Ethernet/IEEE 802.3

Background

The term *Ethernet* refers to the family of local area network (LAN) implementations that includes three principal categories.

- Ethernet and IEEE 802.3—LAN specifications that operate at 10 Mbps over coaxial cable.
- 100-Mbps Ethernet—A single LAN specification, also known as Fast Ethernet, that operates at 100 Mbps over twisted-pair cable.
- 1000-Mbps Ethernet—A single LAN specification, also known as Gigabit Ethernet, that operates at 1000 Mbps (1 Gbps) over fiber and twisted-pair cables.

This chapter provides a high-level overview of each technology variant.

Ethernet has survived as an essential media technology because of its tremendous flexibility and its relative simplicity to implement and understand. Although other technologies have been touted as likely replacements, network managers have turned to Ethernet and its derivatives as effective solutions for a range of campus implementation requirements. To resolve Ethernet's limitations, innovators (and standards bodies) have created progressively larger Ethernet pipes. Critics might dismiss Ethernet as a technology that cannot scale, but its underlying transmission scheme continues to be one of the principal means of transporting data for contemporary campus applications. This chapter outlines the various Ethernet technologies that have evolved to date.

Ethernet/IEEE 802.3 Comparison

Ethernet and IEEE 802.3 specify similar technologies. Both are CSMA/CD LANs. Stations on a CSMA/CD LAN can access the network at any time. Before sending data, CSMA/CD stations "listen" to the network to see if it is already in use. If it is, the station wishing to transmit waits. If the network is not in use, the station transmits. A collision occurs when two stations listen for network traffic, "hear" none, and transmit simultaneously. In this case, both transmissions are damaged, and the stations must retransmit at some later time. *Backoff* algorithms determine when the colliding stations retransmit. CSMA/CD stations can detect collisions, so they know when they must retransmit.

Both Ethernet and IEEE 802.3 LANs are broadcast networks. In other words, all stations see all frames, regardless of whether they represent an intended destination. Each station must examine received frames to determine if the station is a destination. If so, the frame is passed to a higher protocol layer for appropriate processing.

Differences between Ethernet and IEEE 802.3 LANs are subtle. Ethernet provides services corresponding to Layers 1 and 2 of the OSI reference model, while IEEE 802.3 specifies the physical layer (Layer 1) and the channel-access portion of the link layer (Layer 2), but does not define a

logical link control protocol. Both Ethernet and IEEE 802.3 are implemented in hardware. Typically, the physical manifestation of these protocols is either an interface card in a host computer or circuitry on a primary circuit board within a host computer.

Physical Connections

IEEE 802.3 specifies several different physical layers, whereas Ethernet defines only one. Each IEEE 802.3 physical layer protocol has a name that summarizes its characteristics. The coded components of an IEEE 802.3 physical-layer name are shown in Figure 7-1.

Figure 7-1 IEEE 802.3 Physical-Layer Name Components

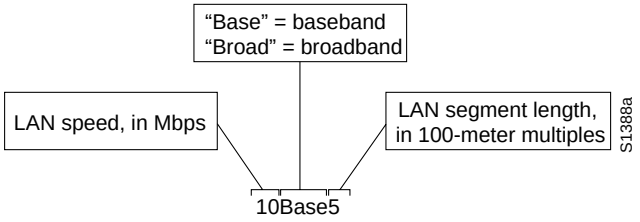

A summary of Ethernet Version 2 and IEEE 802.3 characteristics appears in Table 7-1.

Table 7-1 Ethernet Version 2 and IEEE 802.3 Physical Characteristics

Characteristic	Ethernet Value	IEEE 802.3 Values				
		10Base5	10Base2	1Base5	10BaseT	10Broad36
Data rate (Mbps)	10	10	10	1	10	10
Signaling method	Baseband	Baseband	Baseband	Baseband	Baseband	Broadband
Maximum segment length (m)	500	500	185	250	100 Unshielded twisted-pair wire	1800
Media	50-ohm coax (thick)	50-ohm coax (thick)	50-ohm coax (thin)	Unshielded twisted-pair wire	Unshielded twisted-pair wire	75-ohm coax
Topology	Bus	Bus	Bus	Star	Star	Bus

Ethernet is most similar to IEEE 802.3 10Base5. Both of these protocols specify a bus topology network with a connecting cable between the end stations and the actual network medium. In the case of Ethernet, that cable is called a *transceiver cable*. The transceiver cable connects to a transceiver device attached to the physical network medium. The IEEE 802.3 configuration is much the same, except that the connecting cable is referred to as an *attachment unit interface* (AUI), and the transceiver is called a *medium attachment unit* (MAU). In both cases, the connecting cable attaches to an interface board (or interface circuitry) within the end station.

Ethernet and IEEE 802.3 Frame Formats

Figure 7-2 illustrates the frame fields associated with both Ethernet and IEEE 802.3 frames.

Figure 7-2 Various frame fields exist for both Ethernet and IEEE 802.3.

