Ethernet

Ethernet

April 11, 2006

FUJITSU

Trademarks and Copyrights

Acrobat and Reader are registered trademarks of Adobe Systems, Incorporated

AppleTalk is a registered trademark of Apple Computer

NetBIOS is a registered trademark of the IBM corporation

NetBUEI is a registered trademark of the IBM Corporation

UNIX is a registered trademark of UNIX System Laboratories, Inc., and is exclusively licensed by the X/Open Company, Ltd.

Windows is a registered trademark of the Microsoft Corporation

Windows NT is a registered trademark of the Microsoft Corporation

All other products or services mentioned in this document are identified by the trademarks, service marks, or product names as designated by the companies that market those products or services or own those marks. Inquiries concerning such products, services, or marks should be made directly to those companies.

This document and its contents are provided by Fujitsu Network Communications Inc. (Fujitsu) for guidance purposes only. This document is provided “as is” with no warranties or representations whatsoever, either express or implied, including without limitation the implied warranties of merchantability and fitness for purpose. Fujitsu does not warrant or represent that the contents of this document are error free.

Furthermore, the contents of this document are subject to update and change at any time without notice by Fujitsu, since Fujitsu reserves the right, without notice, to make changes in equipment design or components as progress in engineering methods may warrant. No part of the contents of this document may be copied, modified, or otherwise reproduced without the express written consent of Fujitsu.

Introduction

This self-study tutorial on Ethernet and Ethernet over SONET satisfies a prerequisite needed for attendance at Fujitsu Educational Services training. The tutorial gives a general overview of Ethernet:

- History
- Standards
- Frames
- Access
- Protocols
- Media
- Networks
- Topologies
- Equipment

Additionally, the tutorial provides specific SONET information as it relates to transporting Ethernet over SONET. The tutorial ends with a 25-question review of the information covered in the tutorial.

Any student who completes the tutorial can answer the review questions and, by missing no more than four questions, satisfy a prerequisite requirement for Fujitsu courses. If more than four questions are missed, the student should revisit the tutorial to ensure familiarity with all concepts and terms in the tutorial before attending class.

Distribution Method

The Ethernet tutorial can be viewed using Acrobat® Reader® and is available at the following Internet address:

http:www.fnc.fujitsu.com/services/pdfs/edservethernet.pdf

Additional tutorials are available at these sites:

- ATM:

http:www.fnc.fujitsu.com/services/pdfs/edservatm.pdf

- SONET:

http:www.fnc.fujitsu.com/services/pdfs/edservsonet.pdf

Table 1: Ethernet Standards

Supplement	Year	Description
802.3a	1985	10Base-2 (thin Ethernet)
802.3c	1986	10 Mb/s repeater specifications (clause 9)
802.3d	1987	FOIRL (fiber link)
802.3i	1990	10Base-T (twisted pair)
802.3j	1993	10Base-F (fiber optic)
802.3u	1995	100Base-T (Fast Ethernet and autonegotiation)
802.3x	1997	Full duplex
802.3z	1998	1000Base-X (Gigabit Ethernet)
802.3ab	1999	1000Base-T (Gigabit Ethernet over twisted pair)
802.3ac	1998	VLAN tag (frame size extension to 1522 bytes)
802.3ad	2000	Parallel links (link aggregation)
802.3ae	2002	10-Gigabit Ethernet
802.3ah	2004	Ethernet in the first mile
802.3as	2005	Frame expansion
802.3at	2005	Power over Ethernet Plus

Ethernet

Ethernet, a physical layer local area network (LAN) technology, is nearly 30 years old. In the last three decades, it has become the most widely used LAN technology because of its speed, low cost, and relative ease of installation. This is combined with wide computer-market acceptance and the ability to support the majority of network protocols.

Ethernet History

Robert Metcalfe, an engineer at Xerox, first described the Ethernet network system he invented in 1973. The simple, yet innovative and, for its time, advanced system was used to interconnect computer workstations, sending data between workstations and printers.

Metcalfe's Ethernet was modeled after the Aloha network developed in the 1960s at the University of Hawaii. However, his system detected collisions between simultaneously transmitted frames and included a listening process before frames were transmitted, thereby greatly reducing collisions.

Although Metcalfe and his coworkers received patents for Ethernet and an Ethernet repeater, and Ethernet was wholly-owned by Xerox, Ethernet was not designed nor destined to be a proprietary system. It would soon became a worldwide standard.

Ethernet Standards

The first Metcalfe system ran at 2.94 Mb/s, but by 1980 DEC, Intel, and Xerox (DIX) issued a DIX Ethernet standard for 10 Mb/s Ethernet systems. That same year, the Institute of Electrical and Electronics Engineers (IEEE) commissioned a committee to develop open network standards. In 1985, this committee published the portion of the standard pertaining to Ethernet (based on the DIX standard)—*IEEE 802.3 Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications*. Even though the IEEE title does not mention Ethernet, Metcalfe's original term for his network system had caught on, and IEEE 802.3 was and is referred to as the Ethernet standard.

Note: *The IEEE standard was called 802 because work on it started in February 1980.*

As described in Table 1, many more Ethernet standards have been created since 1985. The IEEE standards have been adopted by the American National Standards Institute (ANSI) and by the International Organization of Standards (ISO). ISO standardization means that companies and organizations around the world use these standards when manufacturing Ethernet products and installing Ethernet network systems.

Fast Ethernet

While 10 Mb/s seemed very fast in the mid-1980s, the need for speed resulted in a 1995 standard (IEEE 802.3u) for 100 Mb/s Ethernet over wire or fiber-optic cable. Although the 100Base-T standard was close to 10Base-T, network designers had to determine which customers needed the extra bandwidth. Because there was a choice of bandwidths, the standard also allowed for equipment that could autonegotiate the two speeds.

In other words, if an Ethernet device was transmitting or receiving from a 10 Mb/s network, it could support that network. If the network operated at 100 Mb/s, the same device could switch automatically to the higher rate. Ethernet networks then could be 10 Mb/s or 100 Mb/s (Fast Ethernet) and connected with 10/100 Mb/s Ethernet devices that automatically switched network speeds.

Gigabit Ethernet

Gigabit Ethernet works much the same way as 10 Mb/s and 100 Mb/s Ethernet, only faster. It uses the same IEEE 802.3 frame format, full duplex, and flow control methods. Additionally, it takes advantage of CSMA/CD when in half-duplex mode, and it supports simple network management protocol (SNMP) tools.

Gigabit Ethernet takes advantage of jumbo frames to reduce the frame rate to the end host. Standard Ethernet frame sizes are between 64 and 1518 bytes. Jumbo frames are between 64 and 9215 bytes. Because larger frames translate to lower frame rates, using jumbo frames on Gigabit Ethernet links greatly reduces the number of packets (from more than 80,000 to less than 15,000 per second) that are received and processed by the end host.

Gigabit Ethernet can be transmitted over CAT 5 cable and optical fiber such as the following:

- 1000Base-CX—Short distance transport (copper)
- 1000Base-SX—850 nm wavelength (fiber optics)
- 1000Base-LX—1300 nm wavelength (fiber optics)

10 Gigabit Ethernet

The operation of 10 Gigabit Ethernet is similar to that of lower speed Ethernets. It maintains the IEEE 802.3 Ethernet frame size and format that preserves layer 3 and greater protocols. However, 10 Gigabit Ethernet only operates over point-to-point links in full-duplex mode. Additionally, it uses only multimode and single mode optical fiber for transporting Ethernet frames.

Note: *Operation in full-duplex mode eliminates the need for CSMA/CD.*

The 10 Gigabit Ethernet standard (IEEE 802.3ae) defines two broad physical layer network applications:

- Local area network (LAN) PHY
- Wide area network (WAN) PHY

LAN PHY

The LAN PHY operates at close to the 10 Gigabit Ethernet rate to maximize throughput over short distances. Two versions of LAN PHY are standardized:

- Serial (10GBASE-R)
- 4-Channel course wave division multiplexing (CWDM) (10GBASE-X)

The 10GBASE-R uses a 64B/66B encoding system that raises the 10 Gigabit Ethernet line rate from a nonencoded 9.58 Gb/s to 10.313 Gb/s. The 10GBASE-X still uses 8B/10B encoding because all of the 2.5 Gb/s CWDM channels it employs are parallel and run at 3.125 Gb/s after encoding.

The MAC to PHY data rate for both LAN PHY versions is 10 Gb/s. Encoding is used so that long runs of ones and zeros that could cause clock and data problems are greatly reduced.

WAN PHY

The WAN PHY supports connections to circuit-switched SONET networks. Besides the sublayers added to the LAN PHY (discussed in the following two pages), the WAN PHY adds another element called the WAN interface sublayer (WIS). The WIS takes data payload and puts it into a 9.58464 Gb/s frame that can be transported at a rate of 9.95328 Gb/s. The WIS does not support every SONET feature, but it carries out enough overhead functions (including timing and framing) to make the Ethernet frames recognizable and manageable by the SONET equipment they pass through.

