
 - topology 1-34

- LocalTalk Link Access Protocol, see LLAP 1-6

- lockpost
 - cable connectors 1-4

- logical link control layer 2-4

- Logical Link Control, see LLC 1-3

- logical ring 1-3

- low-split 1-32

M

MAC frames 1-16
 MAIL 2-25
 Maintenance Operations Protocol 2-23
 Manchester encoding 1-7, 1-14
 manufacturer code
 802.3 address format 1-11
 802.5 address format 1-11
 agency code 1-16
 Ethernet address format 1-11
 PC Network address format 1-11, 1-34
 PC Network original address format 1-11
 StarLAN address format 1-11, 1-26
 token ring address format 1-11, 1-16
 Matchmaker 2-25
 MAU 1-12, 1-13
 medium attachment unit 1-4, 1-6
 multiple access unit 1-12
 twisted pair 1-5
 medium access control
 token ring 1-16
 medium attachment unit 1-4, 1-6
 mid-split 1-32
 mini-hub 1-27
 monitor
 active 1-16
 contention 1-14, 1-16
 standby 1-16
 MOP 2-23
 MOUNT 2-17
 MPR 1-5
 MS NET 2-19
 multiple access unit 1-12, 1-13
 multiport repeater 1-5

N

Name-Binding Protocol 2-28
 NBP 2-11, 2-28
 NCP 2-7, 2-8
 ND 2-17
 NETBIOS 2-6
 NetBIOS 1-8, 2-9, 2-13, 2-25
 protocol 2-11
 SAP 1-9
 NetBIOS header
 frame field 1-8
 NetWare Core Protocol 2-7, 2-8

network

802.3 address format 1-11
 802.5 address format 1-11
 bus topology 1-3, 1-25, 1-33, 1-34
 Ethernet address format 1-11
 PC Network address format 1-11, 1-34
 PC Network original address format 1-11
 ring topology 1-12
 star topology 1-12, 1-25
 StarLAN address format 1-11, 1-26
 token ring address format 1-11, 1-16
 topology 1-3, 1-5, 1-12, 1-25, 1-27, 1-32, 1-33, 1-34
 tree topology 1-33
 Network Basic I/O System, see NetBIOS 1-8
 Network Disk 2-17
 network extension unit 1-25
 Network File Services 2-7
 Network File System 2-16, 2-17
 Network Information and Control Exchange 2-22
 Network Information Services 2-17
 network layer 2-4
 Network Services protocol 2-22
 NFS 2-7, 2-16, 2-17
 NICE 2-22
 NIF 2-37
 NIS 2-17
 nodal priority scheme 1-28
 Novell NetWare
 protocol interpreter suite 2-7
 NR
 HDLC receive sequence number 1-22
 LLC receive sequence number 1-9
 NS
 HDLC send sequence number 1-22
 LLC send sequence number 1-9
 N-series coaxial connector 1-4
 NSP 2-22

O

offset
 frequency difference between two signals 1-32
 Open Systems Interconnection, see OSI 2-3
 order of transmission
 station address 1-11
 OS/2 2-19, 2-25
 OSI 2-3, 2-18
 DECnet 2-21

P

- P1
 - outer addressing 2-18
- P2
 - inner addressing 2-18
- Packet Assembler/ Disassembler Protocol 2-33, 2-37
- Packet Exchange Protocol 2-11
- PAD 2-33
- PAP 2-28
- PC Network
 - access control 1-6, 1-33, 1-35
 - adapter 1-32
 - address format 1-11
 - architecture 1-31
 - data frame format 1-8
 - frequency assignments 1-32
 - NetBIOS 1-8
 - original address format 1-11
 - physical interconnection problems 1-33
 - speed 1-32
 - standard 1-31
 - Sytek 1-31
 - topology 1-3, 1-25, 1-32, 1-34
- PEP 2-11
- physical layer 2-4
- PI 2-3
- PMAP 2-17
- PMF 2-37
- Point-to-Point Protocol 2-33
- Port Mapper 2-17
- PPP 2-33
- preamble
 - frame field 1-7, 1-36
- Presentation 2-19
- presentation layer 2-4
- Printer Access Protocol 2-28
- priority 1-13, 1-27
- Project Athena 2-30
- propagation speed
 - ARCNET 1-27
 - Ethernet 1-3
 - LocalTalk 1-34
 - PC Network 1-32
 - StarLAN 1-25
 - synchronous 1-20
 - token ring 1-12
- protocol
 - interpretation 2-3