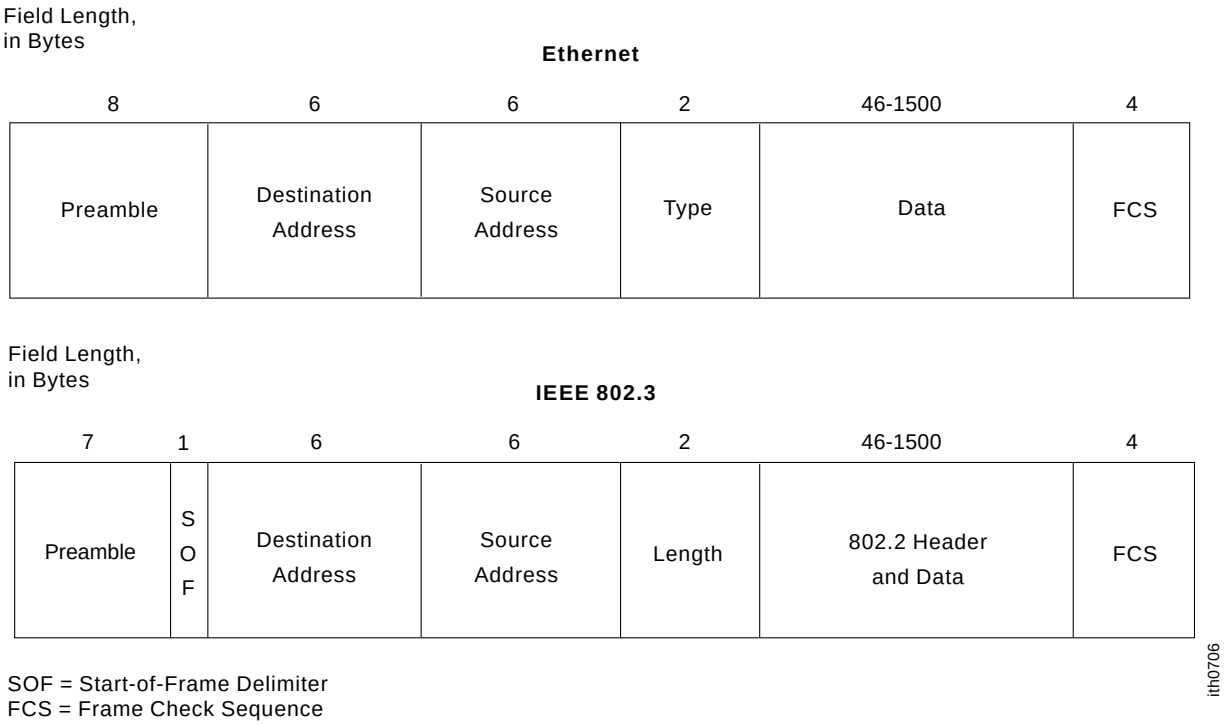

The Ethernet and IEEE 802.3 frame fields illustrated in Figure 7-2 are described in the following summaries:

- *Preamble* — The alternating pattern of ones and zeros tells receiving stations that a frame is coming (Ethernet or IEEE 802.3). The Ethernet frame includes an additional byte that is the equivalent of the Start of Frame (SOF) field specified in the IEEE 802.3 frame.
- *Start-of-Frame (SOF)* — The IEEE 802.3 delimiter byte ends with two consecutive 1 bits, which serve to synchronize the frame-reception portions of all stations on the LAN. SOF is explicitly specified in Ethernet.
- *Destination and Source Addresses* — The first 3 bytes of the addresses are specified by the IEEE on a vendor-dependent basis. The last 3 bytes are specified by the Ethernet or IEEE 802.3 vendor. The source address is always a unicast (single-node) address. The destination address can be unicast, multicast (group), or broadcast (all nodes).
- *Type (Ethernet)* — The type specifies the upper-layer protocol to receive the data after Ethernet processing is completed.
- *Length (IEEE 802.3)* — The length indicates the number of bytes of data that follows this field.
- *Data (Ethernet)* — After physical-layer and link-layer processing is complete, the data contained in the frame is sent to an upper-layer protocol, which is identified in the Type field. Although Ethernet Version 2 does not specify any padding (in contrast to IEEE 802.3), Ethernet expects at least 46 bytes of data.

- *Data (IEEE 802.3)* — After physical-layer and link-layer processing is complete, the data is sent to an upper-layer protocol, which must be defined within the data portion of the frame, if at all. If data in the frame is insufficient to fill the frame to its minimum 64-byte size, padding bytes are inserted to ensure at least a 64-byte frame.
- *Frame Check Sequence (FCS)* — This sequence contains a 4-byte cyclic redundancy check (CRC) value, which is created by the sending device and is recalculated by the receiving device to check for damaged frames.

100-Mbps Ethernet

100-Mbps Ethernet is a high-speed LAN technology that offers increased bandwidth to desktop users in the wiring center, as well as to servers and server clusters (sometimes called server farms) in data centers.

The IEEE Higher Speed Ethernet Study Group was formed to assess the feasibility of running Ethernet at speeds of 100 Mbps. The Study Group established several objectives for this new higher-speed Ethernet but disagreed on the access method. At issue was whether this new faster Ethernet would support CSMA/CD to access the network medium or some other access method.

The Study Group divided into two camps over this access-method disagreement: the Fast Ethernet Alliance and the 100VG-AnyLAN Forum. Each group produced a specification for running Ethernet (and Token Ring for the latter specification) at higher speeds: 100BaseT and 100VG-AnyLAN, respectively.

