

Chapter
1

INTRODUCTION

*Get on the
Fast Track!*



TM

**SYS-ED\
COMPUTER
EDUCATION
TECHNIQUES, INC.**

Objectives

You will learn:

- C Servers and workstations.
- C Network interface cards and drivers: ODI and NDIS.
- C Access protocols.
- C Industry standards: OSI, TCP/IP, SNA, et al.
- C The pros and cons of the most popular LAN technologies.
- C IEEE architectural layers.
- C Merging various LAN technologies into an efficiently supported infrastructure.

1 LAN Components

Hardware

- C Intel based systems.

- C LAN attachments/connections between workstations and servers and workstation to workstation.
 - NIU: Network Interface Units
 - TIU: Terminal Interface Unit
 - NIC: Network Interface Card

Software

- C Device Drivers

- C Operating System & Environments:
DOS, Windows3.11, 95/98, ME, NT/2000, and XP, OS/2, and UNIX

- C Network Operating Systems:
Novell and OS/2 LAN Server

- C Server and Workstations.

Wiring


- C Different types and categories.

Web Servers

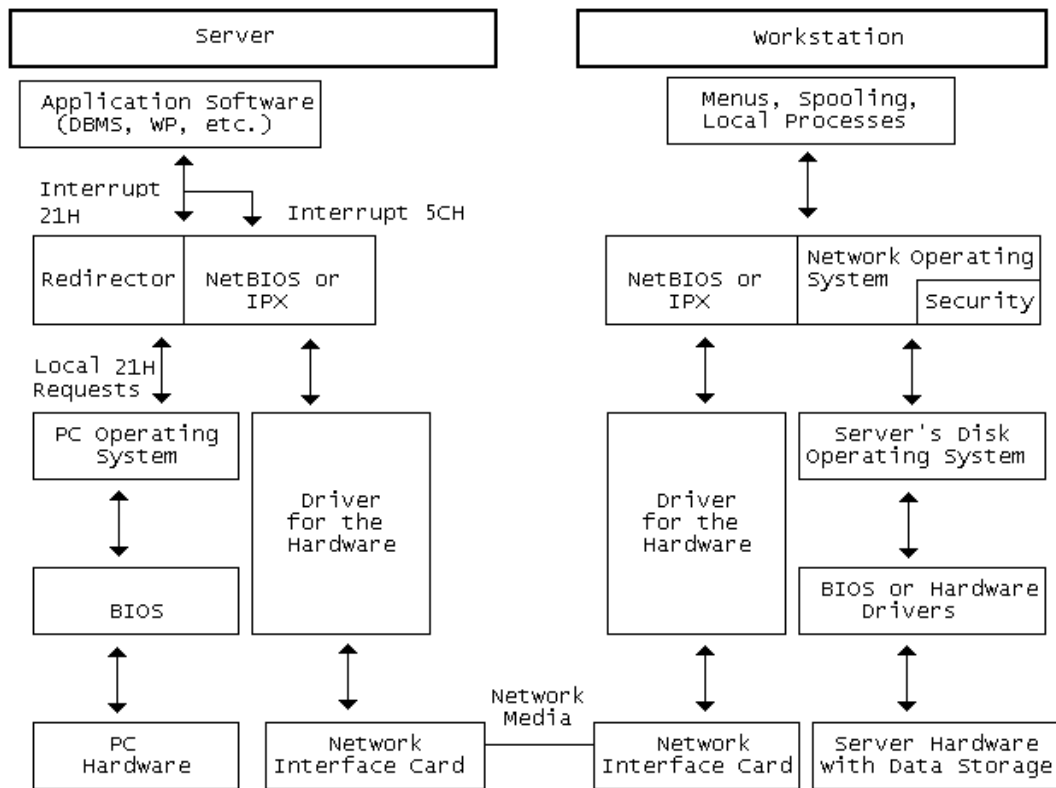
- C Internet Information Servers
- C WebSphere
- C WebLogic
- C Apache
- C JRun

2 Network Software/ Hardware Interaction

LAN Hardware Components



Network Software



3 Network and Terminal Attachment Units

The network interface is the connection between each node and the network link. Everything coming onto the network must be compatible with the network link, so if the signal from the node is not compatible, the network interface must translate or convert it.

In a distributed control network like a ring, where each station has some responsibility for keeping the network running, this signal conversion is done by the "network interface unit," or NIU.

Besides the NIU, the network interface includes the "tap," or "transceiver," which diverts traffic onto and off the network link, and the "drop link," the cable that carries data between the tap and the NIU. In a star network, with a centralized network controller, there are no NIUs and no drop links.

Everything on the other side of the network interface makes up the node, which consists of the "station" and the "station link." The latter is the cable that connects the station to the network interface. Your computer or work station is the "station."

The station might have a built in "terminal interface unit" (TIU), which adapts the signal from your terminal to the network interface. If the TIU isn't built-in, you will see a box and cable, the "terminal link," running between it and your station.

4 Cabling and Media Access Control

Establishes most of the LAN's hardware characteristics:

- C Cable type
- C Topology
- C Access scheme
- C Data transmission (bit) rate

There are 3 major standard protocols for LAN cabling and media access control, and therefore accordingly corresponding categories of network interface cards: Ethernet, Token Ring, and ARCnet.

Ethernet

Ethernet based LANs enable you to interconnect a wide variety of equipment, including UNIX computers, Apple computers, IBM PCs, and IBM clones. You can buy Ethernet cards from dozens of competing manufacturers. Ethernet comes in three varieties (ThinNet, UTP, ThickNet) depending on the thickness of the cabling you use. Thicknet can span a great distance, but they are much more expensive. Ethernet operates at a rate of 100 million bits per second (100 mbps).

Token Ring

Token Ring is the central component of IBM's local and wide area network architecture. IBM provides optional Token-Ring connections on its mainframe computer hardware and software to make PCs and mainframe act as peers on the same network.

IBM offers Token Ring products that operate at either 4 or 16 megabits per second. Several other companies make equipment compatible with IBM Token Ring, and a few companies make Token Ring hardware that operates at different rates or that uses fiber optics.

ARCnet

ARCnet is one of the oldest types of LAN hardware. It was originally a proprietary scheme of the Datapoint Corporation, but today manufacturing make ARCnet compatible cards. ARCnet is known for solid reliability, and ARCnet cable/adaptor problems are easy to diagnose. ARCnet costs less than EtherNet. ARCnet operates similar to Token Ring, but at the slower rate of 2.5 million bits per second (2.5 mbps).