- protocol identification header 1-16
- protocol interpreter suite
 - AppleTalk 2-27
 - Banyan VINES 2-24
 - DECnet 2-21
 - IBM 2-5
 - ISO 2-18
 - Novell NetWare 2-7
 - SMT 2-36
 - SUN 2-16
 - TCP/IP 2-12
 - X Windows 2-30
 - X.25 2-32
- protocol interpreter, see PI 2-3
- PUP Ethertype 1-8

Q

- QLLC 2-33
- Qualified Logical Link Control Protocol 2-33

R

- radio frequency transmission 1-32
- RAF 2-37
- RARP 2-14
- RDF 2-37
- receive not ready
 - HDLC supervisory frame 1-22
 - LLC supervisory frame 1-10
- receive pair 1-13
- receive ready
 - HDLC supervisory frame 1-22
 - LLC supervisory frame 1-10
- reconfiguration 1-28
- REJ
 - HDLC reject 1-22
 - LLC reject 1-10
- reject
 - HDLC supervisory frame 1-22
 - LLC supervisory frame 1-10
- Remote Procedure Call 2-17
- Remote Program Load 2-6
- Remote Unix 2-13
- repeater 1-27
 - network cable 1-4, 1-33
- request-to-send packet 1-35, 1-36
- resistive coupler 1-27
- return traffic 1-32

-
- Reverse ARP 2-14
 - RF 1-32
 - RG-58/U cable 1-4
 - RG-59 cable 1-33
 - RG-62 cable 1-27
 - RI
 - frame field 1-7
 - ring
 - architecture 1-12
 - logical interconnection 1-3, 1-27
 - token ring network topology 1-12
 - RIP 2-9, 2-11, 2-14
 - RJ-11 connector 1-12
 - RJ-18 connector 1-25
 - RJ-45 connector 1-5, 1-25
 - RMT 1-19, 2-36
 - RNR
 - HDLC receive not ready 1-22
 - LLC receive not ready 1-10
 - routing information
 - token ring network 1-15
 - routing information field 1-7
 - Routing Information Protocol 2-9, 2-11, 2-14
 - Routing Table Maintenance Protocol 2-28
 - Routing Update Protocol 2-26
 - RPC 2-17
 - RPL 2-6
 - RR 1-22
 - HDLC receive ready 1-22
 - LLC receive ready 1-10
 - RS-232 1-20
 - connector 1-20
 - RTMP 2-28
 - RTP 2-26
 - RUNIX 2-13