100BaseT is the IEEE specification for the 100-Mbps Ethernet implementation over unshielded twisted-pair (UTP) and shielded twisted-pair (STP) cabling. The Media Access Control (MAC) layer is compatible with the IEEE 802.3 MAC layer. Grand Junction, now a part of Cisco Systems Workgroup Business Unit (WBU), developed Fast Ethernet, which was standardized by the IEEE in the 802.3u specification.

100VG-AnyLAN is an IEEE specification for 100-Mbps Token Ring and Ethernet implementations over 4-pair UTP. The MAC layer is *not* compatible with the IEEE 802.3 MAC layer. 100VG-AnyLAN was developed by Hewlett-Packard (HP) to support newer time-sensitive applications, such as multimedia. A version of HP's implementation is standardized in the IEEE 802.12 specification.

100BaseT Overview

100BaseT uses the existing IEEE 802.3 CSMA/CD specification. As a result, 100BaseT retains the IEEE 802.3 frame format, size, and error-detection mechanism. In addition, it supports all applications and networking software currently running on 802.3 networks. 100BaseT supports dual speeds of 10 and 100 Mbps using 100BaseT fast link pulses (FLPs). 100BaseT hubs must detect dual speeds much like Token Ring 4/16 hubs, but adapter cards can support 10 Mbps, 100 Mbps, or both. Figure 7-3 illustrates how the 802.3 MAC sublayer and higher layers run unchanged on 100BaseT.

Figure 7-3 802.3 MAC and higher-layer protocols operate over 100BaseT.

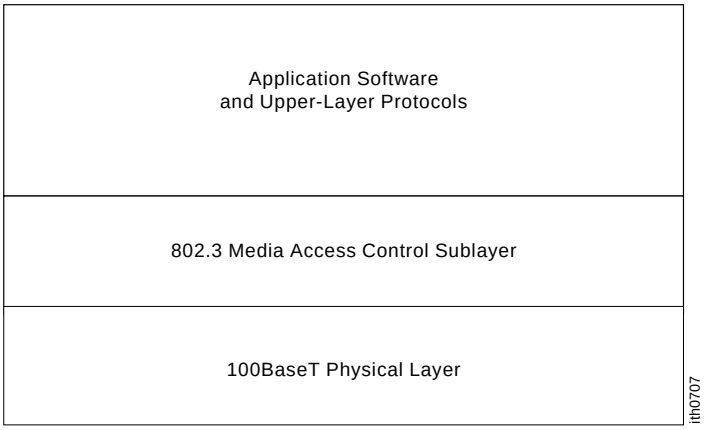

100BaseT Signaling

100BaseT supports two signaling types:

- 100BaseX
- 4T+

Both signaling types are interoperable at the station and hub levels. The Media Independent Interface (MII), an AUI-like interface, provides interoperability at the station level. The hub provides interoperability at the hub level.

The 100BaseX signaling scheme has a convergence sublayer that adapts the full-duplex continuous signaling mechanism of the FDDI physical medium dependent (PMD) layer to the half-duplex, start-stop signaling of the Ethernet Media Access Control (MAC) sublayer. 100BaseTX's use of the existing FDDI specification has allowed quick delivery of products to market. 100BaseX is the signaling scheme used in the 100BaseTX and the 100BaseFX media types. Figure 7-4 illustrates how the 100BaseX convergence sublayer interfaces between the two signaling schemes.

Figure 7-4 The 100BaseX convergence sublayer interfaces two signaling schemes.

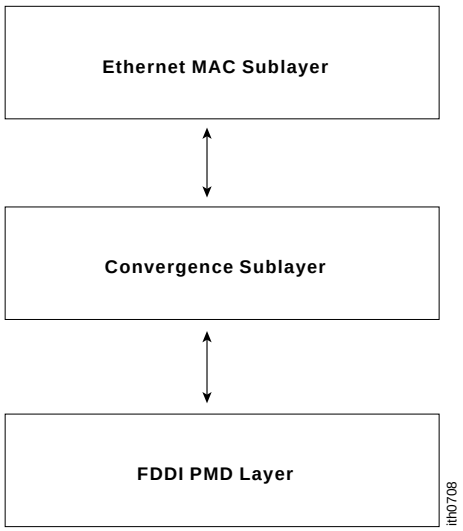

The 4T+ signaling scheme uses one pair of wires for collision detection and the other three pairs to transmit data. It allows 100BaseT to run over existing Category 3 cabling if all four pairs are installed to the desktop. 4T+ is the signaling scheme used in the 100BaseT4 media type, and it supports half-duplex operation only. Figure 7-5 shows how 4T+ signaling requires all four unshielded twisted-pair (UTP) pairs.

Figure 7-5 4T+ requires four UTP pairs.

