Fast Ethernet 100-Mbps Solutions

Introduction

This technology brief describes Fast Ethernet technology, discusses design and migration options for Fast Ethernet deployment, and presents Cisco Systems' Fast Ethernet products.

The growth of local-area networks (LANs) has been driven by the introduction of Ethernet technology as well as the availability of powerful, affordable personal computers and workstations. As a result, applications that once were possible only on mainframe computers are now running on LANs. Network speed and availability are critical requirements. However, existing applications and a new generation of multimedia, groupware, imaging, and database products can easily overwhelm a network running at Ethernet's traditional speed of 10 megabits per second (Mbps). With more applications requiring faster LAN speeds for acceptable performance, network managers face several choices for implementing high-speed LAN technology. For organizations with existing Ethernet installations, increasing the network speed to 100 Mbps is preferable to investing in a completely new LAN technology. This user preference drove the industry's decision to specify a higher-speed Ethernet that operates at 100 Mbps and is known as Fast Ethernet.

Fast Ethernet Development

In July 1993, a group of networking companies joined to form the Fast Ethernet Alliance. The charter of the group was to draft the 802.3u 100BaseT specification of the Institute of Electrical and Electronics Engineers (IEEE) and accelerate market acceptance of Fast Ethernet technology. The final IEEE 802.3u specification was approved in June 1995. Among the other goals of the Fast Ethernet Alliance are: to maintain the Ethernet transmission protocol carrier sense multiple access collision detect (CSMA/CD); to support popular cabling schemes; and to ensure that Fast Ethernet technology will not require changes to the upper-layer protocols and software that run on LAN workstations. For example, no changes are necessary to Simple Network Management Protocol (SNMP) management software or Management Information Bases (MIBs) in order to implement Fast Ethernet.

Cisco made substantial contributions to the development of the basic and optional features of the Fast Ethernet specification through voting representatives in the IEEE 802 Committee and through the Fast Ethernet Alliance membership of Kalpana, a company acquired by Cisco in December 1994. For example, at the 100BaseTX physical layer, Cisco contributed the Multi-Level Transmit (MLT-3) line encoding technology that allows 100-Mbps transmissions such as Fast Ethernet and Fiber Distributed Data Interface (FDDI) to run over Category 5 unshielded twisted-pair (UTP) wire. In addition, Cisco contributed to the specification for the Media Independent Interface (MII), which supports external transceivers at the physical layer and is equivalent to the attachment unit interface (AUI) in 10BaseT. Cisco also contributed a specification for full-duplex operation, first to the 10-Mbps Ethernet standard, and then proposed its extension to the Fast Ethernet standard.

Fast Ethernet Technology

Fast Ethernet, or 100BaseT, is conventional Ethernet but faster, operating at 100 Mbps instead of 10 Mbps. Fast Ethernet is based on the proven CSMA/CD Media Access Control (MAC) protocol and can use existing 10BaseT cabling (See Appendix for pinout diagram and table). Data can move from 10 Mbps to 100 Mbps without protocol translation or changes to application and networking software.

Data-Link Layer

Fast Ethernet maintains CSMA/CD, the Ethernet transmission protocol. However, Fast Ethernet reduces the duration of time each bit is transmitted by a factor of 10, enabling the packet speed to increase tenfold from 10 Mbps to 100 Mbps. Data can move between Ethernet and Fast Ethernet without requiring protocol translation, because Fast Ethernet also maintains the 10BaseT error control functions as well as the frame format and length.

Other high-speed technologies such as 100VG-AnyLAN, FDDI, and Asynchronous Transfer Mode (ATM) achieve 100 Mbps or higher speeds by implementing different protocols that require protocol translation when moving data to and from 10BaseT. This protocol translation involves changes to the frame that typically mean higher latencies when frames are passed through layer 2 LAN switches.

Physical Layer Media Options

Fast Ethernet can run over the same variety of media as 10BaseT, including UTP, shielded twisted-pair (STP), and fiber. The Fast Ethernet specification defines separate physical sublayers for each media type:

- 100BaseT4 for four pairs of voice- or data-grade Category 3, 4, and 5 UTP wiring
- 100BaseTX for two pairs of data-grade Category 5 UTP and STP wiring
- 100BaseFX for two strands of 62.5/125-micron multimode fiber

In many cases, organizations can upgrade to 100BaseT technology without replacing existing wiring. However, for installations with Category 3 UTP wiring in all or part of their locations, four pairs must be available to implement Fast Ethernet.