5 NDIS and ODI**NDIS: Network Driver Interface Specification**

Developed jointly by 3COM and Microsoft, NDIS is a cornerstone of the LAN Server and LAN Manager network operating system products.

A network manufacturer can make its boards work with these products by supplying NDIS-compliant software drivers with the boards.

ODI: Open Datalink Interface

Developed jointly by Novell and Apple computer, ODI performs many of the same functions as NDIS, but NDIS and ODI are incompatible with each other.

A network manufacturer can make its boards work with Netware, the most popular network operating system, by supplying ODI-compliant software drivers with the boards.

Most network adapter manufacturers supply both NDIS and ODI drivers with their products.

6 Interface Cards

The interface card, also known as a NIC or adapter card, takes data from the PC, puts it into the appropriate format, and sends it over the cable to another LAN interface card. This card receives the data, puts it into a form the PC understands, and sends it to the PC.

The interface card's role can be broken into eight tasks:

1. Host-to-Card Communications
2. Buffering
3. Packet Formation
4. Parallel-to-Serial Conversion
5. Encoding and Decoding
6. Cable Access
7. Hand Shaking
8. Transmission and Reception

These steps get the data from the memory of one computer onto the cable, and reversing the steps gets the data off the cable and into the memory of another computer.

7 Methods of Transmission

There are four ways to move data between the PC's memory to the network interface card (NIC):

- C I/O: Input/Output
- C Direct Memory Access
- C Shared Memory
- C MCA: Bus Mastering

7.1 Input/Output (I/O)

I/O is the simplest method. Within this category, the most important types are memory-mapped I/O and program I/O.

Memory-Mapped I/O

In a memory-mapped I/O transfer, the host CPU assigns some of its memory space to the I/O device, in this case the network interface card. Out of the 640KB of RAM that is available for DOS PCs, a few kilobytes are allocated to the network card.

This memory is then treated as if it were the PC's main memory. No special instructions of the CPU are needed to get data from the card since it is like taking data from one part of memory to another.

Program I/O

With program I/O, the CPU is given a set of special instructions to handle the input/output functions. These instructions can be built into the chip or come with software. To send data a request is sent from the network interface card to the CPU. The CPU then moves the data from the card over the bus to the main memory.

Because the CPU is required to handle the I/O process, it cannot perform other tasks while data is being transferred. This makes it slow.

7.2 DMA: Direct Memory Access

All Intel-based computers come with a DMA controller chip that takes care of transferring data from a input/output device to the PC's main memory. In this way the PC's CPU does not have to get involved in the transfer

The controller or processor in the interface card sends a signal to the CPU indicating it wants to transfer information. The CPU also dictates the appropriate memory address at which the DMA controller is to begin putting data in memory. The CPU then relinquishes control of the PC bus to the DMA controller.

Once the DMA controller has command of the bus, it takes the data from the card and places it directly in memory. After all the data is in memory, the DMA controller returns control of the bus to the CPU and tells it how much data has been put in memory.

DMA is generally faster than I/O because the DMA controller removes work from the CPU, so the CPU can perform other functions while data transfer is taking place.

The disadvantage is the CPU cannot access memory while the DMA controller is working.

7.3 Shared Memory

In shared memory, part of host PC's memory is shared by the network interface card's processor.

Shared memory is a very fast transfer method, since no buffering on the card is required. Because the card and the PC do their work on the data in the same place, no transfer is necessary. Although shared memory is the fastest method of moving data between the network interface card and the PC, it is more difficult to build than DMA or I/O. Shared memory takes more PC RAM than other methods.

7.4 Bus Mastering

Bus mastering is a high performance technique which provides for enhanced speeds beyond the 40 megabits per second.

Both EISA and MCA have the capability to implement bus mastering.

8 PC to NIC Buffering

The buffer is a storage place that holds data as it is moving into and out of the NIC.

A buffer is necessary because some parts of data transfer are slower than others. For example, data comes into the card faster than it can be converted from a serial or parallel format, depacketized, read, and sent. This is true in both directions.

To compensate for delays inherent in transmission, a buffer temporarily holds data either for transmission onto the cable or for transfer into the PC. While in the buffer, data may be acted on, such as put into packets, or it may simply wait while the NIC handles other things.

An alternate to buffering is to use PC RAM. This can be less expensive, but is usually slower and it requires memory.

9 Data Packets

The NIC's most important job is packet formation.

Packets are the basic units of transmission. Files and messages are broken into packets for transmission. At the other end the packets are reassembled to form the file or message. Each computer has a unique address on the LAN to which frames can be sent.

The standard functions for data packets included:

- C Opening a communications session with another adapter.
- C Sending data to a PC.
- C Acknowledging the receipt of a data frame.
- C Broadcasting a message to all other adapters.
- C Closing a communications session.

A packet has three sections:

- C Header
- C Data
- C Trailer

Layout of a Generic Packet

SENDER ID	DEST ID	FRAME TYPE	DATA/MESSAGE	CRC
Header			Data	Trailer

Different network implementations define frames in different ways.

Data items common to all frames (Ethernet, Token Ring, and Arcnet) include:

- C The unique network address of the sender.
- C The unique network address of the receiver.
- C Identification of the contents of the frame.
- C Data record or message.
- C Checksum or CRC for error-detection.

10 Packet Sections

Header Section

The header includes an alert to signal that the packet is on its way, the packet's source address, destination address, and clock information to synchronize transmission.

In some networks, header also have preamble bits used for various purposes, like setting up parameters for transmission.

They can also have a control field to direct the packet through the network, a byte count, and a message type field.

Data Section

The data section contains the data being sent such as the numbers in a spreadsheet or words in a document. On some networks, the data section of a packet can be as large as 12KB.

On Ethernet the most widely used access protocol, it is 4KB. Most networks fall between 1KB and 4KB.

Trailer Section

The trailer contains error checking information called a cyclical redundancy check (CRC).

The CRC is a number that is the result of a mathematical calculation. The mathematical calculation is performed both before and after the packet is sent. If the result is the same--all the ones and zeros are in the right place--no errors occurred and the packet is retransmitted.

The trailer, like the header, can hold other information.

11 Data Transmission

Data comes from the PC in parallel form, 8, 16, or 32 bits at a time, depending on the bus width. However, it must travel over the cable in serial form, which is one bit at a time.

The steps involved in transmitting data over a LAN are as follows:

1. Formatting

The network interface card performs a parallel-to-serial conversion.

2. Encoding

This entails converting the data into a series of electrical pulses that convey information.

3. Getting Access to the Cable

The access method protocol, circuitry, and firmware for providing this service reside on the network access card.

C Token Ring and Arcnet use a token to grant network access.