10GBase Interfaces

Just as Fast Ethernet and Gigabit Ethernet have multiple interfaces, 10 Gigabit Ethernet has seven interfaces referred to in Table 2.

Table 2: 10GBASE-x Interfaces

Interface	PHY	Optics
10GBASE-SR	LAN	850 nm serial
10GBASE-LR	LAN	1310 nm serial
10GBASE-ER	LAN	1550 nm serial
10GBASE-LX4	LAN	4 x 1310 nm CWDM
10GBASE-SW	WAN	850 nm serial
10GBASE-LW	WAN	1310 nm serial
10GBASE-EW	WAN	1550 nm serial

Figure 1: LAN PHY/WAN PHY Sublayers

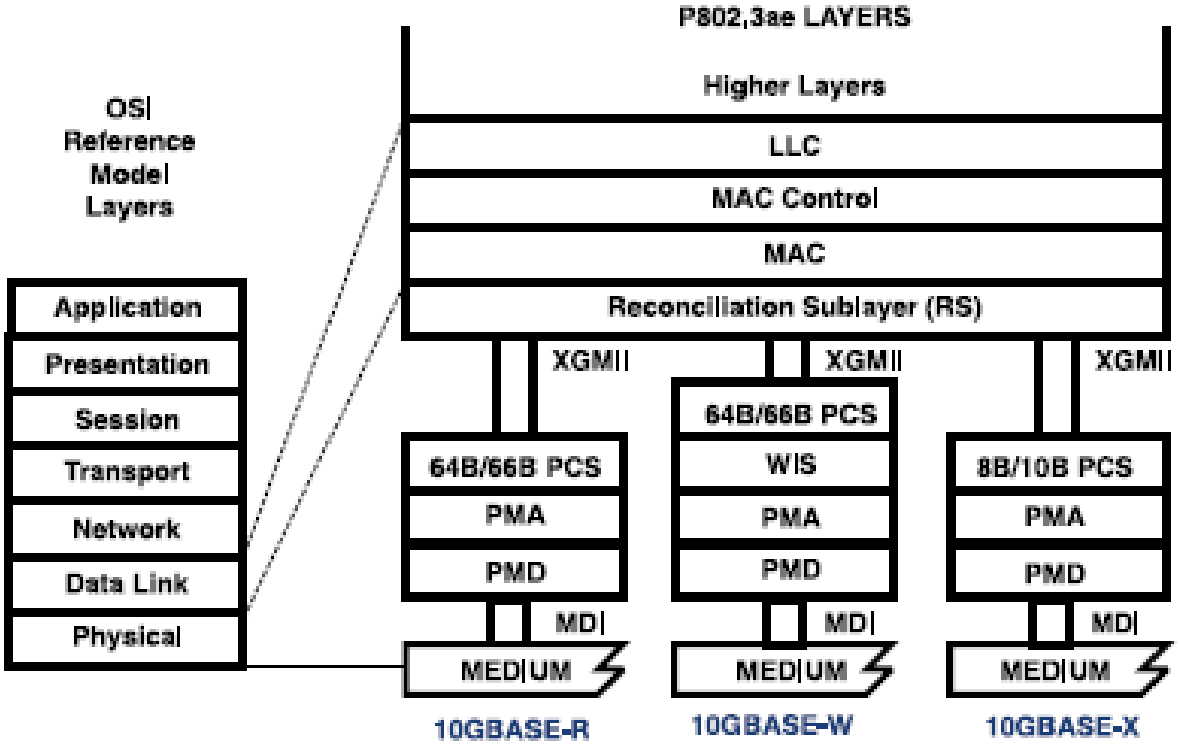

LAN PHY/WAN PHY Sublayers

The PHY is a circuit block at the physical layer that includes the following sublayers (see Figure 1):

- Physical coding sublayer (PCS)
- Physical medium attachment (PMA)
- Physical medium dependent (PMD)

Physical Coding Sublayer

The PCS encodes and decodes the data stream between the MAC and PHY layer. There are three categories for the PCS:

- 10GBASE-R—Serially encoded (64B/66B); 10.3 Gb/s rate not SONET compatible (LAN PHY)
- 10GBASE-X—Serially encoded (8B/10B); used for wavelength division multiplexing (WDM) transmissions (LAN PHY)
- 10GBASE-W—Serially encoded (64B/66B); compatible with SONET standards for a 10 Gb/s WAN (WAN PHY)

Physical Medium Attachment

The PMD is an optional interface for connection to optical modules. The two PMD interfaces are:

- 10 Gigabit Ethernet attachment unit interface (XAUI)
- 10 Gigabit Ethernet 16-bit interface (XSBI)

The XAUI is an interface to specialized 10 Gigabit Ethernet optical modules and system backplanes. It supports 4 SERDES transmit and 4 SERDES receive channels for 8B/10B encoding.

Note: *SERDES stands for SERializer/DESerializer. SERDES is used in high speed communications to convert data from/to serial and parallel data streams.*

The XSBI is a serial optics interface for LAN and WAN PHY applications. Intermediate and long reach optical modules use this interface. It requires more power and more pins than an XAUI.

Physical Medium Dependent

Distance objectives are met by using the physical medium dependent sublayer. Four different PMDs are defined to support single and multimode optical fibers:

- 850 nm serial—MMF, 500 MHz/km, up to 65 meters
- 1310 nm serial—SMF, up to 10 km
- 1550 nm serial—SMF, up to 40 km
- 1310 nm CWDM—MMF, 500 MHz/km, up to 300 meters
- 1310 nm CWDM—SMF, 10 km

Figure 2: Ethernet Frames

DIX Ethernet Frame

Preamble	Destination Address	Source Address	Type	Data	Cyclical Redundancy Check
8 bytes	6 bytes	6 bytes	2 bytes	Up to 1500 bytes	4 bytes

m1619eq_1

IEEE 802.3 Ethernet Frame

Preamble	Start Frame Delimiter	Destination Address	Source Address	Length	LLC	Data	Pad	Frame Check Sequence
7 bytes	1 byte	6 bytes	6 bytes	2 bytes	Up to 1500 bytes			4 bytes

m1619a_2

Ethernet Frames

Ethernet is commonly described as being a packet delivery system. In reality, an Ethernet frame is made up of all the necessary parts to fit the requirements and definition of a packet. An Ethernet frame has a header (Preamble - Length), payload (LLC - Pad), and a trailer (Frame Check Sequence) that are bundled together in a specifically organized manner for transmission (see Figure 2).

Header

- Preamble—Sets bit timing and signals that a frame is being sent (10 Mb/s Ethernet)
- Start Frame Delimiter—8-bit sequence (10101011)

Note: *100 and 1000 Mb/s Ethernet systems signal constantly and do not need preamble or start frame delimiter fields.*

- Destination Address—48-bit receiving hardware media access control (MAC) address
- Source Address—48-bit transmitting hardware address
- Type—Indicates protocol sending the frame (DIX only)
- Length—Indicates the length of data field (number of LLC data bytes) (IEEE 802.3 only)

Payload

- Logical Link Control (LLC)—Governs the assembly of data at the data link (Layer 2) level
- Data—Payload contained in a field (between 46 bytes and just over 1500 bytes in length)
- Pad—0 bits added to the data field if there are fewer than 46 bytes of data in that field

Trailer

- Cyclical Redundancy Check (CRC)—Detects DIX-only transmission errors
- Frame Check Sequence (FCS)—Detects transmission errors and provides quality of service at receiving end

Note: *According to section 3.3 of the IEEE 802.3 standard, each octet of the Ethernet frame, with the exception of the FCS, is transmitted low-order bit first.*

Ethernet Address

Each Ethernet network interface card (NIC) has a unique identifier called a MAC address that is assigned by the card manufacturer. Each manufacturer that complies with IEEE standards can apply to the IEEE Registration Authority for a range of numbers for its cards.

Each MAC address is a 48-bit number, of which the first 24 bits identify the manufacturer. This part of the MAC address (manufacturer ID or organizational unique identifier [OUI]) is assigned by the registration authority. The second half of the address (extension of board ID) is assigned by the manufacturer. The number is usually programmed into the hardware so that it cannot be changed.

Because the MAC address is assigned to the NIC, it moves with the computer. Even if the interface card moves to another location across the world, the user is reached at this address.

Ethernet Access

The CSMA/CD standard defines how Ethernet frames get onto an Ethernet network. Because only one signal at a time can be transmitted on an Ethernet network, every Ethernet device listens to hear if another device is already transmitting (in other words, sense other carriers to determine if one of them is transmitting a signal). If the path is clear, any Ethernet device can transmit because of multiple access to the network. But all devices, even the one transmitting, continue to listen, because they are trying to detect collisions. Frames that collide must be retransmitted.

Ethernet is a shared media, so collisions not only happen, they are expected. Sometimes a transmission is not detected by an Ethernet device, and the device transmits on a busy line. Sometimes two or more devices (after correctly determining the path is clear) transmit simultaneously, resulting in a collision somewhere on the network. However, having too many collisions is not good, so rules were established to minimize these conflicts and protect data integrity.