 - S**
 - SABM
 - LAPB protocol 1-23
 - SABME
 - LAPB protocol 1-23
 - LLC protocol 1-10
 - SAP 2-8
 - codes in LLC protocol 1-9
 - NetBIOS 1-9
 - SNA 1-9
 - SCP 2-22
 - SDEL
 - frame field 1-14
 - SDLC 1-20, 2-6
 - comparison to LAPB 1-21
 - control field 1-22
 - frame format 1-21
 - frame types 1-22
 - information frame type 1-22
 - Sequenced Packet Protocol 2-11, 2-25
 - Sequential Packet Exchange 2-7, 2-9
 - Server Message Block 2-6, 2-11, 2-13, 2-19, 2-22, 2-25
 - Service Advertising Protocol 2-8
 - Session 2-19
 - Session Control Protocol 2-22
 - session layer 2-4
 - set asynchronous balanced mode
 - HDLC unnumbered frame 1-23
 - set asynchronous balanced mode (extended)
 - HDLC unnumbered frame 1-23
 - LLC unnumbered frame 1-10
 - set normal response mode
 - HDLC unnumbered frame 1-23
 - set normal response mode (extended)
 - HDLC unnumbered frame 1-23
 - SIF 2-37
 - signal quality error
 - test 1-6
 - transceiver 1-6
 - signal splitter 1-33
 - Simple Mail Transfer Protocol 2-13
 - SMB 2-6, 2-11, 2-13, 2-19, 2-22, 2-25
 - SMT 1-19, 2-36, 2-37
 - CMT 1-19, 2-36
 - ECF 2-37
 - ESF 2-37
 - frame services 1-19
 - NIF 2-37
 - RAF 2-37
 - RDF 2-37
 - RMT 1-19, 2-36
 - SIF 2-37
 - SMTP 2-13
 - SNA 1-20, 1-21, 2-5
 - SAP 1-9
 - SNAP 1-16, 2-15, 2-23
 - SNDCP 2-33
 - SNMP 2-15

- SNRM
 - SDLC protocol 1-23
 - SNRME
 - SDLC protocol 1-23
 - SoftTalk 2-28
 - source address
 - frame field 1-7, 1-15, 1-30, 1-36
 - source routing
 - token ring network 1-15
 - speed
 - transmission, ARCNET 1-27
 - transmission, Ethernet 1-3
 - transmission, LocalTalk 1-34
 - transmission, PC Network 1-32
 - transmission, StarLAN 1-25
 - transmission, synchronous 1-20
 - transmission, token ring 1-12
 - SPP 2-7, 2-11, 2-25
 - SPX 2-7, 2-9
 - SQE
 - test 1-6
 - transceiver 1-6
 - SRF 2-37
 - SSAP
 - frame field 1-9, 1-16
 - LLC protocol 1-9
 - Standard Microsystems Corporation 1-26
 - standby monitor 1-16
 - star network topology 1-12, 1-25
 - StarLAN 1-3
 - access control 1-6, 1-26, 1-35
 - address format 1-11
 - architecture 1-25
 - speed 1-25
 - standard 1-25
 - topology 1-3, 1-25, 1-32, 1-34
 - StarLAN 10 1-5
 - starting delimiter
 - frame field 1-14
 - station address
 - 2-byte 1-7
 - 6-byte 1-7
 - 802.3 format 1-11
 - 802.5 format 1-11
 - ARCNET network 1-27
 - Ethernet format 1-11
 - PC Network format 1-11, 1-34
 - PC Network original format 1-11
 - StarLAN format 1-11, 1-26
 - token ring format 1-11, 1-16
 - StreetTalk 2-25
 - Sub-Network Access Convergence Protocol 2-26, 2-29
 - Sub-Network Access Protocol 1-16, 2-15, 2-23
 - Subnetwork Dependent Convergence Protocol 2-33
 - subvector 1-16
 - SUN
 - protocol interpreter suite 2-16
 - supervisory
 - HDLC frame type 1-22
 - LLC frame type 1-9
 - synchronous
 - adapter 1-20
 - architecture 1-20
 - frame format 1-21
 - phases of link control 1-24
 - speed 1-20
 - Synchronous Data Link Control, see SDLC 1-20
 - system code 1-31
 - system ID
 - frame field 1-31
 - Systems Network Architecture, see SNA 1-9
 - Sytek
 - standard 1-31
- ## T
- T connector 1-4
 - TCP 2-14
 - SNMP 2-15
 - TCP/IP 2-20, 2-30, 2-36
 - protocol interpreter suite 2-12
 - technical support viii
 - Telnet 2-13
 - TEST
 - HDLC protocol 1-23
 - LLC protocol 1-10
 - test probe
 - LLC unnumbered frame 1-10
 - TFTP 2-13
 - The 2-30, 2-36
 - thin Ethernet
 - cable 1-4
 - tilt compensator 1-33
 - token 1-27
 - free 1-13
 - reestablishment of rotation 1-28