C Ethernet lets any workstation transmit at will and then looks for collisions to see if it must transmit again.

4. Handshaking

In order to send data, a second NIC must be waiting to receive it. A short period of communication between two cards occurs before data is sent.

During this period, the NIC's negotiate the parameters for the upcoming transmission. The transmitting card sends the parameters it wants to use. The receiving card answers with its parameters such as:

C The maximum packet size

C How long to wait for an answer

C Buffer sizes

The card with the slower, smaller, less complicated parameters always wins because more sophisticated cards can "lower" themselves while less sophisticated cards can't "raise" themselves.

5. Placing the Data on the Cable

The NIC transceiver gives the data power to make it down the line. It puts the electrical signal out over the cable, making sure the data can get to the next NIC, repeater, amplifier, or bridge.

At the other end, a transceiver waits to accept the signal and begin the whole process in reverse.

6. Reversing the Process at the Receiving PC

C Modulating the signal through decoding.

C Serial-to-parallel conversion.

C Depacketizing the information into a format readable by the receiving device.

12 Network Interface Card - Selection

More than any other LAN component, the network interface card determines the performance of the LAN.

The speed of the disk drives, file servers, and network operating system are important, but the speed of the interface card and its software driver determines the network speed.

Choosing network cards is a difficult process. Nearly every vendor claims to have the fastest cards.

The following variables should be evaluated:

C Bus Width

A card using a 32-bit bus is faster than one that uses an 8-bit bus.

C Bus Type

Extended Industry Standard Architecture versus Micro Channel Architecture.

C Type of Memory Transfer

The type of memory transfer (shared memory is faster than I/O and DMA), and whether the card can perform bus mastering.

C Network Card Driver

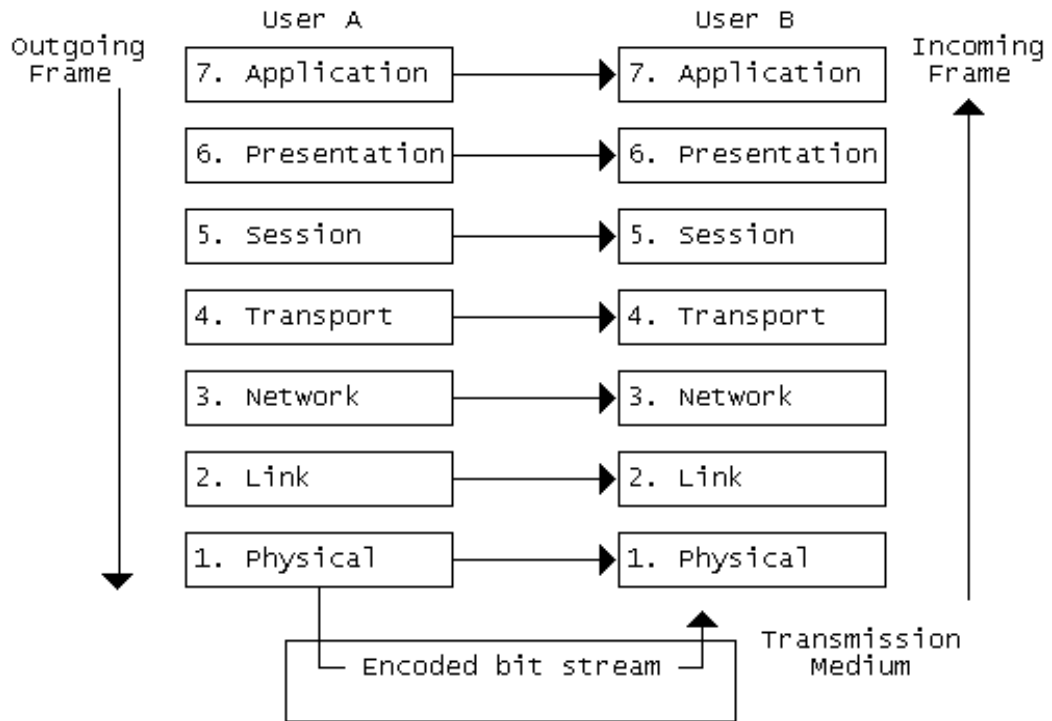
Test the speed of the NIC driver. Also ascertain whether it is NDIS compliant and certified.

13 Standardization and ISO

The ISO is working toward a definition of standard communication architecture called the Open Systems Interconnection (OSI) Reference Model.

Level 7: Application	C	Identifies the transmitting and receiving users.
	C	Sets an agreed-upon level of security and decides which security measures will be invoked.
	C	Makes an agreed-upon user responsible for error recovery.
Level 6: Presentation	C	Adapts for differences in language between users.
	C	Converts information into a format that both users can accept.
	C	Performs such tasks as character set conversion, syntax selection, and encryption.
Level 5: Session	C	Controls data exchange between users; maps names to network addresses.
	C	Assembles physical messages into logical messages.
Level 4: Transport	C	Manages the sequencing and priority of messages
Level 3: Network	C	Provides for end-to-end connection and quality of service.
Level 2: Data Link	C	Determines the best routing for the message and the protocol that devices must use to gain access to the cable and send information. Data is transmitted in packets. The rules governing the size and format of those packets are in the Data Link layer.
Level 1: Physical	C	Makes sure the message reaches the destination in the form originally sent.

14 The Protocol Process



15 OSI Reference Model

The OSI Reference Model divides the communications process into a hierarchy of seven interdependent functional layers. Each layer is responsible for carrying out different functions. The services each layer supplies are defined by a range of standards. A range of standards allows for different ways of performing the same service, which is vital if different users' needs are to be met.

The layers of the Reference Model define the functions of a complete communications network from the cable (Physical layer) to the network application software (Application layer). Each layer has a built-in interface to the adjacent layer. Layer 2 can pass data to Layer 3 or Layer 1, but cannot communicate directly to Layer 3.

The OSI Reference Model does not establish or promote any particular communication technique (protocol). The model's definitions are broad enough to include many protocols.

The reasons for following the OSI layered approach are migration and flexibility. You no longer need to change everything in a communications system just because one component of the system has been superseded by newer technology. For example, if a new type of cable (Physical Layer) is superior to the old cable, you can replace the old cable without making other modifications to the layers above the Physical layer.

Layers 1 and 2 are the hardware layers, providing the fundamental connection on which more sophisticated services are built. The three general protocols in common use are Ethernet (802.3), Token-Bus (802.4), and Token-Ring (802.5). They have defined by the 802 Committee of the Institute of Electrical and Electronic Engineers (IEEE).