Besides the constant listening, there is an enforced minimum quiet time of 9.6 microseconds between frame transmissions. There has to be a break in traffic to allow other devices a chance to get their data moving.

If a collision occurs, retransmission is determined by an algorithm that chooses unique time intervals for resending the frames. The Ethernet interface backs off, or waits, the chosen number of milliseconds and then retransmits automatically if no activity is detected. The process is repeated for that frame if another collision occurs. In fact, as same-frame collisions recur, the process is repeated until the frame collides up to 16 times. Then, after this many tries, it is discarded.

Another way data can be lost is if there is a late collision. Late collisions happen when two devices transmit and the resultant collision is not detected because of a bad cable or too many repeaters on the network. If the time to send a frame from one end of the network to another is longer than the time it takes to put a whole frame on the network, neither device will detect that the other device is transmitting. A late collision can be detected only for frames of 64 or more bytes. As many frames are less than 64 bytes, their collisions go undetected and are not retransmitted.

Segmenting the network and establishing smaller collision domains is a good way to detect and reduce collisions. Another way is to reduce the number of computers and users on the network. And a well-designed and maintained cabling infrastructure can reduce normal collisions and prevent late collisions.

The CSMA/CD protocol is designed to allow fair access by all transmission devices to shared network channels. This fair sharing protocol means that even when all elements are working properly, collisions can occur. For normal Ethernet traffic levels, it is thought that if deferred and retransmitted traffic is less than 5 percent of total network traffic, the Ethernet network is healthy.

Full Duplex

While listening before talking applies to half-duplex systems, the 1997 IEEE 802.3x standard describes full-duplex Ethernet operation between a pair of stations. Simultaneous transmit and receive is over twisted-pair or fiber-optic cables that support two unidirectional paths. Besides the cabling, the Ethernet devices must support simultaneous transmit and receive functions.

The 1997 standard calls for traffic flow control, called MAC Control Protocol and PAUSE. If traffic gets too heavy, the control protocol can pause the flow of traffic for a brief time period.

Ethernet Equipment

The equipment needed to facilitate communication on an Ethernet system is determined, to a large extent, by the size and number of networks transporting Ethernet frames.

Bridges

Bridges transfer MAC-layer packets from one network to another. They serve as gatekeepers between networks and allow only necessary traffic (frames) between the networks they connect. Bridges control traffic by checking source and destination addresses and forwarding only network-specific traffic. Bridges also check for errors and drop traffic that is corrupted, misaligned, and redundant.

Bridges help prevent collisions and create separate collision domains by holding and examining entire Ethernet packets before forwarding them on. This allows the network to cover greater distances and add more repeaters onto the network.

Most bridges can learn and store Ethernet addresses by building tables of addresses with information from traffic passing through them. This is a great advantage for users who move from place to place, but it can cause some problems when multiple bridges start network loops. Spanning Tree Algorithm software is used to prevent these loops.

Note: *More information on Spanning Tree Algorithm is in IEEE 802.1d.*

Routers

Routers work much like bridges and switches by filtering out unnecessary network traffic, but they concentrate on network (OSI Layer 3) functions rather than physical (OSI Layer 1) functions.

They filter network traffic by dividing networks logically into subnetworks and only admit traffic destined to IP addresses on those individual subnets. By acting as a firewall, a router can prevent unwanted packets from either entering or existing a network.

Routers also act as an additional level of security. Routers can be configured with access lists that define which protocols and hosts have access to a network.

Switches

Ethernet switches took the bridge concept and made it bigger by linking multiple networks. Switches can increase network performance by eliminating extraneous traffic on network segments.

Two varieties of switches are commonly in use:

- Cut through
- Store and forward

Cut through switches use algorithms that read the destination address on the frame header and immediately forward the frame to the switch port attached to the destination MAC address. This is very fast because once the header information is read, the rest of the frame is transmitted without inspection.

Store-and-forward switches, like a bridge, hold the frame until the whole packet is inspected. This type of switch makes sure the frame is fit to travel before transmitting it. Fortunately, store and forward switches are nearly as fast as cut-through switches, so the extra work does not require significant additional time.

Hubs/Repeaters

Active hubs and repeaters connect LAN segments, and the terms often are used synonymously. Active hubs are described as the central device in a star topology, and while it connects multiple cable segments at one point, it also repeats the incoming signals before transmitting them on to their destinations.

Hubs can be designed to support Simple Network Management Protocol (SNMP) so that network management software can administer and reconfigure the hub.

Repeaters are used when the distance between computers is great enough to cause signal degradation. The repeaters, too, must listen and pass their signals only when the line is clear.

Table 3: ISO 7-Layer Reference Model

Layer Number	Layer Name	Layer Purpose
7	Application	Supports file transfers, virtual terminals, remote file access
6	Presentation	Provides communication services that mask differences in dissimilar system' data formats
5	Session	Manages dialog
4	Transport	Defines information exchange rules and provides delivery management, including error recovery and flow control
3	Network	Controls data transfer between computers
2	Data Link	Checks protocols and procedures for operating communication lines
1	Physical	Ensures electrical, mechanical, and functional control of data circuits

Table 4: TCP/IP 5-Layer Reference Model

Layer Number	Layer Name	Layer Specification
5	Application	Particular network applications
4	Transport	Reliable data transport
3	Internet	Frame formatting and routing
2	Network	Frame organization and transmittal
1	Physical	Network hardware

Ethernet Protocols

Network protocols are established standards that describe how computers communicate with each other. These protocols include:

- Computer identification
- In-transit data format
- Final destination data processing
- Lost or damaged transmission handling

Different computer manufacturers have different operating systems that support these protocols. Some examples of common network protocols are:

- TCP/IP (UNIX®, Windows® NT, Windows 2000)
- DECnet (formerly Digital Equipment Corporation)
- AppleTalk® (Macintosh® computers)
- NetBIOS/NetBEUI® (LAN Manager, Windows NT networks)

While these network protocols are diverse and varied, they all share something in common—the same physical cabling. Because the network devices are compatible at the physical and data link levels, they can run different operating systems over the same medium and communicate successfully with each other. This concept is commonly referred to as protocol independence.

Transmission Control Protocol/Internet Protocol (TCP/IP) are only two of several Internet protocols; however, they have come to represent all of them. They are commonly called TCP over IP because IP operates at the Layer 3 (network) level and TCP is over that layer, operating from the Layer 4 (transport) level.

Message delivery is guaranteed by TCP. It allows cooperating computers to share resources across a network. TCP works with IP because IP is responsible for addressing and routing the messages.

TCP/IP is popular for several reasons:

- It is an excellent client-server application platform, especially when used in a wide-area network (WAN) environment.
- It provides avenues for sharing a wealth of information.
- It is generally available throughout the world. Even manufacturers of Ethernet peripherals make their products TCP/IP compliant.

The Open Systems Interconnect (OSI) 7-layer reference model (refer to Table 3) was revised to accommodate TCP/IP (refer to Table 4) and to meet the needs of protocol designers. By tying into all layers of the OSI reference model through TCP/IP, Ethernet is able to greatly increase its information-sharing capabilities.

Table 5: IP Addresses

Class	Identifiers	Network Address	Host Address
A	0/1 through 126	7 bits (first byte)	24 bits (last three bytes)
B	10/128 through 191	14 bits (first two bytes)	16 bits (last two bytes)
C	110/192 through 223	21 bits (first three bytes)	8 bits (last byte)
D	1110/224 through 239	Ranges from 224.0.0.0 through 239.255.255.255	
E	11110/240 through 255	Reserved	

IP Addresses

An IP address has two parts (refer to Table 5):

- Network
- Host (also known as local or node)

Each network has an Internet address. Each network also must know the address of every other network with which it communicates. A unique network number is requested from the Network Information Center (NIC) and becomes part of that network's IP address.

After the network is identified, the specific host or node must be specified. A unique host address for the particular network is added to the end of the IP address.

Subnet Mask

Subnetting enables the network administrator to further divide the host part of the address into two or more subnets. The subnet mask is the network address plus the bits reserved for identifying the subnetwork.The rightmost bits in a valid subnet address must be set to 0. (An IP address must accompany a subnet address.)

Network Classes

Networks vary in size. The network number issued by the NIC identifies the size of the network:

- Class A—Large networks
- Class B—Medium-sized networks
- Class C—Small networks with less than 256 devices
- Class D—Multicasting

Dot Address

A dot address is the IP address expressed by four bytes that are separated from each other by three dots, such as the following:

205.245.172.72

The dots make the address easier to read. They also help identify the IP address network classes and separate the network address from the host address.

If the first byte of a dot address has a range of numbers between 1 and 126, it is a Class A network address. The remaining three bytes will contain the host address. There are fewer large networks (short addresses), but there are many more nodes (long addresses) in these large networks.

According to Table 5, the dot address example would apply to a small Class C network. There are many small networks so the network address is much longer than the Class A address and takes up the first three bytes. The node address (one 256 nodes or less) is contained in the last byte of the address.