The MII layer of 100BaseT couples these physical sublayers to the CSMA/CD MAC layer (see Figure 1). The MII provides a single interface that can support external transceivers for any of the 100BaseT physical sublayers. For the physical connection, the MII is implemented on Fast Ethernet devices such as routers, switches, hubs, and adapters, and on transceiver devices using a 40-pin connector (See Appendix for pinout and connector diagrams). Cisco Systems contributed to the MII specification.

Physical Layer Signaling Schemes

Each physical sublayer uses a signaling scheme that is appropriate to its media type. 100BaseT4 uses three pairs of wire for 100-Mbps transmission and the fourth pair for collision detection. This method lowers the 100BaseT4 signaling to 33 Mbps per pair, making it suitable for Category 3, 4, and 5 wiring.

100BaseTX uses one pair of wires for transmission (125-MHz frequency operating at 80-percent efficiency to allow for 4B5B encoding) and the other pair for collision detection and receive. 100BaseFX uses one fiber for transmission and the other fiber for collision detection and receive. The 100BaseTX and 100BaseFX physical signaling channels are based on FDDI physical layers developed and approved by the American National Standards Institute (ANSI) X3T9.5 committee. 100BaseTX uses the MLT-3 line encoding signaling scheme, which Cisco developed and contributed to the ANSI committee as the specification for FDDI over Category 5 UTP. Today MLT-3 also is used as the signaling scheme for ATM over Category 5 UTP.

Figure 1 MII Location in the MAC Layer

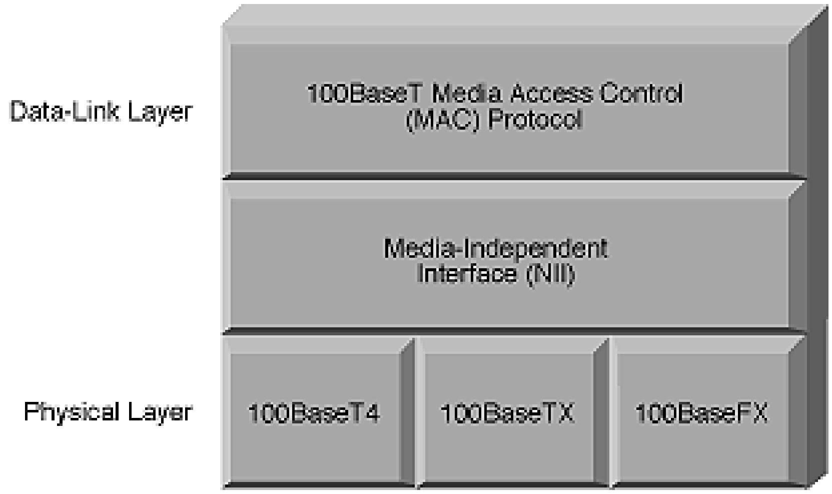

Cabling Distances

While the 100BaseTX and 100Base T4 specifications maintain the same 100-meter limit from the wiring closet to the desktop as 10BaseT, 100BaseFX can exceed the 100-meter limit because it uses fiber instead of UTP. However, 100BaseFX is used primarily between wiring closets and campus buildings to better leverage its support for longer cables. Figure 2 shows a typical wiring topology for a single building.

Just as with 10-Mbps Ethernet, different wiring types can be connected through a repeater. The 100BaseT standard defines two classes of repeaters: Class I and Class II. At most, a collision domain can include one Class I or two Class II repeaters. Fast Ethernet is implemented in a star topology, but even with repeaters, the network diameter is proportionately smaller than 10-Mbps Ethernet, given Fast Ethernet's tenfold increase in packet speed. For example, using two Class II repeaters, the maximum distance using copper wire is 100 meters to the Class II repeater, 5 meters between Class II repeaters, and 100 meters to the desktop.