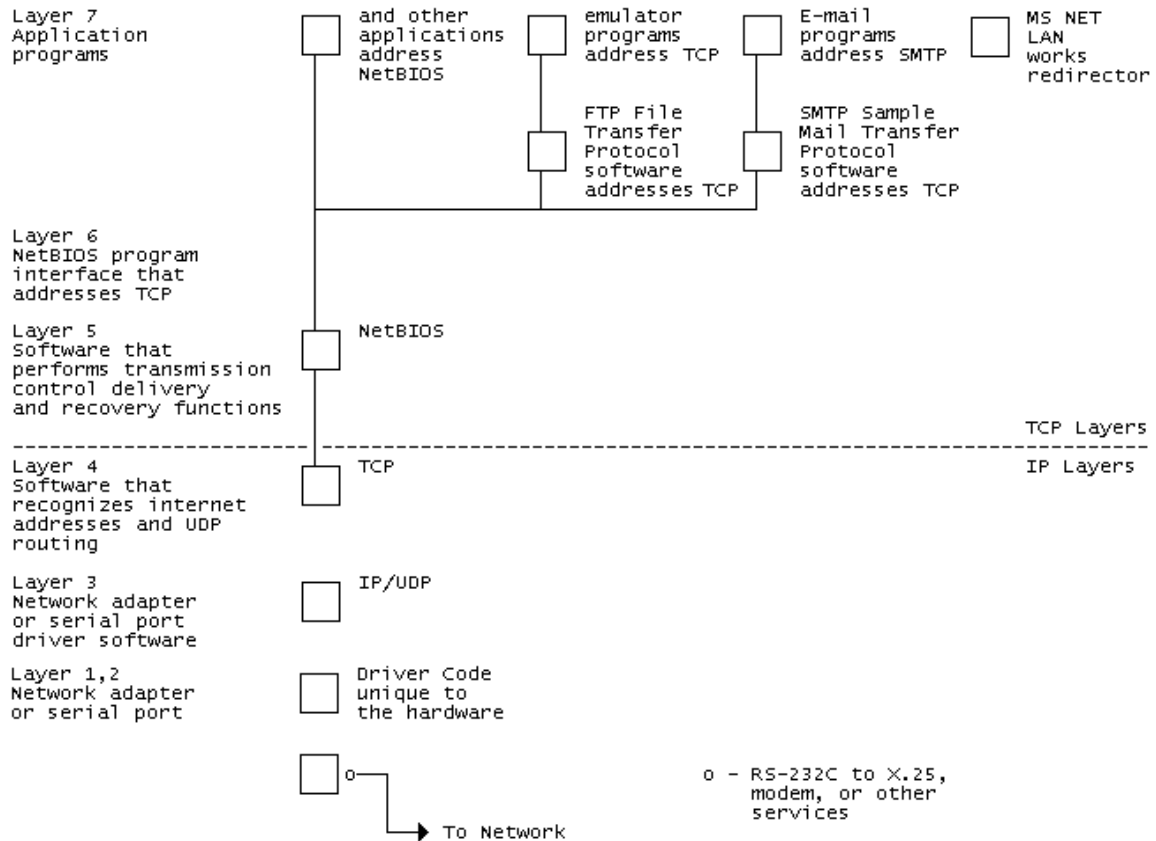
Layers 3 - 5 form the network's transport level - the software controlling network communications. This level is also known as the "subnet." Software at this level establishes and manages the temporary link between sender and receiver (virtual connection), and how communication application software (gateways, for example) attaches itself to the network.

Layer 6, the Presentation layer, is where almost all other application programs (LAN operating systems) should attach to the network.

Layer 7 defines network applications.

16 TCP/IP and the OSI Protocol

The upward path through software and hardware within the OSI seven-layer stack and the manner in which TCP/IP serves as a viable means of data transfer among dissimilar machines linked in a network.



17 TCP/IP Architecture**Layers**

Process Layer
FTP SMTP TELNET
Host-to-Host Layer
TCP
Internet Layer
IP
Network Access Layer

Comparison: TCP and OSI

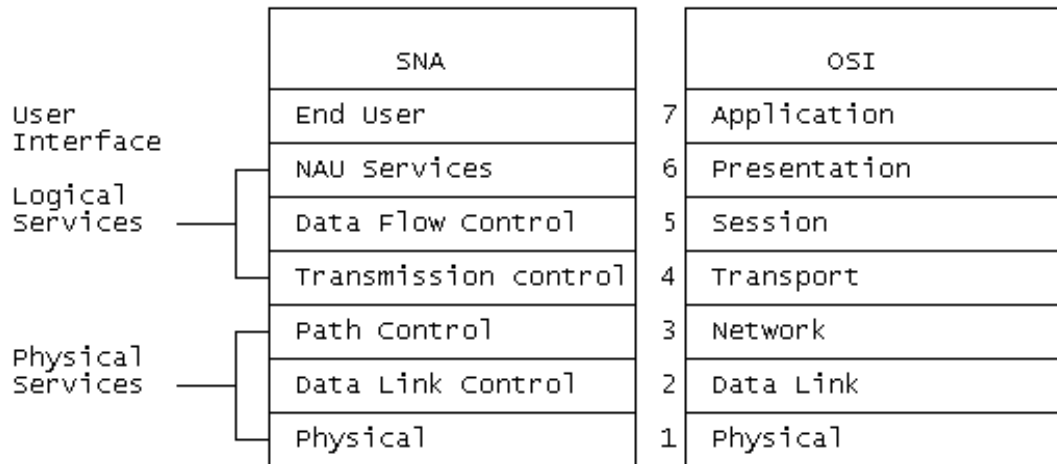
The TCP/IP network-access layer services correspond to those provided by the physical, data-link, and parts of the network layers in the OSI model.

OSI	TCP/IP
Application	Process
Presentation	
Session	
Transport	Host-to-Host Internet
Network	
Data Link	Network Access
Physical	

18 SNA and OSI Layers

Although SNA is a seven-layer protocol, in practice it provides functionality in only Layers 4 through 6.

These three offer functions roughly equivalent to those provided by Layers 4 through 6 in the OSI reference model.



19 Access Protocols: Purpose and Function

In many LANS, several attached devices must share a single transmission channel.

An access protocol establishes:

- C WHERE control of access resides in the network.
- C WHO gets the channel.
- C HOW MUCH channel capacity a device can have.

Access protocols are categorized as deterministic or contention based.