Classless Interdomain Routing

Because of the proliferation of nodes within networks, classless interdomain routing (CIDR) adds an IP network prefix to the end of the IP address such as the following:

205.245.172.72/16

CIDR addresses reduce the size of routing tables and make more IP addresses available within each network.

Spanning Tree Protocol

The Spanning Tree Protocol (STP) prevents loops on an Ethernet network. It is a protocol developed by IEEE to allow bridges to have multiple connections between networks. STP reduces problems that occur on redundant links by:

- Assigning priorities to each bridge (this produces a logical tree topology of the bridges)
- Defining (and redefining, when necessary) a single path through a network

An algorithm used by the STP establishes a root bridge by using bridge priority and MAC address information that is transmitted by each bridge in a Bridge Protocol Data Unit (BPDU). The bridge with the lowest MAC address and highest priority becomes the root bridge. The root bridge now designates one of its ports as the root port.

Note: *Once selected, only the root bridge continues to generate BPDUs (also called hello packets). The other bridges then update and forward the packets they receive.*

The other bridges determine the least cost path to the root bridge to determine a bridge for each LAN. The cost of the path is determined by such factors as wire speed, distance from the root bridge, and available bandwidth.

Note: *The designated bridge assigns one of its ports as the LAN root port, and that port is the only one that forwards BPDUs on that LAN.*

The spanning tree algorithm places all bridge ports in one of four states as the path discovery is in process:

- Blocking—Does not receive nor pass frames
- Listening—Listens to BPDUs to ensure no loops occur before passing frames
- Learning—Learns MAC addresses and other forwarding information but does not pass frames
- Forwarding—Receives and transmits frames

Once the spanning tree is complete, the ports on the network are placed in either blocking or forwarding states. Loops are prevented because a single path over selected bridges and through designated ports is established. If a failure occurs, the process is repeated and a new pathway is generated.

Rapid Spanning Tree Protocol

Rapid Spanning Tree Protocol (RSPT) works in much the same way as STP. It eliminates loops by reducing the bridged network to a single spanning tree topology and reconfigures redundant connections if a link or component failure occurs. The major difference is the significantly shorter reconvergence time with RSPT. An STP can take between 30 to 50 seconds to reconverge while RSTP takes less than a second.

The RSTP takes different processing steps than STP that allow this time difference, including the following:

- Faster filtering database aging
- BPDU generation
- Edge port recognition

The STP bridges do not discard their MAC address databases when a topology change is detected. Instead, the bridge sends a change notice to the root bridge and the root bridge informs other bridges of the change. This can take several seconds. Any RSTP bridge that detects a topology change sends a BPDU to inform all other switches, not just the root bridge, that a topology change occurred. The detecting bridge immediately deletes the old database, as do the receiving switches without waiting for a BPDU from the root bridge.

If an STP switch does not hear from the root bridge, it waits for a maximum age time (20 second default) to declare the root bridge dead. The RSTP switch waits only 3 hello times (about 2 seconds each hello time) before declaring a root bridge dead. This RSTP switch takes over as the root bridge and sends out a BPDU that starts the reconvergence process.

Edge ports that connect directly with end user equipment cannot create loops. RSTP recognizes this distinction and places all edge ports in a forwarding state.

When a topology change occurs, STP requires time to wait for the maximum age time to expire and for time for the spanning tree ports to listen and learn before the new spanning tree is created. The RSTP topology change does not require a long maximum age time or a reconvergence of the entire network. The RSTP monitors the MAC states and deletes nonfunctional ports. Then it processes BPDUs from all the switches, not just the root bridge, to detect topology changes. When a change is detected, RSTP quickly places an alternate port in the forwarding state.

Table 6: Ethernet Media

Transmission	Media	Cabling	Connector
10Base-2	Thin Ethernet	Thinwire (RG-58) coax	BNC
10Base-5	Thick Ethernet	Thickwire (10mm) coax	RG-8
10Base-T[1]	Unshielded twisted pair (UTP)	Category 3 (Cat. 3)	RJ-45
100Base-T	UTP and fiber optics	Cat. 3 to Cat. 5	RJ-45 and FC, ST, SC, LC
1000Base-T	UTP	Cat. 5 and greater	RJ-45
10Base-FX[2]	Fiber optic	Single-mode fiber (SMF) or multimode fiber (MMF)	FC, ST, SC, LC
1000Base-SX 1000Base-LX	Short wavelength Long wavelength		

Ethernet Media

A big part of designing and installing an Ethernet system is using the right medium. A large selection of cables is a result of Ethernet evolution over the years and Ethernet flexibility (refer to Table 6).

The three major types of media in use today are:

- Coax
  - Thick wire 10Base-5
  - Thin wire 10Base-2
- Twisted pair
  - 10Base-T
  - 100Base-T (twisted pair and fiber-optic cable)
  - 1000Base-T
- Fiber optic
  - 10Base-F
  - 1000Base-X

The first Ethernet systems used 10Base-5 coax cable. The name identifies different functions and limitations of the cabling:

- 10—Mb/s bandwidth of the signal
- Base—Cable carries only Ethernet
- 5—500 meters is the longest distance a cable segment can carry the signal

The expensive thick wire cable (10 mm) was soon replaced by less expensive, but less efficient, 10Base-2 thin wire cable. The maximum range for thin wire is 185 meters (that rounded up to 2 for the name of the cable).

The most popular Ethernet wiring is UTP cable that comes in a wide variety of grades (categories), performance levels, and prices. The UTP cable is similar to telephone plug cable, except it has eight connectors that are either straightthrough or crossover. Straightthrough connections are used in most Ethernet networks, while crossover cables can connect two computers directly to each other without the use of a hub.

Cat. 3 to Cat. 5 UTP is used for 10Base-T and 100Base-T electrical signals. Cat. 5 or better is used for 1000Base-T.

Fiber optic cables are gaining more in popularity each year. While it is more expensive than electrical cable, it is unaffected by many environmental conditions that can be problems for coax or UTP. Additionally, SMF optic cable provides practically unlimited bandwidth expandability and the ability to transmit Ethernet signals considerably longer distances.

The most commonly used cable today is 10Base-FL (L stands for link), which is part of the 1993 IEEE 802.3j standard. The 1993 standard updated the older DIX standard for Fiber Optic Interrepeater Link (FOIRL) that described Ethernet segments, or links, between repeaters.

The newest and largest Ethernet networking configuration is 1000Base-X or Gigabit Ethernet, with 10 Gig Ethernet looming on the horizon (IEEE 802.3ae was released in June 2002). Short wavelength (1000Base-SX) and long wavelength (1000Base-LX) are based on an 8B/10B block encoding scheme.

> ***Note:*** *8B/10B stands for 8 bits of data transmitted in a 10-bit sequence where the last two additional bits are used for signal and control functions.*

The proliferation of media not only exemplifies Ethernet's growth and flexibility, it is also an example of Ethernet's complexity.

Ethernet Networks

There are several types of Ethernet networks in use throughout the world. They are similar in that they are collections of independent computers that communicate and share resources by using transmission hardware; devices to connect and control signals; and software to decode and format data, as well as detect and correct errors. Two of the most important Ethernet networks in use today are:

- Local area network
- Wide area network

LAN

A LAN is usually confined to a geographical area such as a building or a campus. While closely confined geographically, a LAN can be complex, consisting of hundreds of computers and thousands of users. Local area networks are ubiquitous because Ethernet (the most popular physical layer LAN technology) is relatively inexpensive and Ethernet standards are so widely accepted.

WAN

A WAN is a collection of LANs that are connected in multiple ways using a variety of services. Geography does not limit the LAN intercommunication of WANs. Everything from phone lines to satellite links are used to create WANs.

Private Networks

The Metro Ethernet Forum (MEF) has specified two Ethernet services that replace traditional static permanent virtual circuits with Ethernet virtual circuits (EVCs). These two services are:

- ELINE—Point-to-point
- ELAN—Multipoint-to-multipoint

Both services are offered as private or virtual. Private services are transported over dedicated lines and connections at fixed and predefined wire speeds. Virtual services are multiplexed statistically and transported over shared lines and connections.

ELINE

This point-to-point service is transparent to layer-2 protocols and supports scalability and bandwidth management. ELINE also supports CoS and VLAN tagging.

Two ELINE connectivity services are:

- Ethernet Private Line (EPL)—Port-based, dedicated bandwidth
- Ethernet Virtual Private Line (EVPL)—VLAN-based, shared bandwidth

ELAN

The multipoint-to-multipoint service allows hosts to be connected and disconnected dynamically. Each ELAN host is given its own bandwidth profile, which specifies its CoS.

Two ELAN connectivity services are:

- Ethernet Private LAN (EPLAN)—Port-based, dedicated bandwidth
- Ethernet Virtual Private LAN (EVPLAN)—VLAN-based, shared bandwidth

Virtual private networks (VPNs) use public telecommunications networks for private data communications. VPNs usually follow a client/server approach to authenticate users, encrypt data, and manage transmissions.