Figure 2 Fast Ethernet Topology Guidelines

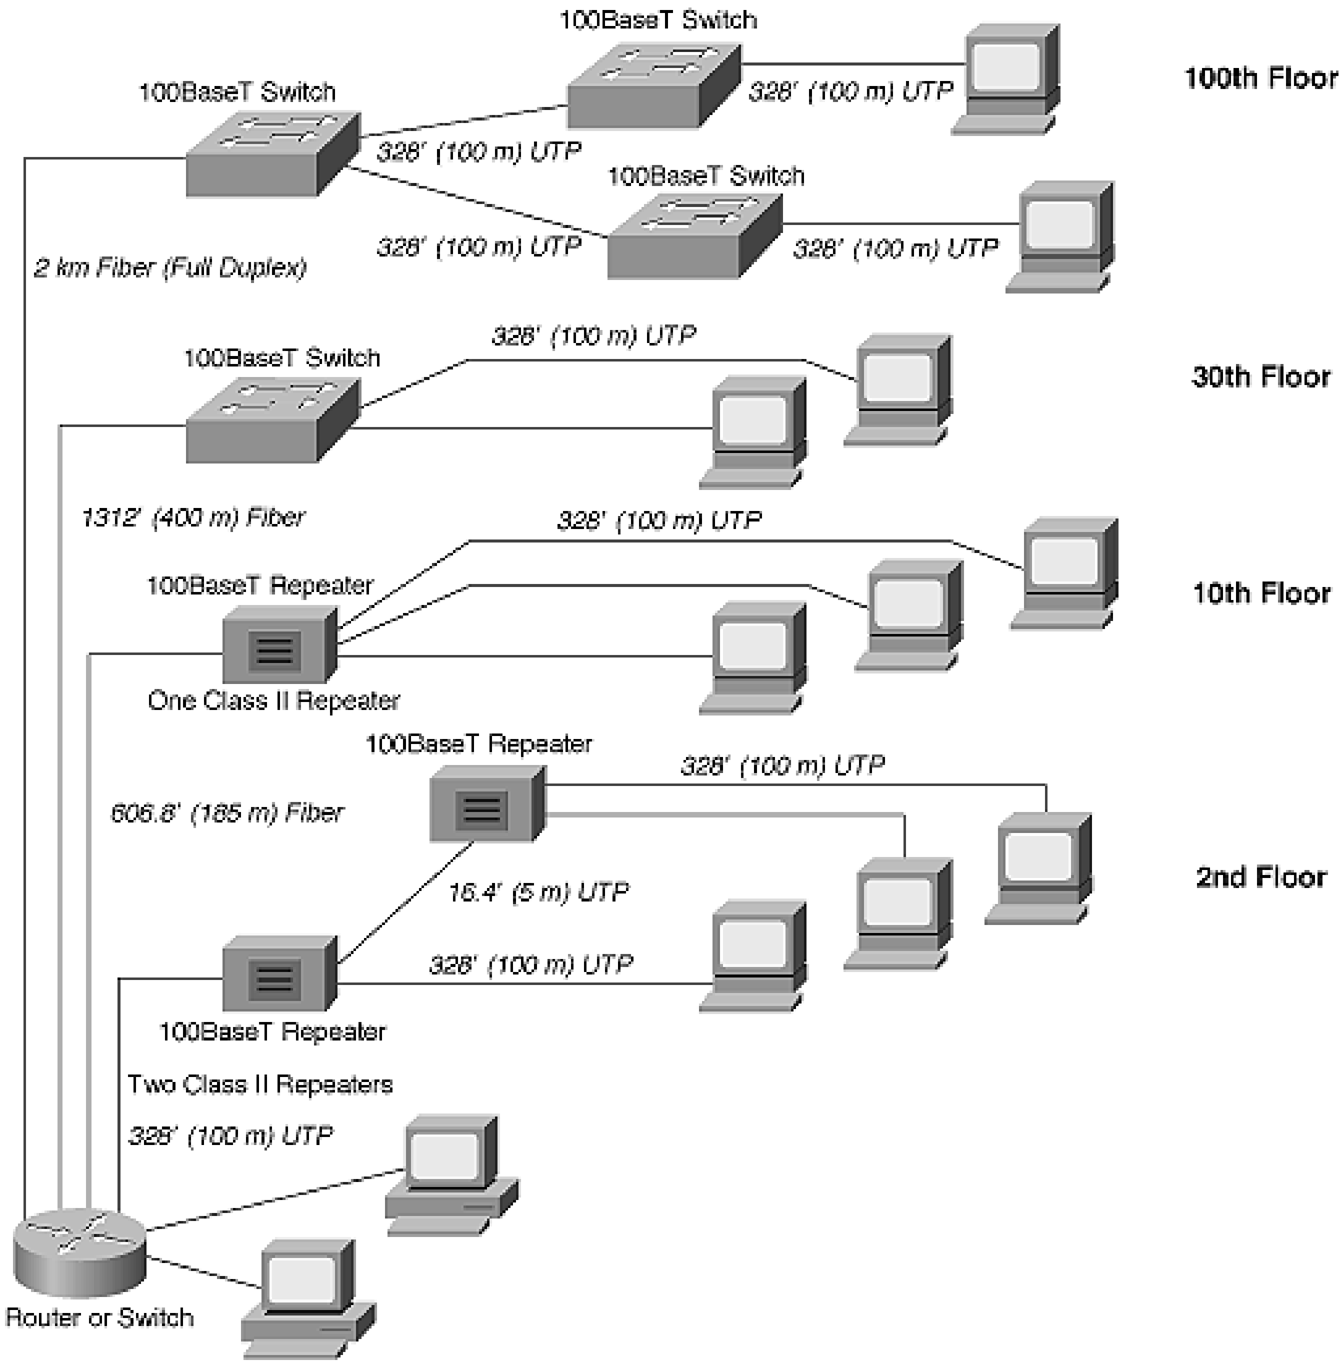

Table 1 lists the maximum network diameter, given different mixes of media types and repeaters. Switches and routers can be used to extend the network diameter between wiring closets.

Table 1 Maximum Distance Parameters

Model	Copper	Copper and Fiber[1]	Fiber
DTE-DTE (no Repeater)	328' (100)	—	1312' (400)
One Class I Repeater	656' (100)	754.4' (230)	787.2' (240)
One Class II Repeater	656' (200)	934.8' (285)	1043' (318)
Two Class II Repeaters	672.4' (205); (656' [200] Category 3)	695.3' (212)	741.2' (226)

1. Note: Mixed copper and fiber assumes a 328' (100 m) copper link

Fast Ethernet Optional Features

Full Duplex

Full-duplex technology delivers up to 200-Mbps bandwidth because it provides bidirectional communication—meaning that 100 Mbps is available for transmission in each direction. Full duplex also increases the maximum distance supported for fiber cables between two data terminal equipment (DTE) devices up to 2 km (see Table 2).

Table 2 Cabling and Distance Parameters

Physical Sublayer Option		Cable Specification	Length (meters)
100BaseTX		Category 5 UTP, two pairs;	100 half/full duplex
		Types 1 and 2 STP, two pairs	100 half/full duplex
100BaseT4		Categories 3, 4, and 5 UTP four pairs	100 half/full duplex
100BaseFX		62.5/125 multimode fiber/two strands	400 half duplex; 2000 full duplex
20	+5V	40	+5V

Full-duplex communication is implemented by disabling the collision detection and loopback functions, which are necessary to ensure smooth communication in a shared network. Only switches can offer full duplex to directly attached workstations or servers. Shared 