The two factors which affect the design and nature of an access protocol are:

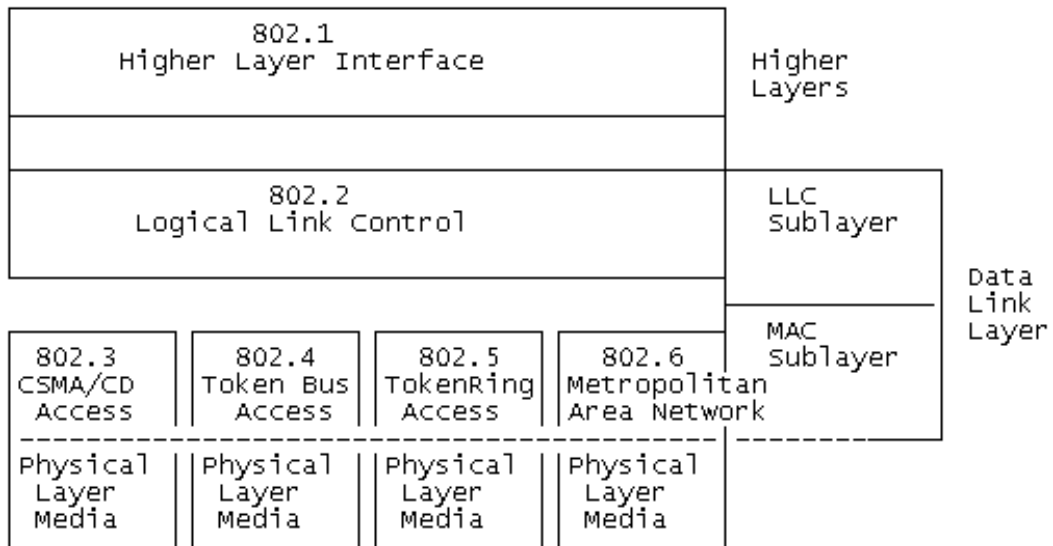
1. Control Strategy

The control strategy is either centralized or distributed.

2. Access Method

The access method specifies how a node can gain access to the communication path.

20 The IEEE 802 Architectural Layers



21 Control Access Protocols

In a controlled access network, a predetermined algorithm awards access on the basis of several factors, some of which might be message priority and type of device requesting access.

21.1 Polling

The polling protocol requires a central controller (server) and is frequently associated with a star topology. The central controller polls each node in turn, sensing whether or not the node is requesting access. If it is not, the controller moves on to poll the next node in order. If access is requested, the message is transmitted.

- C By its nature, polling eliminates the possibility of interference but may spend considerable overhead in the polling process, particularly when there are many inactive or seldom used devices online.
- C Access may be based on a list with an arbitrary order, or could be based on the physical location of the nodes.
- C Some nodes could be polled more frequently than others.
- C TDM is a form of polling.
- C Nodes may be forced to send all messages to central controller, or send them directly to the destination.

21.2 Token Passing

Token passing is similar in effect to polling, but operates in a distributed system without need of a central controller. Token passing is frequently associated with both ring and bus topologies.

This technique could be compared to passing a single microphone around a conference table. Those who have nothing to say pass the microphone (token) to the next seat. If that person has something to say, he completes his statement (perhaps under a time limit rule), then passes the microphone (token) to the next person.

The token consists of a number of bits used to control transmission. The token is transferred from one node to the next, providing each node in turn with the opportunity to transmit a message. When transmission is completed, the token is passed to the next node. Various controls can be built into the strategy to limit or specify the length of transmission.

Token passing has different implications when used on a ring than when used with another topology.

Token passing on a ring is implicit. Each node accepts the token from the node on one side, then passes it to the node on the other side. A node always passes the token to the next node physically in line.

Adding or deleting a node on a ring does not disturb the network operation. The next node may be a new or different node, but the logic implicit in the network maintains order. The ring requires a means of protection against a node that fails to pass the token. This is usually accomplished with a simple timeout circuit.

In another topology, such as a bus, the passing strategy is explicit addressing. The logic programmed into each node tells it to pass the token only to a specifically addressed node. In effect, a logical ring is created.

If a node were to malfunction or be removed from the network, or a new node were added, the network would require changing to maintain the logical ring. For this reason, adding or deleting nodes could be more difficult in a token passing bus.

21.3 Deterministic: Token Passing

- C A form of distributed polling.
- C The right to use the channel passes around the network in an orderly fashion.
- C Access is controlled by a special bit pattern known as a token.
- C Holder of token has exclusive right to transmit.
- C When transmission is complete, token is passed along.
- C Sender rather than receiver of message typically passes token to next node in the network.
- C For rings, the sequence in which the token is passed is reflected by the physical layout of the ring. Token is passed to next device on the ring.
- C In buses, the sequence may have no relationship to the physical layout of the network. Token is passed to a node based on next address in sending node table.

21.4 Slotted Access

This access method can be compared to a traffic signal. A green light allows traffic from one street to enter the intersection. An entire funeral line may not be able to get through the intersection with the same green signal, so they will have to wait for the next cycling of lights.

Data packet windows slots circulate from node to node to be filled with packets from transmitting stations. Each slot is preceded by a flag, which, like the token in token passing, is changed from free to busy status by the using station.

All the slots in one network are the same size. Messages too long for a slot must be broken into shorter segments, or packets for transmission.

Slotted access is most often used in ring topologies. It is most useful in a network where nearly all transmissions will be the same length.

Token passing and slotted access are popular in distributed control networks, like the ring, where each station is more or less equal partners in keeping the whole network running.

In centrally controlled networks, where network responsibility resides in a single host or controlling node, it makes more sense to let that node function the way a traffic cop would at the intersection we have been visualizing.

21.5 Slotted Ring

Slotted ring is another form of distributed polling.

Features and characteristics of slotted ring include:

- C A number of frames, usually of fixed size, circulate around the network.
- C A node that wishes to transmit can fill an empty frame.
- C Control information comprising sender address, receiver address and a bit to indicate that the frame is full.
- C Also referred to as a Pierce Ring.
- C Can have mix of dedicated and non-dedicated frames.

22 Demand Access Protocols

In a demand access protocol, a node gains access on demand. The mechanism for gaining access, once the node has requested it, depends on whether control is centralized or distributed.

22.1 Circuit Switching

This protocol is adaptable to any star topology.

When a node demands access, the central controller connects, by means of a switch, the calling node and the called node. If the answering node is communicating already, access is denied by the controller. If the connection is made, the caller's line and the answering line may be dedicated, that is, unavailable for other use until the two nodes complete their communication, no matter how long that might be, and regardless of the rate of communication.

The switching process itself can take considerable time, particularly since much control time is occupied with uncompleted calls. This results in high system overhead.