Metro Ethernet Forum

The MEF is a non-profit organization that was chartered to accelerate the worldwide adoption of carrier-class Ethernet networks and services. ELINE and ELAN are two MEF priorities.

The MEF defined service attributes and parameters for successful implementation. It does not create standards, but tries to leverage current standards through the following:

- Marketing—White papers, demonstrations, case studies
- Test procedures—Methodologies and procedures
- Technical specifications—Defining new standards
- Position papers—Request for additional technical work by standards bodies (such as IEEE, IETF, ITU)

Figure 3: Tagged Frame Format

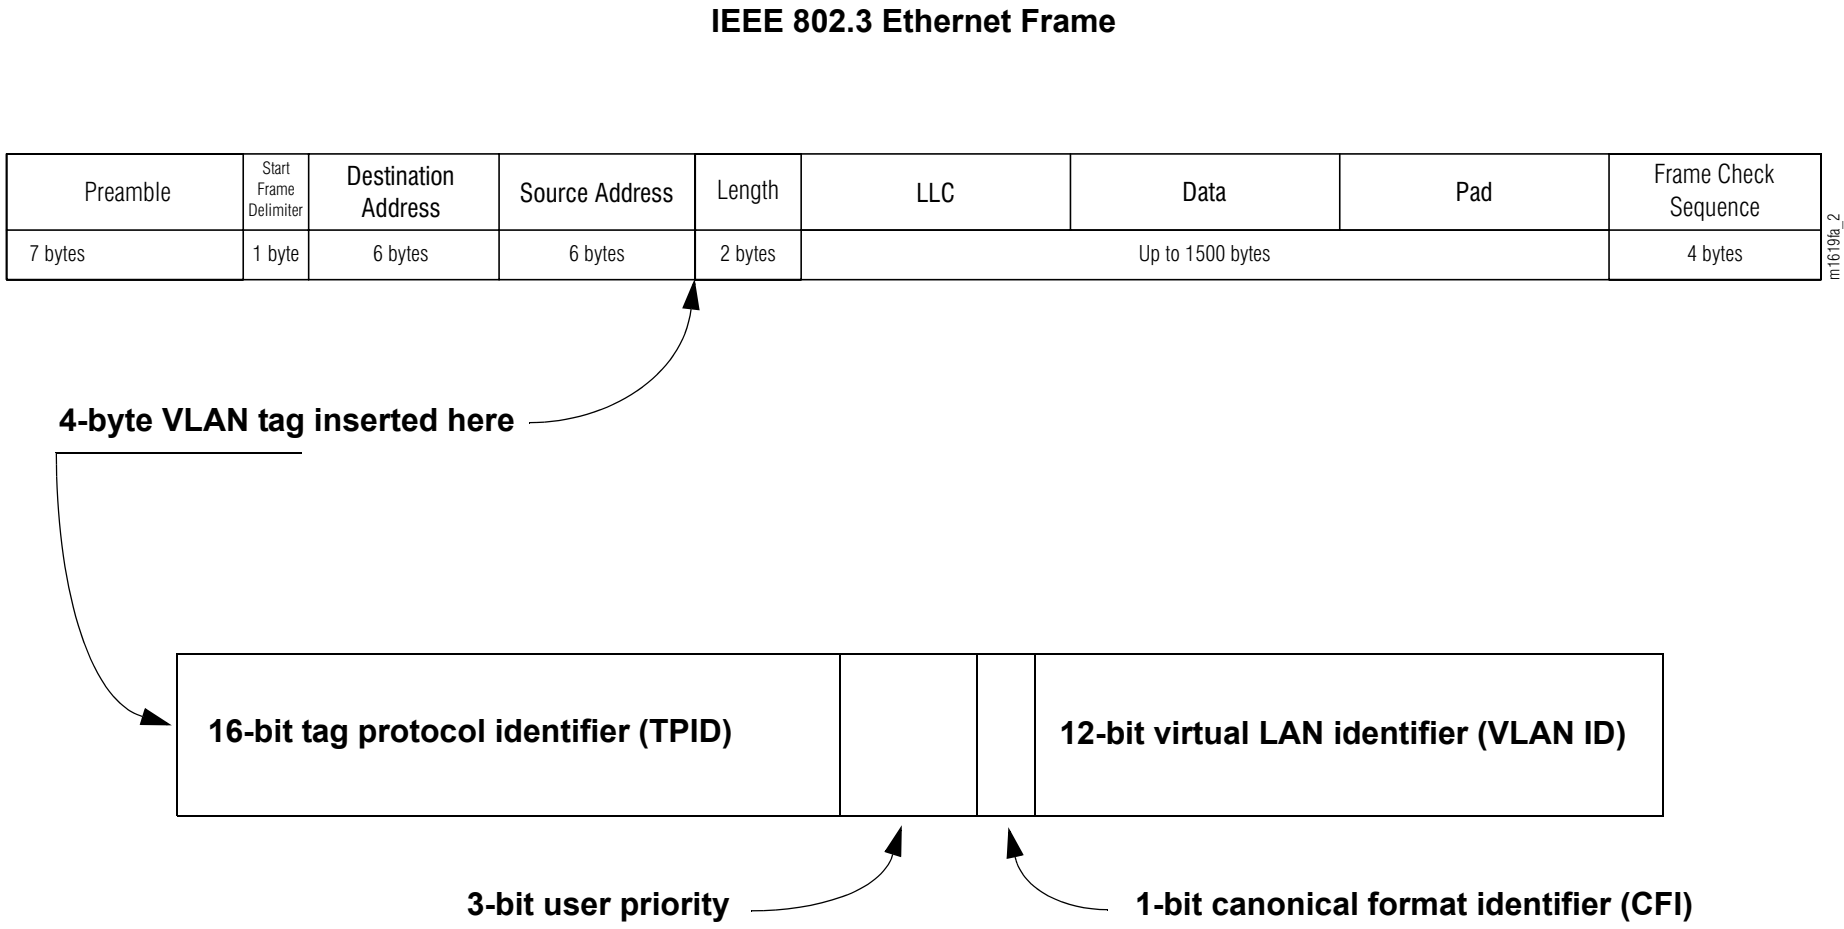

VLAN

A virtual LAN (VLAN) is a group of PCs, servers, and other network resources that behave as if they were connected to a single network segment—even though they may not be. Personnel who are physically separated from each other are virtually connected on a VLAN. A VLAN is often easier to design, administer, and manage than a physical LAN. Furthermore, a VLAN provides better quality control and security because the network can be segmented more effectively.

VLAN Tagging

A VLAN tag is 4-bytes long and is inserted in the Ethernet frame after the source address (see Figure 3). The first two bytes or 16 bits of the VLAN tag identify the tag as an 802.1q protocol tag. This tag is always set at 0x8100 and serves as an indication that the length/type field has moved four bytes further down the frame.

Note: *The addition of the VLAN tag extends the maximum length of an Ethernet frame from 1518 bytes to 1522 bytes.*

The last two bytes of the VLAN contain tag control information. The first three bits of the third byte in the VLAN tag make up the user priority field that assigns a priority level to the Ethernet frame. Binary coding is used to specify the class of service the VLAN frame should be given. Eight classes of service have been identified for VLAN frames.

The next bit in the VLAN tag is the canonical format indicator (CFI). The CFI, if set to 1, indicates the presence of a Routing Information Field (RIF). The RIF indicates the route the frame will take through the network and consists of the following:

- Routing control field
- Route descriptor field

The routing control field contains broadcast indicators that lay out one of the following paths the frame will take:

- Along a specified path (nonbroadcast)
- Through all bridges (all route broadcast)
- Through designated bridges for single appearance per network (single-route broadcast)

The routing control field also has a direction bit. This bit tells the bridge how to read the route descriptor (first to last or last to first). The route descriptor shows the path the frame will travel by calling out the next bridge number.

The last 12 bits in the VLAN tag make up the VLAN ID field that identifies the virtual circuit to which the frame belongs. The VLAN ID field allows the assignment of up to 4096 VLAN numbers to distinguish the different VLAN tagged packets.