100BaseT hubs must operate at half duplex to detect collisions among end stations. Users will see greater performance improvements when full-duplex 100BaseT is implemented on a backbone connection rather than on a client connection, because client/server applications primarily transmit read/write asymmetrical traffic.

Cisco's Workgroup Business Unit pioneered the development of full-duplex features for 10BaseT switching technology and has brought them forward for 100BaseT switching. Full-duplex 10BaseT technology is readily available and deployed by customers today. Adapter, hub, and switch vendors including Cisco and others have tested and verified the interoperability of their full-duplex 10BaseT products.

Autonegotiation

The 100BaseT specification describes a negotiation process that allows devices at each end of a network link to automatically exchange information about their capabilities and perform the configuration necessary to operate together at their maximum common level. For example, autonegotiation can determine whether a 100-Mbps hub is connected to a 10-Mbps or 100-Mbps adapter and then adjust its mode of operation accordingly.

This autonegotiation activity is performed out-of-bank using Fast Link Pulse (FLP) burst to identify the highest physical-layer technology that can be used by both devices, such as 10BaseT, 100BaseT, 100BaseTX, or 100BaseT4. The autonegotiation definition also provides a parallel detection function that allows half- and full-duplex 1BaseT, half- and full-duplex 100BaseTX, and 100BaseT4 physical layers to be recognized, even if one of the connected devices does not offer autonegotiation capabilities.