22.2 Carrier Sense Multiple Access/Collision Detection

CSMA/CD is the access method most often used in bus topologies.

It can be thought of as being similar to a party line telephone system without a central exchange. All subscribers are connected at all times and listening at all times for messages intended for them. Before talking, they listen to hear whether the line is in use. If two users begin talking at once, one stops, waiting for a quiet period to try again.

The network operates in much the same way. Before initiating a transmission, a node listens for a carrier signal (carrier sense) that would indicate the bus is in use. If so, it waits until the bus is free.

Since all nodes use the same carrier sensing technique, the only situation where two nodes might interfere is when two begin transmission at the same instant. This event is called a collision. Each node is equipped to detect this event (collision detection), and the protocol includes rules that determine how long the nodes will cease transmission and wait for another transmission opportunity.

The greater the number of the nodes in the network and the more active the nodes, the higher the rate of collisions. Since each collision takes some finite time to resolve, large, active systems can experience a reduction in throughput during periods of high activity.

22.3 Collision Detection

- C Node continues to listen while transmitting.
- C If the signal detected on the channel is different from the one being transmitted, a collision is detected.
- C When a collision is detected, the node backs off for a randomly chosen interval and then tries to retransmit.
- C Since access is statistical, there is no guarantee of access.
- C Large numbers of collisions will degrade channel capacity.
- C More efficient for systems transmitting large messages on short cables.
- C To insure that all collisions are detected, a minimum packet size is required.

23 Contention Protocols: CSMA/CD

23.1 General

- C Nodes compete for access to the medium.
- C Messages may collide and be destroyed, so a method for detecting or avoiding, collisions or lost data must be incorporated.

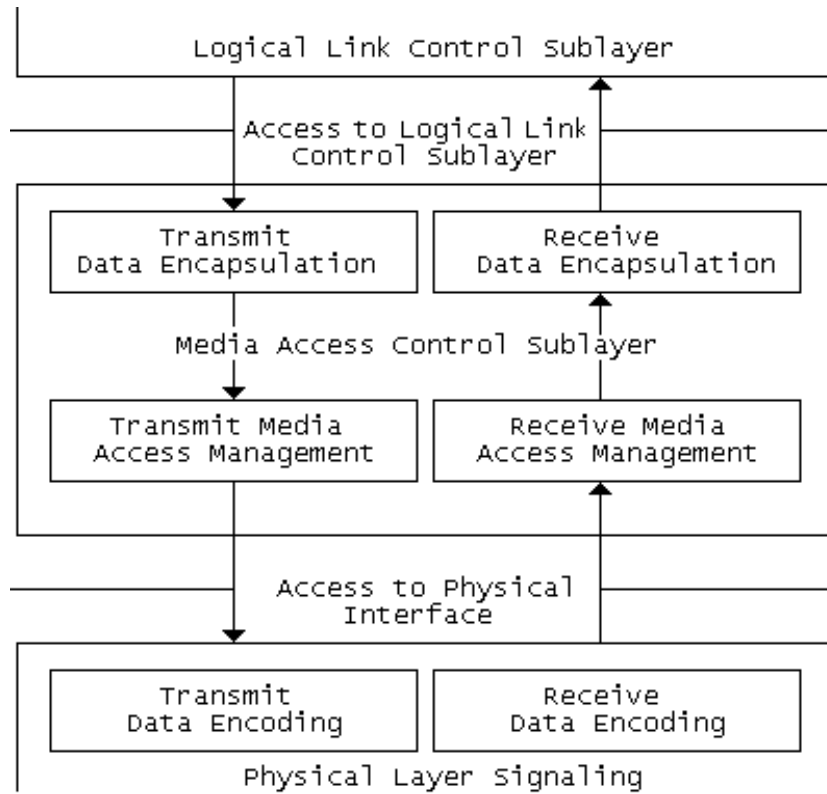
23.2 Aloha

- C Used to connect terminals throughout the islands of Hawaii to a central computer facility on Oahu.
- C Terminals can not 'hear' others sending, so they transmit whenever they have a message to send.
- C Required acknowledgment protects against loss due to collisions.
- C Maximum channel utilization is approximately 18%.

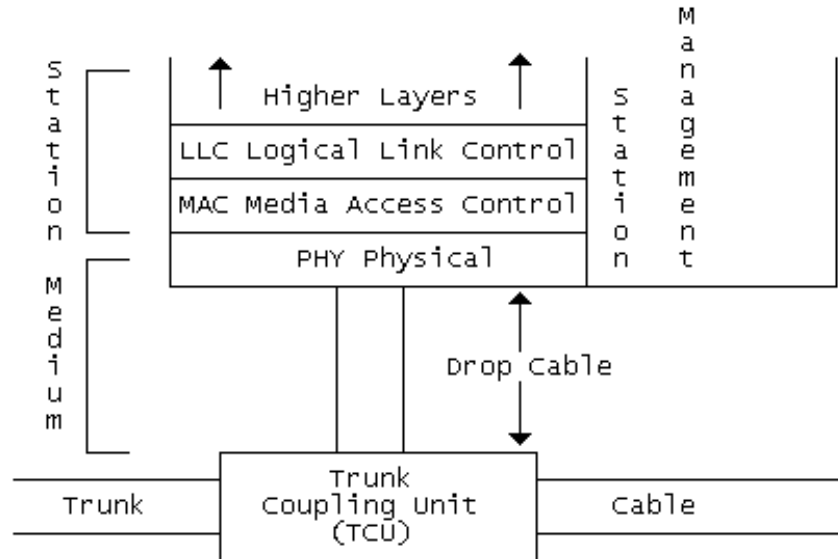
23.3 Slotted Aloha

- C Only allows messages to be transmitted at the beginning of a specified time slot.
- C Reduces likelihood of collisions.
- C Collision does not waste as much of the channel's capacity.
- C Maximum channel utilization is approximately 36%.

24 CSMA/CD Media Access Control Function



25 Physical Hardware Partitioning For Token-Passing Bus Network



26 Performance Comparisons

26.1 CSMA/CD

Efficiency depends on many factors including:

- C Length of bus
- C Speed of bus
- C Packet size
- C Number of active users

Because of the issue of collisions, the maximum delay to deliver a packet can not be calculated.

With light traffic load, CSMA/CD is superior to Token Passing because it does not have the time delay of the Token Passing overhead.

The larger the size of the packet, the more efficient CSMA/CD is.

26.2 Token Passing Ring

Efficiency depends on many factors including:

- C Number of users
- C Number of signed on users
- C Speed of ring
- C Length of ring
- C Early token release

With heavy traffic load and number of active users, Token Passing is superior to CSMA/CD because there are no collisions.