Note: *More information is in the VLAN tag standard, IEEE 802.3ac.*

Figure 4: Ethernet Topologies

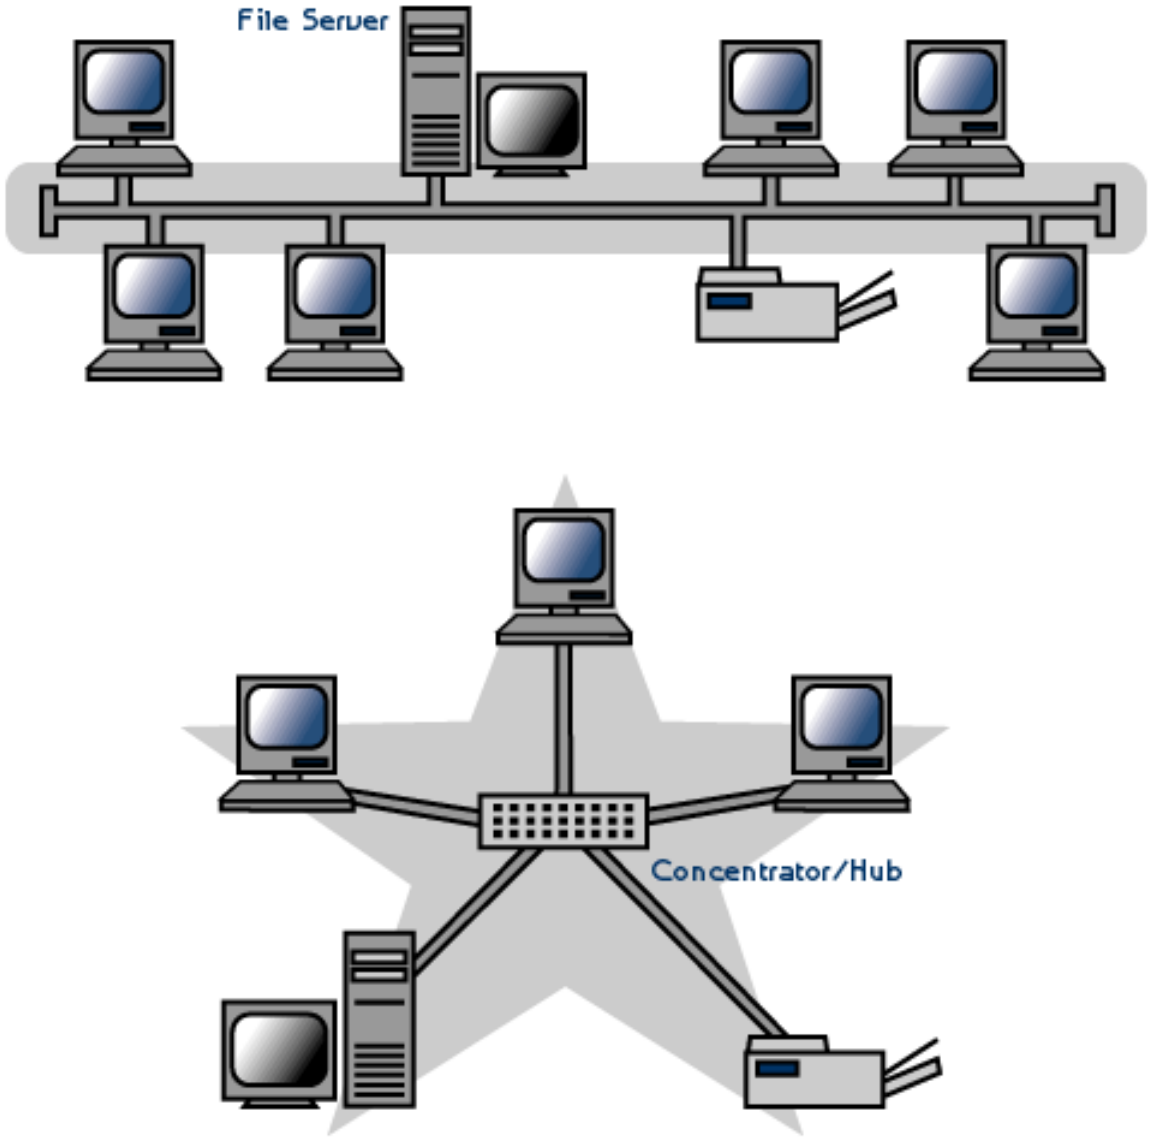

Ethernet Topologies

A network topology is the geometric arrangement of nodes and cable links in a LAN. Devices on an Ethernet network are arranged in either a bus or star topology.

A bus topology (see Figure 4, top) is configured so that all devices on the network connect to one trunk cable. This type of arrangement is easy to configure and install, and it is relatively inexpensive. Bus topology requires no amplification or regeneration equipment, and all devices on the network have access to the bus. After waiting to determine an open line, a signal can be sent by any device along the network. It is expected that any signal is terminated at the end of the trunk.

One potential cost of this less expensive topology is that all devices are affected if the trunk cable fails. For this reason, more recent Ethernet networks are configured as stars.

Separate cables connect each device in a star topology (see Figure 4, bottom) with a central device—usually a hub. Because of the separate cabling requirement, only one transmitting device is affected if a cable fails. Other advantages of a star topology over a bus are:

- Ease of expansion
- Localized troubleshooting
- Support of multiple cable types

The centralized hub can be passive or active. If it is active, the hub regenerates the signal and extends the length of the network. Installing a centralized hub is made easier because it can be housed in a closet with telephone equipment, and cables attached to the hub can run through the same conduits as phone lines.

Tree Topology

A tree topology connects multiple star topologies through a bus topology. A tree topology allows for the expansion of an existing network while supporting point-to-point wiring for individual segments.

Tree topology usually applies the 5-4-3 rule in its Ethernet protocol. Because each Ethernet signal must reach its destination within a specified length of time, the 5-4-3 rule is exercised:

- A maximum of 5 segments are allowed between two nodes.
- A maximum of 4 repeaters/concentrator connections are allowed.
- A maximum of 3 segments can be populated if supported by coaxial cable.

Note: *A populated segment has one or more nodes attached to it.*

Figure 5: SONET Frame

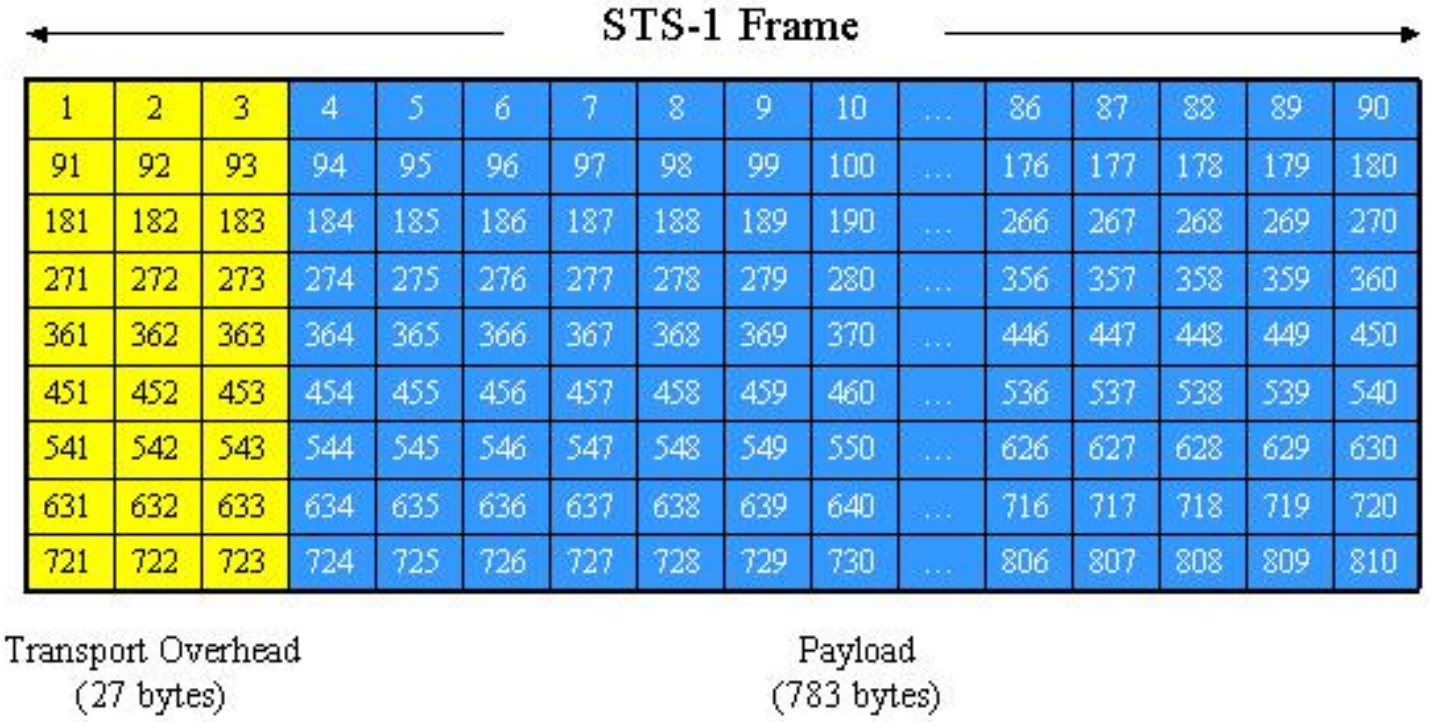

Ethernet Over SONET

Because of the proliferation of LANs that require the ability to intercommunicate with other LANs across the country and around the world, Ethernet carriers have struggled with the best solution for handling all-Ethernet traffic. One new means of transporting Ethernet traffic is metro area networks (MANs). These are being developed as stepping stones between LANs and WANs.