Flow control can be implemented on a link-by-link or end-to-end basis and allows all devices on the path to reduce the amount of data they receive. Because flow control has implications beyond full duplex and the MAC sublayer of the data link layer, methods and standards are still under consideration by the IEEE 802.3x committee.

Link-to-link flow control examines only an individual link between switches or stations. When the recipient of a transmission becomes busy, it will send a signal to the directly linked transmitter. If that transmitter is not the originator of the traffic, this signal would have to be propagated back through each link in order to reach the originator of the traffic. A more complete solution would be to identify the specific traffic causing congestion, but this may require upper-level protocols.

End-to-end flow control means that the switches at each end of the link communicate to throttle back the end stations that are originating the traffic. Until this information is propagated, packets must be stored or dropped, meaning that flow control reduces but does not eliminate the need for buffers.

As is the case for all high-speed networking technologies, managers should choose products that have sufficient buffer space to keep up with the speed of their networks. Large buffers ensure that data received from a network can be stored until a workstation can accept it. Methods and standards for flow control are still under consideration by the IEEE 802.3u committee.

Alternative High-Speed Technologies

Network managers who plan to migrate to high-speed technologies can select from several different options, such as Copper Distributed Data Interface (CDDI)/FDDI, ATM, and 100GV-AnyLAN. Most networks will deploy a combination of these technologies, which are compared in Table 3.

Table 3 Comparisons of High-Speed LAN Technologies

Feature	100BaseT Fast Ethernet	100VG-AnyLAN	CDDI/FDDI	ATM	LAN Switching
Data Rate	100 Mbps	100 Mbps	100 Mbps	25 to 622 Mbps	10 or 4/16 Mbps
Access Method	CSMA/CD	Demand priority	Token passing	Cell based	LAN-based switching
Frame Size	64 to 1500 bytes	64 to 16 KB	64 to 4500 bytes	53 bytes	64 to 8 KB
Services	Asynchronous	Asynchronous and synchronous	Asynchronous and synchronous	Isochronous, asynchronous, and synchronous	Asynchronous
Network Diameter	672.4' (205)	984' (300)	328' (100) to 18.6 mi. (30 km)	328' (100)to multiple miles (km)	—
Management	SNMP and Ethernet MIBs	SNMP MIB	SMT, SNMP	Proprietary MIBs and SNMP	SNMP and Ethernet MIBs
Cost	Low cost	Low cost	Declining cost	Higher cost	Low cost
Fault Tolerance	Spanning tree		Dual-homing MAC ring	Multiple paths	Spanning tree
Application	Desktop, workgroup, and backbone	Desktop, backbone, and multimedia	Desktop, workgroup, and backbone	Backbone, WAN, LAN, multimedia, and desktop	Desktop, workgroup, and backbone

CDDI/FDDI is a 100-Mbps technology that uses token passing over copper wiring (CDDI) or fiber (FDDI). 100VG-AnyLAN is a 100-Mbps technology that uses a demand-priority protocol over four pairs of Category 3 UTP cable. ATM is a cell-based technology operating at 155 Mbps or higher that is designed to be a single network for voice, data, and video—namely multimedia traffic.

Designing with 100BaseT

Fast Ethernet technology can be designed to an internetwork in several different ways. The following LAN designs can be used in isolation or combined. The first design is 10-Mbps LAN switching that coexists with hubs and routers. For example, switches can segment hubs to increase performance on shared-media LANs without changing a network's topology. To enhance client/server performance across the enterprise, servers can connect directly to Fast Ethernet interfaces on the LAN switches that aggregate traffic from numerous switched 10-Mbps Ethernet segments. These switched groups connect to the enterprise through routers with Fast Ethernet interfaces.

Network managers should choose a LAN switching platform that offers at least two switched Fast Ethernet interfaces for connecting to a server and router. In addition, they should choose a switching platform that supports both fiber (10Base-FL) for connecting to the wiring closet and long-distance fiber (100BaseFX) connections between closets and campus buildings. (See Figure 3.)