- token bus
 - topology contrasted to token ring 1-27
- token frame 1-14
- token ring
 - address format 1-11
 - architecture 1-12
 - monitor 1-14
 - NetBIOS 1-8
 - network topology 1-12
 - RI field 1-7
 - speed 1-12
 - standard 1-12
 - topology contrasted to Ethernet 1-3
 - topology contrasted to LocalTalk 1-34
 - topology contrasted to PC Network 1-32
 - topology contrasted to StarLAN 1-25
- topology
 - ARCNET network 1-27
 - network 1-3, 1-5, 1-12, 1-25, 1-26, 1-32, 1-33, 1-34
- TOPS 2-27
- TP 2-19
- traffic
 - forward 1-32
 - return 1-32
- Trailer format 2-15
- transceiver 1-27
 - collision detection 1-6
 - Ethernet 1-4, 1-6
 - jabber detection 1-7
 - twisted pair 1-5
- translator 1-32
- Transmission Control Protocol 2-14
- transmission order
 - address 1-11
- transmission rate 1-27
 - Ethernet 1-3
 - LocalTalk 1-34
 - PC Network 1-32
 - StarLAN 1-25
 - synchronous 1-20
 - token ring 1-12
- transmit pair 1-13
- transport layer 2-4
- Transport Protocol 2-19
- tree network topology 1-33
- Trivial File Transfer Protocol 2-13
- TRLR 2-15
- twisted pair Ethernet 1-5
 - LattisNet 1-5
 - StarLAN 10 1-5
- twisted-pair cable 1-3, 1-5, 1-34
- twisted-pair MAU 1-5
- type code
 - frame field 1-30
- U**
- UA
 - HDLC protocol 1-23
 - LLC protocol 1-10
- UDP 2-14
- UI
 - HDLC protocol 1-23
 - LLC protocol 1-10
- unacknowledged connectionless service 1-11
- UNIX 2-16, 2-24
- unnumbered
 - HDLC frame type 1-23
 - LLC frame type 1-10
- unnumbered acknowledgement
 - HDLC unnumbered frame 1-23
 - LLC unnumbered frame 1-10
- unnumbered information
 - HDLC unnumbered frame 1-23
 - LLC unnumbered frame 1-10
- US Defense Communication Agency 1-8
- User Datagram Protocol 2-14
- V**
- V.35 1-20
- vampire tap 1-5
- Virtual Terminal Protocol 2-19
- V-Series standards 1-20
- VTP 2-19
- W**
- WAN 1-20
- wide area network 1-20
- X**
- X Windows
 - protocol interpreter suite 2-30, 2-36
- X.25 1-20, 1-21, 2-33
 - protocol interpreter suite 2-32
- X.400 2-18

Xerox

Etherypes 1-8

Xerox Network Systems 2-25

protocol interpreter suite 2-10

Xerox Network Systems Protocol 2-7, 2-9, 2-11

XID

HDLC protocol 1-23

LLC protocol 1-10

XNS 2-7, 2-9, 2-11, 2-25

protocol interpreter suite 2-10

X-Series standards 1-20

Y

Yellow Pages 2-17

YP 2-17

Z

ZIP 2-28

Zone Information Protocol 2-28



We simplify network complexity.™

Network General Corporation

4200 Bohannon Drive
Menlo Park, CA 94025 USA
+1 (415) 688-2700

Network General Europe

Belgicastraat 4
1930 Zaventem, Belgium
+32 (2) 725-6030

Network General Canada, Ltd.

2275 Lakeshore Blvd., West, 5th Floor
Etobicoke, Ontario Canada M8V 3Y3
+1 (416) 259-5022