As a deterministic protocol, it is possible to calculate the delay, average, and maximum response times.

27 IEEE 802 Committee Standards

The IEEE assigns numbers to its active committees.

Committee 802 is a very large organization that includes members from industry and academia interested in a broad range of wide-area and local-area network systems.

Subcommittees of Committee 802 develop and maintain standards for several LAN topologies. The subcommittees use decimal numbers to identify their work.

27.1 802.3

These standards govern the use of the CSMA/CD (Carrier Sense Multiple Access/Collision Detection) network access method used by Ethernet networks. Although CSMA/CD and Ethernet are used interchangeably, there are technical differences.

Networks conforming to IEEE standard 802.3 use a carrier sense multiple access (CSMA) media-access control scheme on an bus or star topology.

This standard leaves room for several wiring options, including thin coaxial cable and unshielded twisted-pair wiring.

27.2 802.4

These standards govern the use of the token-passing, bus-based, media access method. In token-passing networks, an electrical signal, or token, is passed from one station to another. Only when a station has possession of the token can it transmit data.

27.3 802.5

These standards govern the use token-passing, ring-based, media-access method. Token Ring networks are wired electronically, or logically, in a closed circle, or ring. Physically, however, Token Ring networks are wired in a star configuration, with a centralized wiring hub providing the electrical circuits that form the logical ring.

28 Protocol Stacks

LAYER	ISO	DoD	IBM
7. Application	FTAM X.400 JTAM X.500 VT CASE	SMTP FTP NFS Telnet SNMP	
6. Presentation	8923		
5. Session	8327	UDP	NetBIOS APPC
4. Transport	8073 (TPO) 8602 (CONS)	XNS TCP	NetBIOS APPC
3. Network	8208 (X.25) 8473 (CLNS) 9542 (ES-IS) 8348 (CONS)	IP	APPC
2. Data-Link	8802.2 LLC 8802.3/4/5	LLC Ethernet	LLC HDLC SDLC MAC
1. Physical	8802.3 Ethernet 8802.4 Token Bus 8802.5 Token Ring	Ethernet FDDI	Token Ring Ethernet FDDI

29 LAN Standards: The State of the Market

Each LAN standard combines physical and logical topologies, signaling, and media-access control techniques in different ways.

There are three major protocol standards for LAN cabling and media-access control:

C Ethernet	C Token-Ring	C ARCnet
------------	--------------	----------

29.1 Ethernet

Ethernet was one of the first LAN architectures. This network cabling and signaling scheme entered the market in the late 1970s and is still a respected standard.

The Ethernet standard provides high-speed transmission at an economical price, offering a broad base of support for a variety of LAN and micro-to-mainframe applications.

Ethernet is a specification which describes a method for computers and data systems to connect and share cabling. Ethernet encompasses what the International Standards Organization calls the physical and data-link layers of data communications.

IEEE 802.3 Ethernet can have either of the two common physical topologies:

C Daisy Chain	C Star
---------------	--------

Ethernet is a broadcast system - each station broadcasts data into the network for all other station to hear. Regardless of how the wires run, the logical topology remains the same.

The primary characteristics of the physical Ethernet link are:

- C Data rate of 100 megabits per second.
- C Maximum station separation of 2.8 kilometers.
- C Shielded coaxial cable connecting the stations.
- C Electrical signaling on the cable is known as Manchester-encoded digital baseband which describes the electrical signals that make up the digital 0s and 1s that are constantly passing over the network.

30 Ethernet Intersystem Connectivity

Ethernet offers efficient ways for connecting to:

C	DEC	C	Hewlett Packard
C	IBM	C	Xerox
C	Many other computer systems.		

The latest growth area is in Ethernet adapters operating at data rates of 100 megabits per second.

30.1 Ethernet Cabling

The coaxial cabling scheme found most often in installation of PC-based networks uses a thin, 52-ohm coaxial cable between each pair of network stations. This cable, commonly called thin Ethernet and sometimes called "cheapernet," is typically limited to 305 meter (1,000 feet) between repeaters, although an IEEE specification limits it to 600 feet.

The oldest Ethernet cabling scheme uses heavily shielded coaxial cable (informally named "frozen yellow garden hose," to serve as a backbone among the clusters of nodes scattered around a building. The maximum length of cable between repeaters is 500 meters (1,640 feet), and the cable attaches to devices called transceivers, which transform the cable's connections into something more suitable for a PC or terminal.

A flexible transceiver cable made up of a shielded twisted-pair wire runs between the transceiver and the AUI port on the network adapter. Transceiver cables can be up to 15 meters (45 feet) long; they connect to the network card through a 15-pin D connector.

30.2 Ethernet Packaging and Moving Data

Ethernet uses a communications concept called datagrams to get messages across the network. The CSMA/CD media-access technique makes sure the two datagrams aren't sent out at the same time, and it serves as a method of arbitration if they are.

Ethernet datagram concept is supported by the simple premise that a communicating node will make its best effort to get a message across. The datagram concept does not include a guarantee that a message will arrive at any specific time or will be free of errors or duplications. Ethernet datagrams do not even guarantee that delivery will occur.

To provide a mechanism for assuring the receipt and accuracy of data, you will have to do the implementation via the higher-level software.

The Ethernet datagrams take the form of self-contained packets of information. These packets have fields that contain information about their destination and origin and the sort of data they contain. Large messages must traverse the network in multiple packets.

31 IEEE 10BaseT

The IEEE 802.3 family of standards describes carrier sense multiple access signaling, like Ethernet, used over various wiring systems.

10BaseT is the IEEE designation for Ethernet running on unshielded twisted-pair (UTP) wire in a physical star topology. The UTP can run directly to the adapters in each node or to an unshielded twisted-pair medium attachment unit (MAU) connected to the node through AUI cable.

The name 10BaseT indicates a signaling speed of 10 megabits per second, a baseband signaling scheme, and twisted-pair wiring in a physical star topology.

31.1 IEEE 10BaseT Advantages

- C Gives LAN managers the option of using installed telephone wiring thus saving installation costs and problems.
- C The technology of twisted-pair wiring, unlike coaxial Ethernet alternatives and Token-Ring's shielded twisted-pair, is familiar to technicians.
- C The star based wiring scheme of 10BaseT - running a single wire from the central wiring hub to the desktop - improves the reliability of the system over the daisy-chain wiring scheme.
- C The commonality of 10BaseT products.

10BaseT cards can be mixed and matched and wiring hubs from many companies can be used together on the same network. This translates into multiple sources of supply, competitive pricing, and confidence in long-term support.
- C Consistent performance.