Figure 3 Front-End Hubs with Switching, Fast Ethernet Server, and Backbone Links

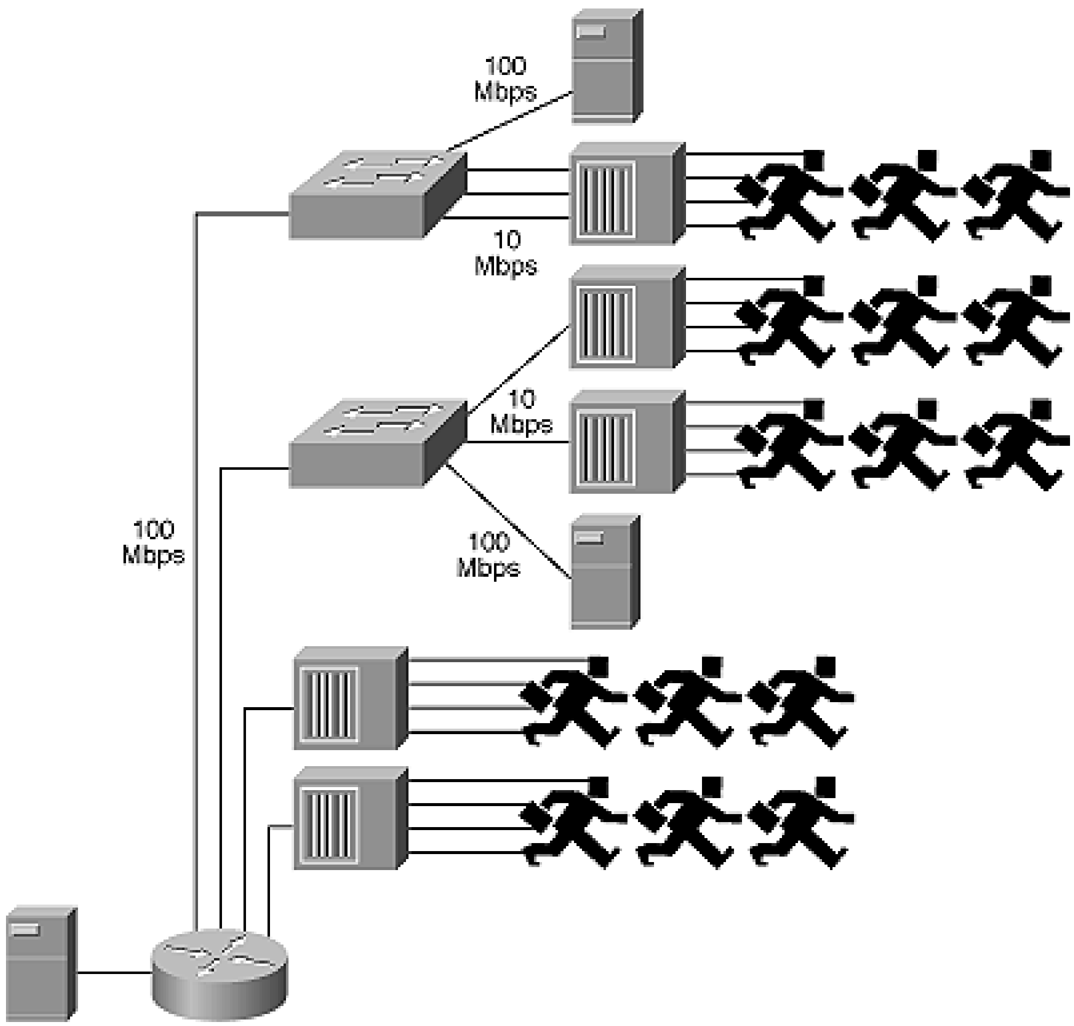

In the second design, 10-Mbps LAN switching is integrated to the desktop, providing dedicated bandwidth and virtual LAN (VLAN) services to the end stations. LAN switches can be linked together (cascaded) using Fast Ethernet interfaces to build high-density networks, as shown in Figure 4.