After weighing cost, distance, bandwidth, and traffic management requirements, several solutions were suggested and implemented:

- Ethernet over wavelengths (EOW)
- Ethernet over SONET/SDH (EOS)
- Optical Ethernet (native Ethernet over long-haul fiber)
- Ethernet in the first mile (EFM) over copper or fiber

It soon became clear that EOS was emerging as the most popular choice as the MAN connection. Flexible and guaranteed bandwidth puts EOS ahead of the others.

Note: *Much of the discussion about Ethernet over SONET also applies to Ethernet over Synchronous Digital Hierarchy (SDH).*

The MAN/WAN Connection

Ethernet over SONET MAN/WAN connections can solve many business problems associated with moving outside LAN environments, such as:

- Connecting remote offices so that everyone has access to the same information at the same high LAN speed using flexible and guaranteed bandwidth.
- Accessing application service providers and remote storage sites with enough bandwidth to quickly upload and download files.
- Multiuser access to Web sites is important so customers can transact business quickly and easily. Employee access to the Internet also requires flexible and guaranteed bandwidth.

Note: *Ethernet frames fit into the payload side of an STS-1 frame (see Figure 5). Multiple STS-1 frames are used for Fast Ethernet and Gigabit Ethernet frames.*

Figure 6: VT-Mapped EOS

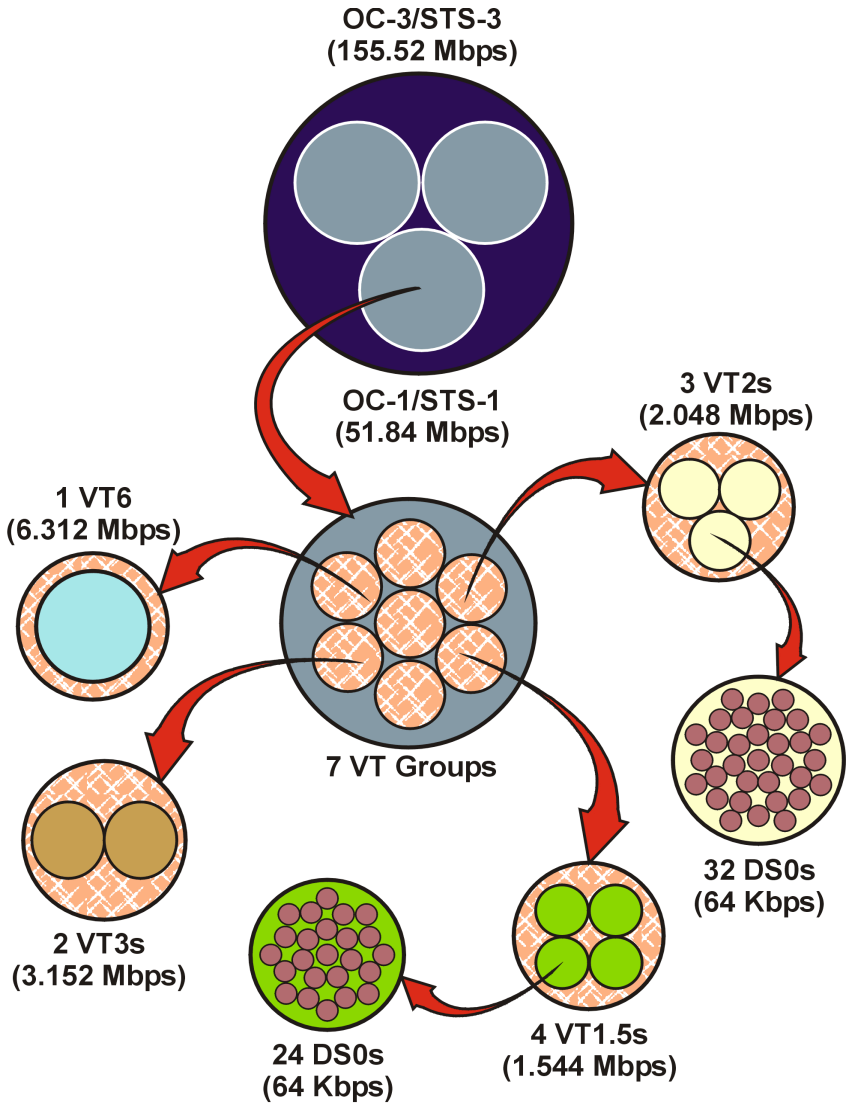

Encapsulated Ethernet

The Synchronous Optical Network (SONET) is structured to allow for a flexible synchronous optical hierarchy to carry signals at various rates. The basic synchronous transport signal-1 (STS-1) consists of lower order virtual tributaries (VTs) that can be subdivided into as many as 28 VT1.5s. Conversely, STS-1s can be put together individually or even concatenated for transport in higher-level optical carrier (OC) networks.

Two concatenation methods are used for higher-order paths:

- Pointer based—For an STS-3c, three STS-1s are joined together and are transported in identical phase. (Path overhead from only one STS is needed, allowing for more payload capacity; the pointers needed to separate the individual STSs reside in the line overhead).
- Virtual concatenation (VC)—Portions of an STS-1 payload are split over multiple STSs and can follow physically separate routes to the final destination.

Virtual tributaries come in a variety of sizes and bandwidths (see Figure 6):

- VT1.5—Carries a T1, a DS1, or 24 DS0s (28 in an STS)
- VT2—Carries an E1 or 32 DS0s (21 in an STS)
- VT3—Essentially a concatenated VT1.5 (14 in an STS)
- VT6—Carries a DS2 (seven in an STS)

SONET has specific concatenation rules; however, they apply only at the path or end user connections on the network. Virtual concatenation bends these rules to efficiently accommodate Ethernet signals transported over SONET networks.

Concatenated VTs

To allow finer bandwidth mapping of SONET transport to the data rate requirement, VT mapping supports the transport of up to 28 VT1.5s with a bandwidth of 43 Mb/s but scalable to 1.544 Mb/s.

Systems are available that provide point-to-point transport over existing SONET networks for native LAN interfaces at rates of 10 Mb/s and 100 Mb/s, using half-duplex or full-duplex applications. Up to 28 individual VT1.5s can be provisioned to specific ports on an Ethernet bridge to support 100 Mb/s networks, or seven VT1.5s can be used for 10 Mb/s systems.

Unlike normal VT concatenation, virtually concatenated VTs are provisioned and transported independently of each other. However, they are permanently connected on the VT switch fabric on an OC line unit, across the backplane, and over a point-to-point connection to a VT-mapped EOS module.

Differential Delay

Because VTs travel separately across the SONET network, they arrive at different times and out of order. Two methods are used at the receiving end to reduce or resolve this problem:

- A sliding window algorithm processes and resequences VTs.
- Up to 4 ms of delay is acceptable between receipt of each VT.

Provisioning plays an important role in the efficient transport of concatenated VTs. When possible, VTs should be provisioned as a whole to ensure transport over the same physical path. All automatic protection systems should be set to revertive to keep all VTs on a primary path and stay within differential delay limits.

Figure 7: EOS Protocol Frames

Link Address Procedure - SDH

Flag	Address	Control	SAPI	Ethernet MAC Frame	FCS	Closing Flag
1 byte	1 byte	1 byte	2 bytes	64-1522 bytes	4 bytes	1 byte

m1619ff_2

Generic Framing Procedure

PLI	cHEC	Type	tHEC	Extend	eHEC	Ethernet MAC Frame	FCS
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	64-1522 bytes	4 bytes

m1619jk_1

EOS Protocols

There are two standardized protocols for transmitting Ethernet over a SONET network:

- Link Access Procedure—SDH (LAPS)
- Generic Framing Procedure (GFP)

LAPS

Simply put, LAPS is a method of communication with no single controller. Either of the communication devices connected together on the line can initiate a data transfer operation.

LAPS is a high-level data link control (HDLC)-like protocol that fits into a SONET payload. Its frames are formatted for high-layer (network and transport) protocol encapsulation.

The LAPS frame structure is illustrated in Figure 7 and includes:

- Flags—01111110 at the start and end of each frame
- Address and Control—1 byte each
- Source Access Point Identifier (SAPI)—Connection point between data link and network layers
- Ethernet frame (*Refer to page 8.*)
- Frame check sequence (FCS)—Error checking

GFP

Generic Framing Procedure (GFP) is an efficient generic enveloping protocol for adapting multiple types of packet traffic to SONET. The adaption mechanisms GFP uses are:

- Frame-mapped GFP (a protocol data unit [PDU]-oriented adaption for IP/point-to-point [IP/PPP] or Ethernet MAC frames)
- Transparent GFP (a block code-oriented adaptation for 8B/10B users such as Gigabit Ethernet)

Frame-Mapped GFP

At the Ethernet/SONET interface, the mandatory idle time in each Ethernet MAC address is stripped out and the frames are bundled more tightly together and fit into virtual containers and transmitted over the SONET network. At the SONET/Ethernet interface, the bundles are separated and the idle time is put back into the Ethernet MAC frames.

Transparent GFP

Transparent GFP takes 8B/10B encoded signals and compresses them into 64B/66B super blocks to more efficiently use virtual container bandwidth. This compression method minimizes latency across the SONET network.