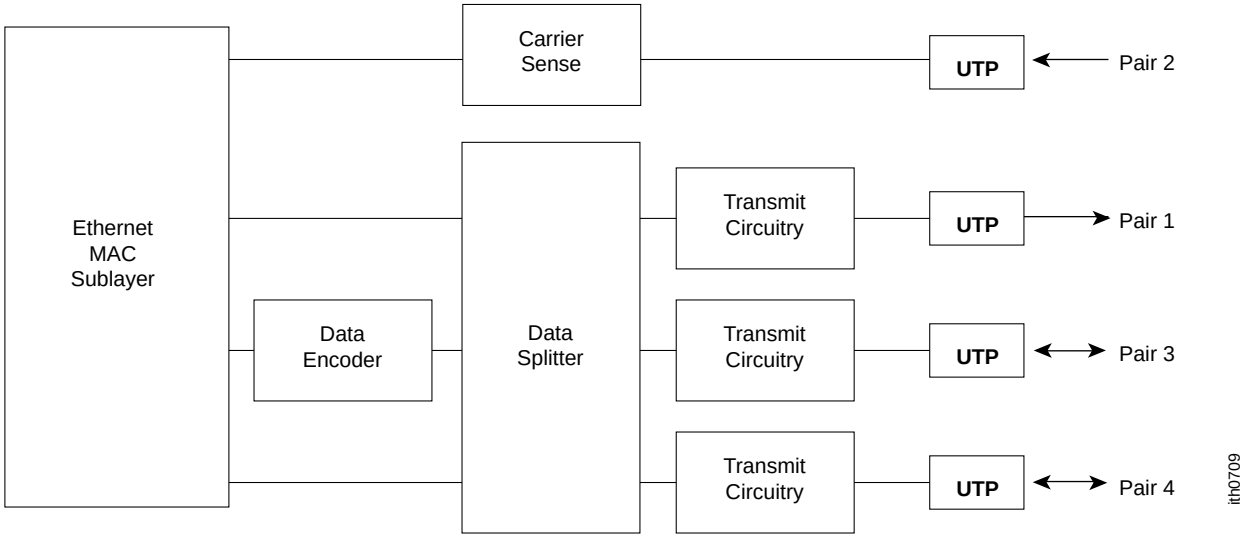

100BaseT Hardware

Components used for a 100BaseT physical connection include the following:

- *Physical Medium*—This device carries signals between computers and can be one of three 100BaseT media types:
- 100BaseTX
- 100BaseFX
- 100BaseT4
- *Medium-Dependent Interface (MDI)*— The MDI is a mechanical and electrical interface between the transmission medium and the physical-layer device (PHY).
- *Physical-Layer Device (PHY)* — The PHY provides either 10- or 100-Mbps operation and can be a set of integrated circuits (or a daughter board) on an Ethernet port, or an external device supplied with a Medium Independent Interface (MII) cable that plugs into an MII port on a 100BaseT device (similar to a 10-Mbps Ethernet transceiver).
- *Medium-Independent Interface (MII)*— The MII is used with a 100-Mbps external transceiver to connect a 100-Mbps Ethernet device to any of the three media types. The MII has a 40-pin plug and cable that stretches up to 0.5 meters.

Figure 7-6 depicts the 100BaseT hardware components.

Figure 7-6 100BaseT requires several hardware components.

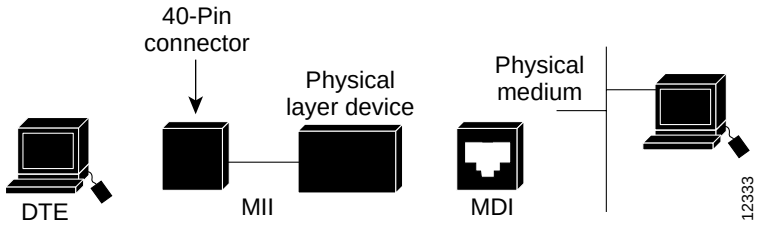

100BaseT Operation

100BaseT and 10BaseT use the same IEEE 802.3 MAC access and collision-detection methods, and they also have the same frame format and length requirements. The main difference between 100BaseT and 10BaseT (other than the obvious speed differential) is the network diameter. The 100BaseT maximum network diameter is 205 meters, which is approximately 10 times less than 10-Mbps Ethernet.

Reducing the 100BaseT network diameter is necessary because 100BaseT uses the same collision-detection mechanism as 10BaseT. With 10BaseT, distance limitations are defined so that a station knows while transmitting the smallest legal frame size (64 bytes) that a collision has taken place with another sending station that is located at the farthest point of the domain.

To achieve the increased throughput of 100BaseT, the size of the collision domain had to shrink. This is because the propagation speed of the medium has not changed, so a station transmitting 10 times faster must have a maximum distance that is 10 times less. As a result, any station knows within the first 64 bytes whether a collision has occurred with any other station.

100BaseT Fast-Link Pulses (FLPs)

100BaseT uses pulses, called fast-link pulses (FLPs), to check the link integrity between the hub and the 100BaseT device. FLPs are backward-compatible with 10BaseT normal-link pulses (NLPs). But FLPs contain more information than NLPs and are used in the autonegotiation process between a hub and a device on a 100BaseT network.

100BaseT Autonegotiation Option

100BaseT networks support an optional feature, called autonegotiation, that enables a device and a hub to exchange information using (100BaseT FLPs) about their capabilities, thereby creating an optimal communications environment.

Autonegotiaton supports a number of capabilities, including speed matching for devices that support both 10- and 100-Mbps operation, full-duplex mode of operation for devices that support such communications, and an automatic signaling configuration for 100BaseT4 and 100BaseTX stations.