Figure 4 10-Mbps Switching to the Desktop with Fast Ethernet

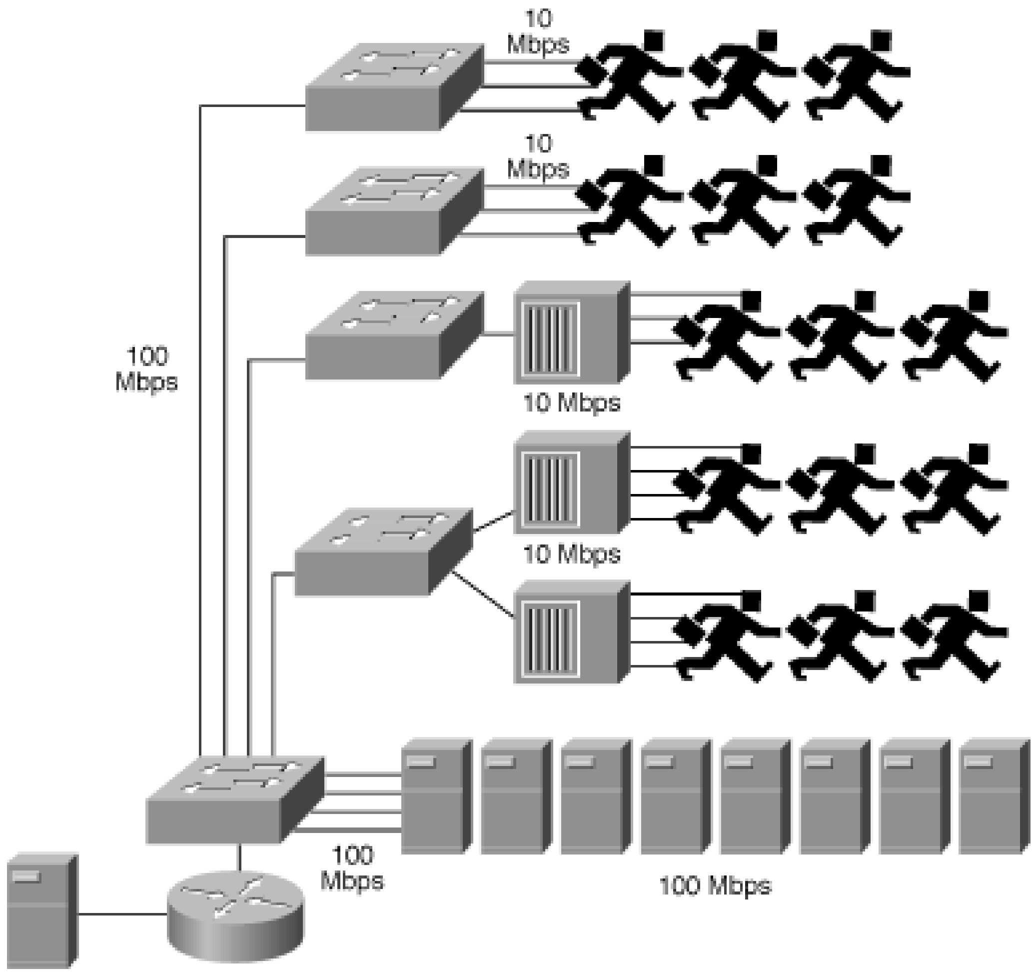

In the data center, the switching platform should offer high-density Fast Ethernet and the ability to extend VLANs over any Fast Ethernet interface for use as a collapsed switch. In addition, high-density Fast Ethernet is necessary for connecting multiple servers in the data center for applications such as "server farms."

VLANs can be extended across the enterprise with network managers choosing switches and routers that support a common set of VLAN protocols. Network managers can further enhance performance by choosing switches and routers that support full duplex on Ethernet and Fast Ethernet connections.

In the third design, shared Fast Ethernet bandwidth is integrated to the desktop, allowing user applications to obtain burst data speeds of 100 Mbps. When only a few users require this speed, the Fast Ethernet bandwidth can be shared with minimal contention. Figure 5 illustrates this example.