GFP Frame Structure

- PDU length indicator (PLI) and core header error check (cHEC)—Used to delineate frames and support Level 2 functions (see Figure 7)
- Type and tHEC—Indicates whether this frame is a data or management frame
- Extension and eHEC—Indicates if the frame is a null header, a linear frame or a ring frame
- Ethernet frame
- FCS—Error checking

GFP is considered a better option than LAPS because it is more robust, more bandwidth-efficient, provides more control, and is better suited for Ethernet over SONET traffic.

EOS Advantages

The best physical layer choice is SONET because it meets all the requirements:

- Reliability—SONET puts overhead bytes in every frame, constantly monitoring for problems that might affect the health of the signal. SONET splits up its overhead among section, line, and path, so fault identification is easier and recovery is quicker.
- Recovery—The SONET fault-protection scheme is well-defined. Automatic protection switching has a standardized mandatory rate of less than 50 ms.
- Consistent service—Even over long distances, jitter is controlled because of tough specifications for all SONET-standard components. The result is sustained available bandwidth.
- Bandwidth management—The SONET time-division multiplexing (TDM) structure allows bandwidth to be added and dropped at every node in the network. With SONET, provisioning and management is relatively simple and economical.
- Bandwidth flexibility—While this was not a strong suit for SONET, virtual concatenation has solved many flexibility issues. Without having to bundle virtual tributaries together, they are virtually connected even though they travel independently to their common destination.

Ethernet and SONET complement each other:

- Ethernet, relatively inexpensive for everyone to afford, is tied into a billion dollars worth of deployed SONET infrastructure.
- Ethernet is scalable from 1 Mb/s to 10 Gb/s and SONET can handle all of it.
- Ethernet alone may not be quickly restored, while the SONET restoration rate is less than 50 ms.

New Ethernet and SONET technologies will enable these protocols to coexist and commingle well in the future.

Ethernet Acronyms

ANSI	American National Standards Institute
BPDU	Bridge protocol data unit
CFI	Canonical format indicator
cHEC	Core header error check
CIDR	Classless interdomain routing
CRC	Cyclical redundancy check
CSMA/CD	Carrier Sense Multiple Access with Collision Detection
CWDM	Course wave division multiplexing
DIX	Digital, Intel, Xerox
eHEC	Extension header error check
EFM	Ethernet in the first mile
EOS	Ethernet over SONET
EOW	Ethernet over wavelength
EPL	Ethernet private line
EPLAN	Ethernet private LAN
EVC	Ethernet virtual circuit
EVPL	Ethernet virtual private line
EVPLAN	Ethernet virtual private LAN
FCS	Frame check sequence
FOIRL	Fiber-optic interrepeater link
Gb/s	Gigabits per second
GFP	Generic framing procedure
HDLC	High-level data link control
HEC	Header error check
IEEE	Institute of Electrical and Electronics Engineers
IP	Internet Protocol
ISO	International Organization for Standardization
LAN	Local area network
LAPS	Link access procedure—SDH
LLC	Logical link control
MAC	Medium access control
MAN	Metro area network
MDI	Medium dependent interface
MEF	Metro Ethernet Forum
MMF	Multimode fiber
ms	Millisecond
NIC	Network interface card
OC	Optical carrier
OSI	Open Systems Interconnection
OUI	Organizational unique identifier
PCS	Physical coding sublayer
PDU	Protocol data unit
PLI	PDU length indicator
PMA	Physical medium attachment

Ethernet Acronyms (Cont)

PMD	Physical medium dependent
PPP	Point-to-Point Protocol
RIF	Routing information field
RS	Reconciliation sublayer
RSTP	Rapid Spanning Tree Protocol
SAPI	Source access point identifier
SDH	Synchronous Digital Hierarchy
SERDES	Serializer/deserializer
SMF	Single mode fiber
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
STP	Spanning Tree Protocol
STS	Synchronous Transport Signal
TCP	Transmission Control Protocol
TDM	Time-division multiplexing
tHEC	Type header error check
UTP	Unshielded twisted pair
VC	Virtual channel
VLAN	Virtual LAN
VT	Virtual tributary
WAN	Wide area network
WIS	WAN interface sublayer
XAUI	10 Gigabit Ethernet attachment unit interface
XGMII	10 Gigabit Ethernet media-independent interface
XSBI	10 Gigabit Ethernet 16-bit interface

Tutorial Review

Instructions

Please read the questions below and mark your answers on the answer sheet.

1. A group of PCs, servers and other network resources that behave as if they are connected to a single network segment is called a _____.
   a. LAN
   b. MAN
   c. WAN
   d. VLAN

2. The line rate for the first Ethernet was _____.
   a. 54-kb/s
   b. 9600 baud
   c. 10 Mb/s
   d. 2.94 Mb/s

3. _____ is configured with one trunk cable.
   a. Monocabling
   b. Centralized topography
   c. Unitrunking
   d. Bus topology

4. The Fast Ethernet standard applies to _____ Ethernet.
   a. 10 Mb/s
   b. 1000 Mb/s
   c. 100 Mb/s
   d. 10 gigabit

5. The central device in a star topology is a _____.
   a. router
   b. hub
   c. bridge
   d. switch

6. Another term for a hub is a _____.
   a. repeater
   b. transmitter
   c. switch
   d. bridge

7. A healthy Ethernet network should have less than ___% of its traffic dedicated to retransmissions from collisions.
   a. 3
   b. 5
   c. 13
   d. 15

8. The unique Ethernet address is embedded in the _____ by the manufacturer.
    a. NIC
    b. MAC
    c. HEC
    d. OUI

9. _____ concentrate on network function instead of physical functions.
    a. Bridges
    b. Switches
    c. Hubs
    d. Routers

10. There are ____ VT1.5s in an STS-1.
    a. 32
    b. 21
    c. 28
    d. 24

11. How many VT1.5s are required for 10 Mb/s Ethernet over SONET?
    a. One
    b. Seven
    c. Ten
    d. None of the above

12. IP resides in the OSI _____ layer.
    a. physical
    b. network
    c. data link
    d. transport

13. VTs that travel without being joined together are _____.
    a. virtually concatenated
    b. concatenated
    c. virtual containers
    d. virtual tributaries

14. An STS-1 contains enough VT1.5s to support up to ___ 10 Mb/s Ethernet network(s).
    a. one
    b. two
    c. three
    d. four

15. Concatenated STSs are separated using pointers carried in _____ overhead in the STS frame.
    a. section
    b. line
    c. path
    d. VT

16. A sliding window algorithm is used to process and resequence VTS and reduce _____.
    a. differential delay
    b. collisions
    c. link access processing
    d. set nullification

17. 28 VT1.5s provide _____ Mb/s of usable bandwidth.
    a. 54
    b. 43
    c. 50
    d. 28

18. The process that compresses 8B/10B encoded signals into 64B/66B super blocks is called _____.
    a. LAPS
    b. transparent GFP
    c. frame-mapped GFP
    d. PLI

19. SONET APS must occur within _____.
    a. 5 ms
    b. 50 ms
    c. 100 ms
    d. 5 min.

20. Which is not a way to reduce collisions on an Ethernet network?
    a. Reduce the number of users
    b. Reduce the number of PCs
    c. Reduce the number of collision domains
    d. Reduce the number of Ethernet frames

21. _____ is used to prevent loops on an Ethernet network.
    a. Classless interdomain routing
    b. Spanning Tree Protocol
    c. A dot address
    d. A MAC address

22. A unique network number is requested from the _____.
    a. Metro Ethernet Forum
    b. IEEE
    c. ITU-T
    d. Network Information Center

23. VLAN tagging adds _____ to the IEEE 802.3 Ethernet frame.
    a. 2 bytes
    b. 8 bytes
    c. 4 bytes
    d. 1 byte

24. A _____ topology is created when bus and star topologies are connected.
    a. tree
    b. ring
    c. root
    d. branch

25. An IP address has two parts: _____ and _____.
    a. network and node
    b. host and node
    c. host and class
    d. domain and network

The answers and reference pages for all questions are listed on the next page.

Review Answers

Question	Answer	Reference Page
1	**d. VLAN**	23
2	**c. 10 Mb/s**	19
3	**d. Bus topology**	25
4	**c. 100 Mb/s**	4
5	**b. hub**	13
6	**a. repeater**	13
7	**b. 5**	10
8	**a. NIC**	9
9	**d. Routers**	12
10	**c. 28**	29
11	**b. Seven**	29
12	**b. network**	15
13	**a. virtually concatenated**	29
14	**d. four**	29
15	**b. line**	29
16	**a. differential delay**	29
17	**b. 43**	29
18	**b. transparent GFP**	31
19	**b. 50 ms**	32
20	**c. Reduce the number of collision domains**	10
21	**b. Spanning Tree Protocol**	18
22	**d. Network Information Center**	17
23	**c. 4 bytes**	25
24	**a. tree**	27
25	**a. network and node**	17