100BaseT Media Types

100BaseT supports three media types at the OSI physical layer (Layer 1): 100BaseTX, 100BaseFX, and 100BaseT4. The three media types, which all interface with the IEEE 802.3 MAC layer, are shown in Figure 7-7. Table 7-2 compares key characteristics of the three 100BaseT media types.

Figure 7-7 Three 100BaseT media types exist at the physical layer.

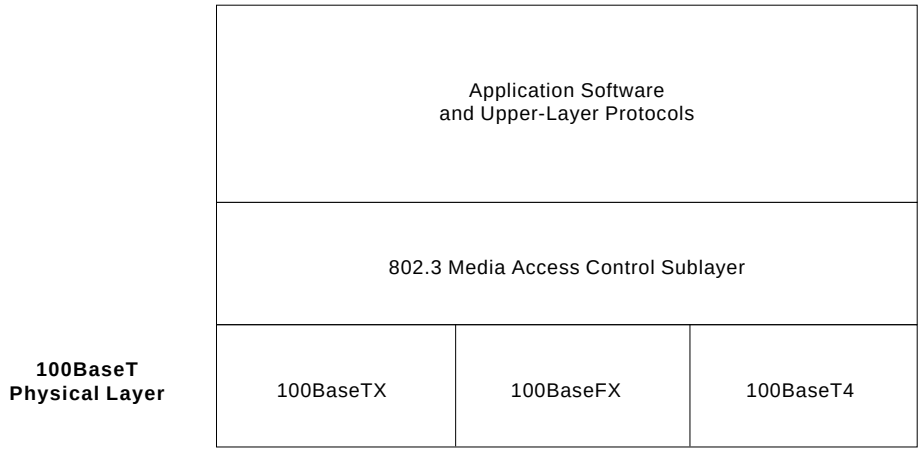

Table 7-2 Characteristics of 100BaseT Media Types

Characteristics	100BaseTX	100BaseFX	100BaseT4
Cable	Category 5 UTP, or Type 1 and 2 STP	62.5/125 micron multi-mode fiber	Category 3, 4, or 5 UTP
Number of pairs or strands	2 pairs	2 strands	4 pairs
Connector	ISO 8877 (RJ-45) connector	Duplex SCmedia-interface connector (MIC) ST	ISO 8877 (RJ-45) connector
Maximum segment length	100 meters	400 meters	100 meters
Maximum network diameter	200 meters	400 meters	200 meters

100BaseTX

100BaseTX is based on the American National Standards Institutes (ANSI) Twisted Pair-Physical Medium Dependent (TP-PMD) specification. The ANSI TP-PMD supports unshielded twisted-pair (UTP) and shielded twisted-pair (STP) cabling. 100BaseTX uses the 100BaseX signaling scheme over 2-pair Category 5 UTP or STP.

The IEEE 802.3u specification for 100BaseTX networks allows a maximum of two repeater (hub) networks and a total network diameter of approximately 200 meters. A link segment, which is defined as a point-to-point connection between two Medium Independent Interface (MII) devices, can be up to 100 meters. Figure 7-8 illustrates these configuration guidelines.

Figure 7-8 The 100BaseTX is limited to a link distance of 100 meters.

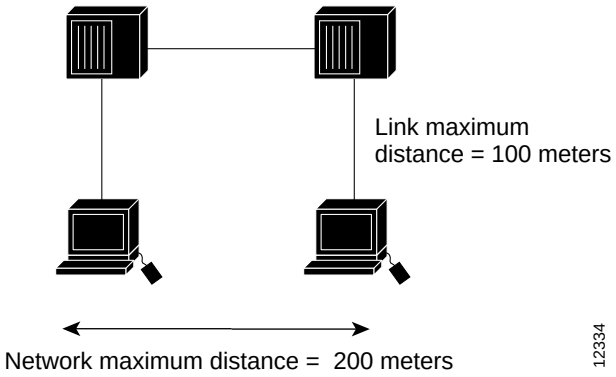

100BaseFX

100BaseFX is based on the ANSI Twisted Pair-Physical Medium Dependent (TP-PMD) X3T9.5 specification for Fiber Distributed Data Interface (FDDI) LANs. 100BaseFX uses the 100BaseX signaling scheme over two-strand multi-mode fiber-optic (MMF) cable. The IEEE 802.3u specification for 100BaseFX networks allows data terminal equipment (DTE)-to-DTE links of approximately 400 meters, or one repeater network of approximately 300 meters in length. Figure 7-9 illustrates these configuration guidelines.

Figure 7-9 The 100BaseFX DTE-to-DTE limit is 400 meters.