According PC Magazine throughput test, 10BaseT twisted-pair, on a par with that of coaxial-cable Ethernet wiring.
- C The star wiring scheme provides both reliability and centralized management.

If one wire run is broken or shorted, that node is out of commission but the network remains operational.
- C The central wiring hub is an ideal place to install a monitoring microprocessor and network management software.

32 Token-Ring

Token-Ring is a set of standards developed by the IEEE 802.5 subcommittee which describes a token-passing network in a logical ring topology.

Token-Ring is a proprietary standard owned of IBM and actively promoted by IBM.

In conjunction with the IEEE 802.5 standards, IBM also put identical standards into place within the structure of the European Computer Manufacturers' Association.

The primary characteristics of the Token-Ring are:

- C The use of the token-passing media-access control system prevents messages from interfering with one another by guaranteeing that only one station at a time transmits.
- C Cabling containing two pairs of twisted wire covered by a foil shield.
- C 4-megabit per-second signaling with an option to implement a faster 16-megabit signaling is also part of the standard; the 100 megabit signaling is on the horizon in the near future.
- C Maximum length of cable between the Token-Ring hub and the attachment point for the network node is 150 feet (45 meters).
- C This streaming of data makes Token-Ring networks better suited to fiber-optic media than broadcast-type systems like Ethernet or ARCnet. Optical media typically carry one-way transmission, and the token travels in only one direction around the ring, therefore there is no need for optical mixers that divide power, or for expensive active repeaters.
- C The ability to implement TCP/IP over Token-Ring provides seamless connection to wide range of different networks.

32.1 Token Ring Speed

In 1989, IBM released a Token-Ring version using 16-megabit-per-second signaling. The 16-megabit adapters also work at 4 megabits on networks with the slower adapters.

Installing 16-megabit Token-Ring over unshielded twisted-pair wire does however introduce new problems. The allowable cable lengths and number of nodes in each ring are determined by a complex chart.

The faster signals are much more difficult to decode and more easily marked by cumulative noise on the cable system.

33 TCP/IP for Multi Platform Networking

As the number of PCs in an organization grows, so does the need to link those PCs to any of several minicomputer and mainframe systems. One of the most difficult problems system integrators and managers face is connecting different types of computers in a network.

The common denominator for linking a wide variety of mainframe, minicomputer, and PC systems is TCP/IP.

TCP/ICP is a \$250 to \$500 software package that provides easy file transfers and simple electronic-mail services between PCs and many kinds of dissimilar computer systems.

A packet driver, the software interface between application software and network adapter hardware, handles the job of routing packets - or messages on the network - to the various applications registered with it.

By allowing network protocol stacks to talk to the packet driver instead of the network interface card, packet drivers lessen the programmer's problem of supporting multiple protocols and applications that need to share the same LAN adapter card simultaneously.

Many corporations and almost all Federal-government organizations and universities in the U.S. have taken advantage of the availability and standardization of TCP/IP software.

Plans have been made for TCP/IP to evolve into something called, Transport Class 4 or TP4 under the ISO's program.

33.1 TCP/IP On-line

The basic challenge in resolving the interconnectivity issue is providing computers with different operating systems and architectures with the instructions for translating and interpreting data from foreign machines.

TCP/IP provides those instructions; the packets receive standardized handling when they arrive, regardless of the operating environment on the receiving side.

The TCP/IP module used by each machine must be customized for the computer and its operating system but standardized for the network. TCP/IP modules are available for hundreds of mainframe and minicomputer systems and for many PC networks.

If your network has a great deal of interaction between different types of machines, it makes sense to give every PC its own TCP/IP module. The penalties you pay for putting the software on every machine are greater RAM use and increased network overhead.

33.2 TCP/IP Gateway

Setting up a TCP/IP gateway is frequently the best solution for a homogeneous network of PCs that sometimes need access to a specific TCP/IP network on machine.

The connection between the gateway and the TCP/IP system can be through:

- C Ethernet coaxial cable
- C Public data networks
- C Private networks

When a TCP/IP gateway connects two Ethernet networks, it is called an Internet router.

33.3 TCP/IP: Internet Address

The heart of the IP portion of TCP/IP is the Internet address which is a 32-bit number assigned to every node on the network.

There are various types of addresses designed for different-sized networks, but you can write every address in base 10 using this form: 128.22.5.13. These numbers identify the major network and subnetworks a node is on.

The address identifies a particular node and provides a path that gateways can use to route information from one machine to another.

Although data-delivery systems like Ethernet or X.25 bring their packets to any machine electrically attached to the cable, the IP modules must know each others' Internet addresses to communicate.

A machine acting as a gateway between different TCP/IP networks will have a different Internet address on each network. Internal look-up tables and software based on another standard called the Address Resolution Protocol are used to route the data between networks through a gateway.

Another piece of software works with the IP-layer programs to move information to the right application on the receiving system. This software follows a standard called the User Data Protocol (UDP).

It is helpful to think of the UDP software as creating a data address in the TCP/IP message that details exactly what application the data block is supposed to contact when it reaches the address described by the IP software. The UDP software provides the final routing for the data within the receiving system.

33.4 TCP/IP: Transmission Control Protocol

The TCP or Transmission Control Protocol portion of TCP/IP comes into operation once the packet is delivered to the correct Internet address and application port.

Software packages that follow the TCP standard run on each machine, establish a connection to each other, and manage the communications exchanges.

Neither IP nor UDP knows anything about recovering packets that aren't successfully delivered, but TCP structures and buffers the data flow, loads for responses, and takes action to replace missing data blocks.

Conceptually, software that supports the TCP protocol stands alone. It can work with data received through a serial port, over a packet-switched network, or from a network system like Ethernet.

In concept, software doesn't need or even know about IP or UDP, but in practice TCP is an integral part of the TCP/IP equation and is most frequently used with IP and UDP.

34 Token Technique

In a token-passing ring network, a stream of data called a token circulates through the network stations when they are idle. This serves to define both the sequential logical topology and the media-access control protocol.

A station with a message to transmit waits until it receives a free token. It then changes the free token to a busy token, and transmits a block of data called a frame immediately following the busy token.

The frame contains all or part of the message the station has to send. The system does not operate by having one station accept a token, read it, and then pass it on. Instead, the stream of bits that make up a token or a message might pass through as many as three stations simultaneously.

When a station transmits a message, there is no free token on the network, so other stations wishing to transmit must wait. The receiving station copies the data in the frame, and the frame continues around the ring, making a complete round trip back to the transmitting station. The transmitting station purges the busy token and inserts a new free token on the ring.