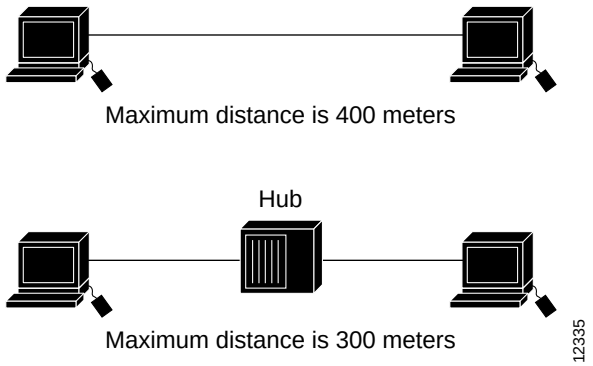

100BaseT4

100BaseT4 allows 100BaseT to run over existing Category 3 wiring, provided that all four pairs of cabling are installed to the desktop. 100BaseT4 uses the half-duplex 4T+ signaling scheme. The IEEE 802.3u specification for 100BaseT4 networks allows a maximum of two repeater (hub) networks and a total network diameter of approximately 200 meters. A link segment, which is defined as a point-to-point connection between two Medium Independent Interface (MII) devices) can be up to 100 meters. Figure 7-10 illustrates these configuration guidelines.

Figure 7-10 The 100BaseT4 supports a maximum link distance of 100 meters.

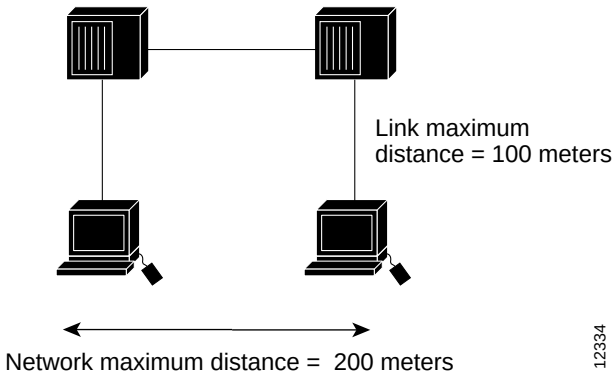

100VG-AnyLAN

100VG-AnyLAN was developed by Hewlett-Packard (HP) as an alternative to CSMA/CD for newer time-sensitive applications, such as multimedia. The access method is based on station demand and was designed as an upgrade path from Ethernet and 16-Mbps Token Ring. 100VG-AnyLAN supports the following cable types:

- 4-pair Category 3 unshielded twisted-pair (UTP)
- 2-pair Category 4 or 5 UTP
- Shielded twisted-pair (STP)
- Fiber-optic

The IEEE 802.12 100VG-AnyLAN standard specifies the link-distance limitations, hub-configuration limitations, and maximum network-distance limitations. Link distances from node to hub are 100 meters (Category 3 UTP) or 150 meters (Category 5 UTP). Figure 7-11 illustrates the 100VG-AnyLAN link distance limitations.

Figure 7-11 100VG-AnyLAN link-distance limitations differ for category 3 and 5 UTP links.

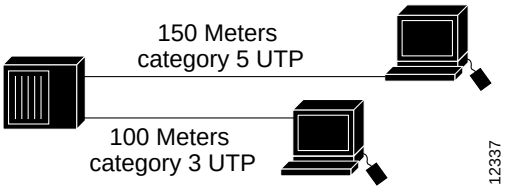

100VG-Any LAN hubs are arranged in a hierarchical fashion. Each hub has at least one uplink port, and every other port can be a downlink port. Hubs can be cascaded three-deep if uplinked to other hubs, and cascaded hubs can be 100 meters apart (Category 3 UTP) or 150 meters apart (Category 5 UTP). Figure 7-12 shows the 100VG-AnyLAN hub configuration.

Figure 7-12 100VG-AnyLAN hubs are arranged hierarchically.

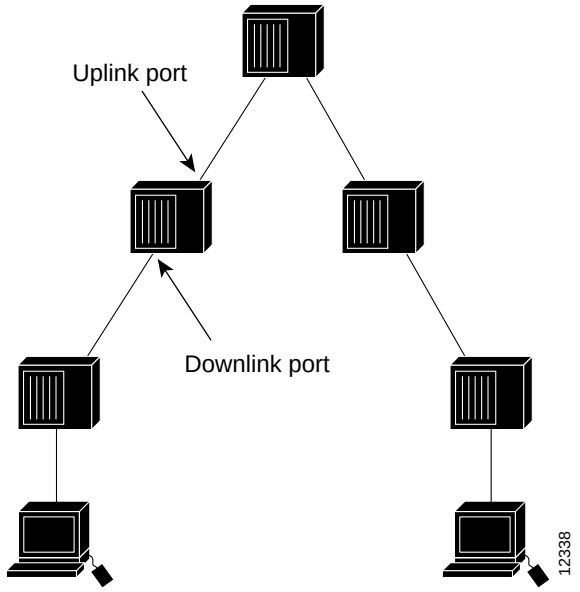

End-to-end network-distance limitations are 600 meters (Category 3 UTP) or 900 meters (Category 5 UTP). If hubs are located in the same wiring closet, end-to-end distances shrink to 200 meters (Category 3 UTP) and 300 meters (Category 5 UTP). Figure 7-13 shows the 100VG-AnyLAN maximum network-distance limitations.

Figure 7-13 End-to-end distance limitations differ for 100VG-AnyLAN implementations.

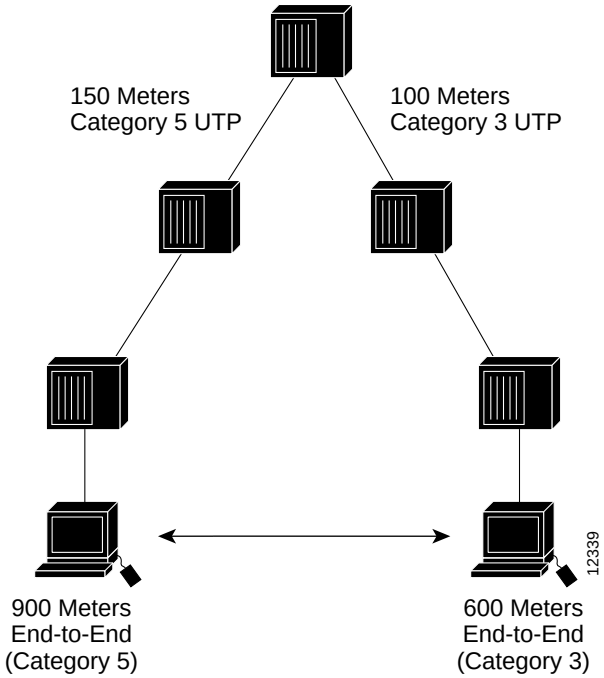

100VG-AnyLAN Operation

100VG-AnyLAN uses a demand-priority access method that eliminates collisions and can be more heavily loaded than 100BaseT. The demand-priority access method is more deterministic than CSMA/CD because the hub controls access to the network.

The 100VG-AnyLAN standard calls for a level-one hub, or repeater, that acts as the root. This root repeater controls the operation of the priority domain. Hubs can be cascaded three-deep in a star topology. Interconnected hubs act as a single large repeater, with the root repeater polling each port in port order.

In general, under 100VG-AnyLAN demand-priority operation, a node wanting to transmit, signals its request to the hub (or switch). If the network is idle, the hub immediately acknowledges the request and the node begins transmitting a packet to the hub. If more than one request is received at the same time, the hub uses a round-robin technique to acknowledge each request in turn. High-priority requests, such as time-sensitive video-conferencing applications, are serviced ahead of normal-priority requests. To ensure fairness to all stations, a hub will not grant priority access to a port more than twice in a row.

Gigabit Ethernet

Gigabit Ethernet is an extension of the IEEE 802.3 Ethernet standard. Gigabit Ethernet offers 1000 Mbps of raw-data bandwidth while maintaining compatibility with Ethernet and Fast Ethernet network devices. Gigabit Ethernet provides for new, full-duplex operating modes for switch-to-switch and switch-to-end-station connections. It also permits half-duplex operating modes for shared connections by using repeaters and CSMA/CD. Furthermore, Gigabit Ethernet uses the same frame format, frame size, and management objects used in existing IEEE 802.3 networks. In general, Gigabit Ethernet is expected to initially operate over fiber-optic cabling but will be implemented over Category 5 unshielded twisted-pair (UTP) and coaxial cabling as well.

The Gigabit Ethernet Alliance is a multi-vendor open forum that promotes industry cooperation in the development of Gigabit Ethernet. The Alliance supports the Gigabit Ethernet standards activities being conducted in the IEEE 802.3 working group and also contributes technical resources to facilitate convergence and consensus on technical specifications. In addition, it provides resources to establish and demonstrate product interoperability, as well as to foster two-way communication between potential suppliers and consumers of Gigabit Ethernet products.

The IEEE 802.3 Working Group has formed the 802.3z Gigabit Ethernet Task Force, which will develop a Gigabit Ethernet Standard that adheres to a number of requirements. The standard must permit full- and half-duplex operation at 1000 Mbps. Conforming implementations will use the standard IEEE 802.3/Ethernet frame format, as well as the CSMA/CD media access method. Gigabit Ethernet implementations also will be backwards-compatible with 10BaseT and 100BaseT. Furthermore, the IEEE standard will specify support of a multi-mode fiber-optic link with a maximum length of 500 meters; a single-mode fiber-optic link with a maximum length of 2 km; and a copper-based link with a maximum length of at least 25 meters. The Gigabit Ethernet standard will be a supplement to the existing 802.3 Ethernet/Fast Ethernet standards.

Gigabit Ethernet Specification

Current standards efforts are based on Fibre Channel and other high-speed networking components. Initial Gigabit Ethernet implementations will use high-speed, 780 nm (short wavelength) Fibre Channel optical components for signaling over optical fiber. 8B/10B encoding and decoding schemes will be used for serialization and deserialization. Fibre Channel technology currently runs at 1.063 Gbps but is being enhanced to run at 1.250 Gbps to provide a full 1000-Mbps data rate. For longer link distances, 1300 nm (long wavelength) optical components will be specified. To

accommodate future advances in silicon technology and digital-signal processing, a logical media-independent interface will be specified between the MAC and PHY layers that will enable Gigabit Ethernet to operate over unshielded twisted-pair (UTP) cabling. This logical interface will enable encoding schemes more suited for use on UTP cabling to be implemented independent of Fibre Channel encoding. Figure 7-14 illustrates the functional elements of Gigabit Ethernet.

Figure 7-14 Functional elements of Gigabit Ethernet.

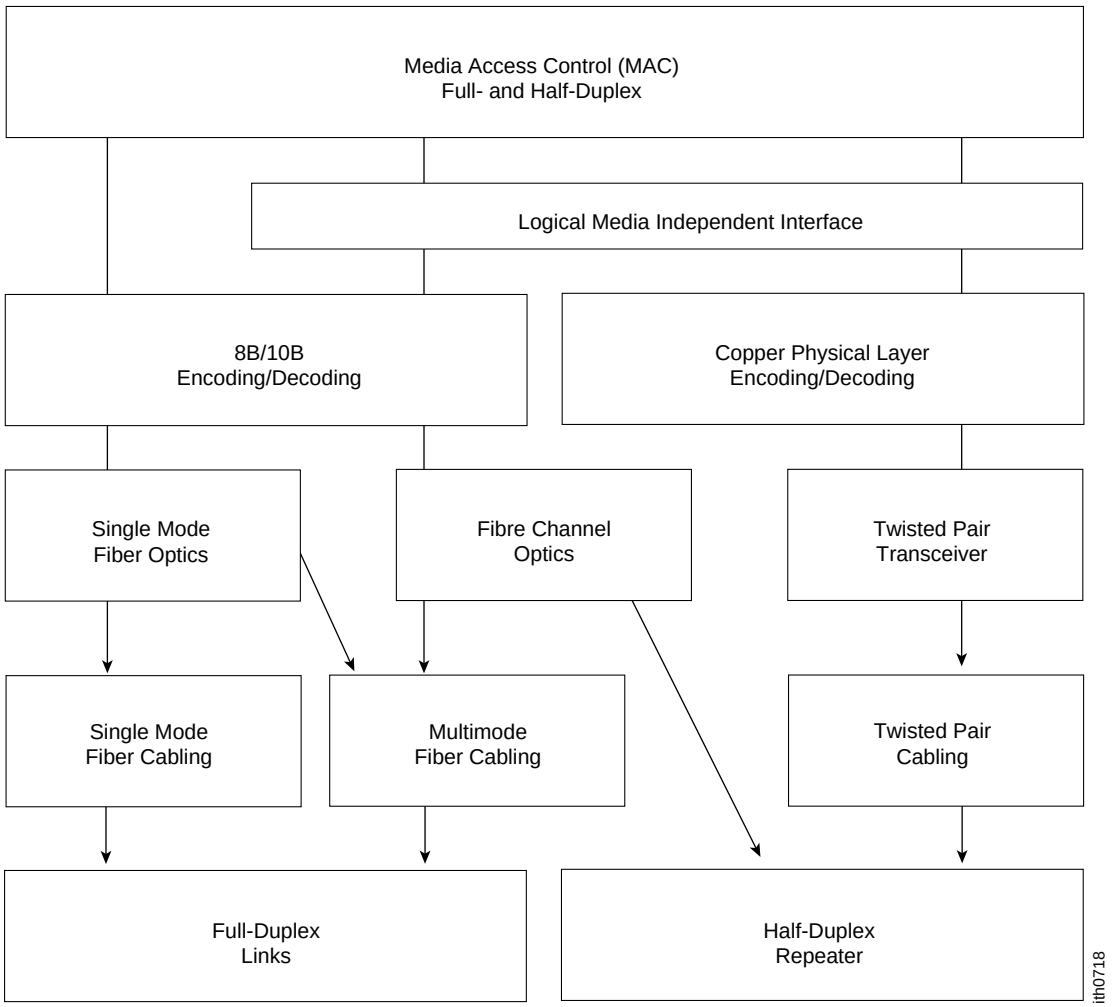

Migrating to Gigabit Ethernet

Migration to Gigabit Ethernet will occur gradually, and initial implementation will be in the backbone of existing Ethernet LANs. Next, server connections will be upgraded, and eventually upgrades will reach the desktop as well. Some of the likely implementation actions include:

- Upgrading switch-to-switch links — 100-Mbps links between Fast Ethernet switches or repeaters can be replaced with 1000-Mbps links, speeding communication between backbone switches and allowing the switches to support a greater number of switched and shared Fast Ethernet segments.

- Upgrading switch-to-server links — 1000-Mbps connections can be implemented between switches and high-performance servers. This upgrade would require the servers to be outfitted with Gigabit Ethernet NICs.
- Upgrading a Fast Ethernet backbone — A Fast Ethernet backbone switch with attached 10/100 switches can be upgraded to a Gigabit Ethernet switch supporting multiple 100/1000 switches, as well as routers and hubs with Gigabit Ethernet interfaces and Gigabit repeaters.

  This action would allow servers to be connected directly to the backbone with Gigabit Ethernet NICs, increasing throughput to the severs for users with high-bandwidth applications. A Gigabit Ethernet network could support a greater number of segments, more bandwidth per segment, and hence a greater number of nodes per segment.
- Upgrading a shared FDDI backbone — An FDDI backbone can be upgraded by replacing the FDDI concentrator, hub, or Ethernet-to-FDDI router with a Gigabit Ethernet switch or repeater. The only upgrade required is the installation of new Gigabit Ethernet interfaces in the routers, switches, or repeaters.
- Upgrading high-performance desktops — Gigabit Ethernet NICs can be used to upgrade high-performance desktop computers to Gigabit Ethernet. These desktop computers would be connected to Gigabit Ethernet switches or repeaters.