Figure 5 Front-End Fast Ethernet Hubs with Switching

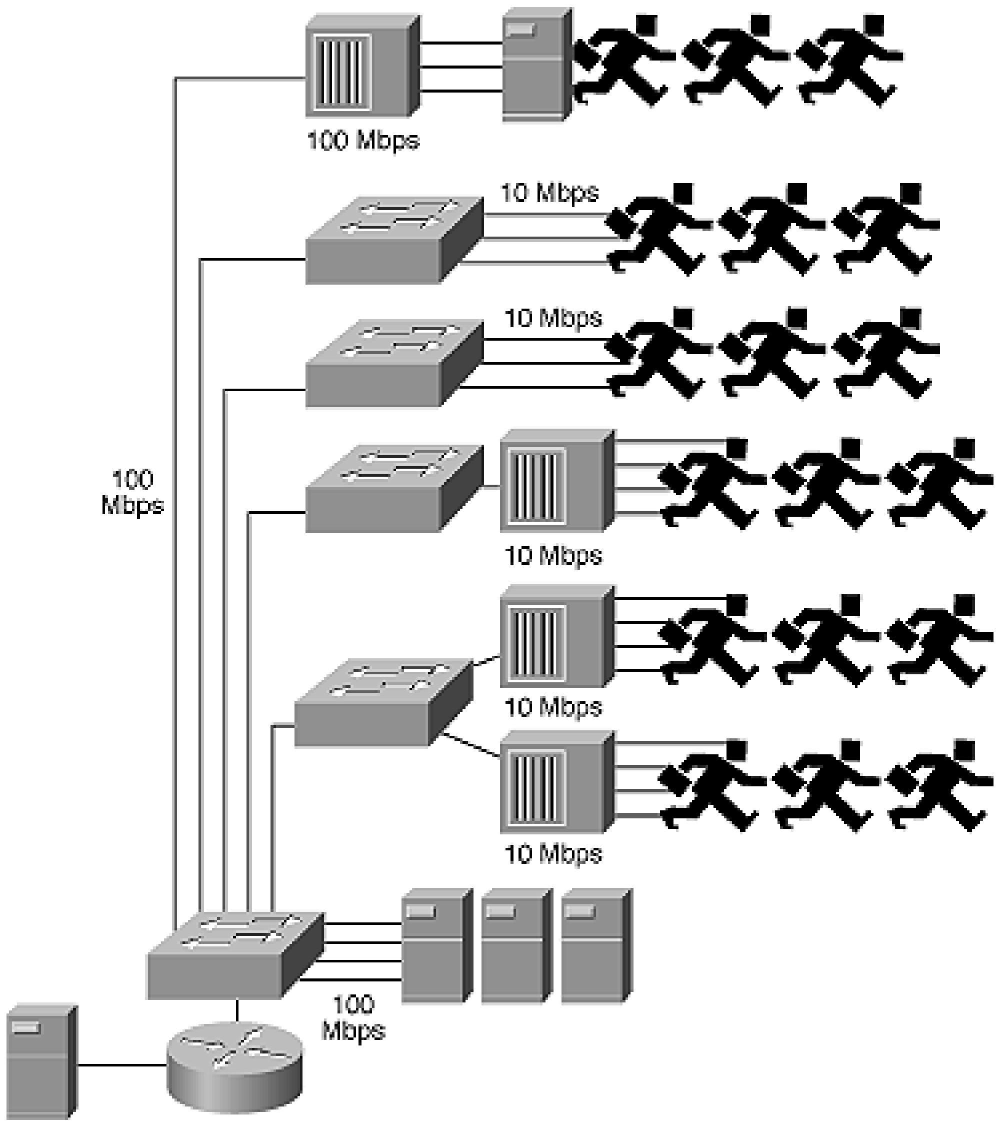

In the fourth design, Fast Ethernet switching is integrated to the desktop, giving high-performance PC and workstation users 100-Mbps access to the server. Several adapter vendors are currently delivering Fast Ethernet products at less than two times the cost of 10-Mbps Ethernet adapters. This design is shown in Figure 6. The LAN switching platform should be capable of adding one or more ATM uplinks to an interconnected core of ATM switches.

Cisco Systems Products

Cisco Systems offers a complete family of routing and switching products for constructing cost-effective Fast Ethernet enterprise networks that can coexist with established 10BaseT networks. Cisco provides critical components for a phased migration to Fast Ethernet through its interoperable platforms, software, and management tools.

Tying together these products is the Cisco Internetwork Operating System (Cisco IOS®) software, Cisco's underlying networking software. Cisco IOS software enables a variety of routing and switching functions across internetworks, including those with a mix of protocols and LAN types.

Cisco products are tested at the University of New Hampshire, U.S.A., for conformance with 10BaseT, 10BaseFL, and Fast Ethernet standards. The participation of Cisco Systems and other vendors at this lab fosters interoperability among the products tested. Cisco's quality commitment extends to its own lab, which tests the compatibility of adapters from other vendors with Cisco's LAN switches. In the future, Cisco will continue to enhance its Fast Ethernet technology and provide leadership for the definition of Fast Ethernet features.

Figure 6 Fast Ethernet to the Desktop

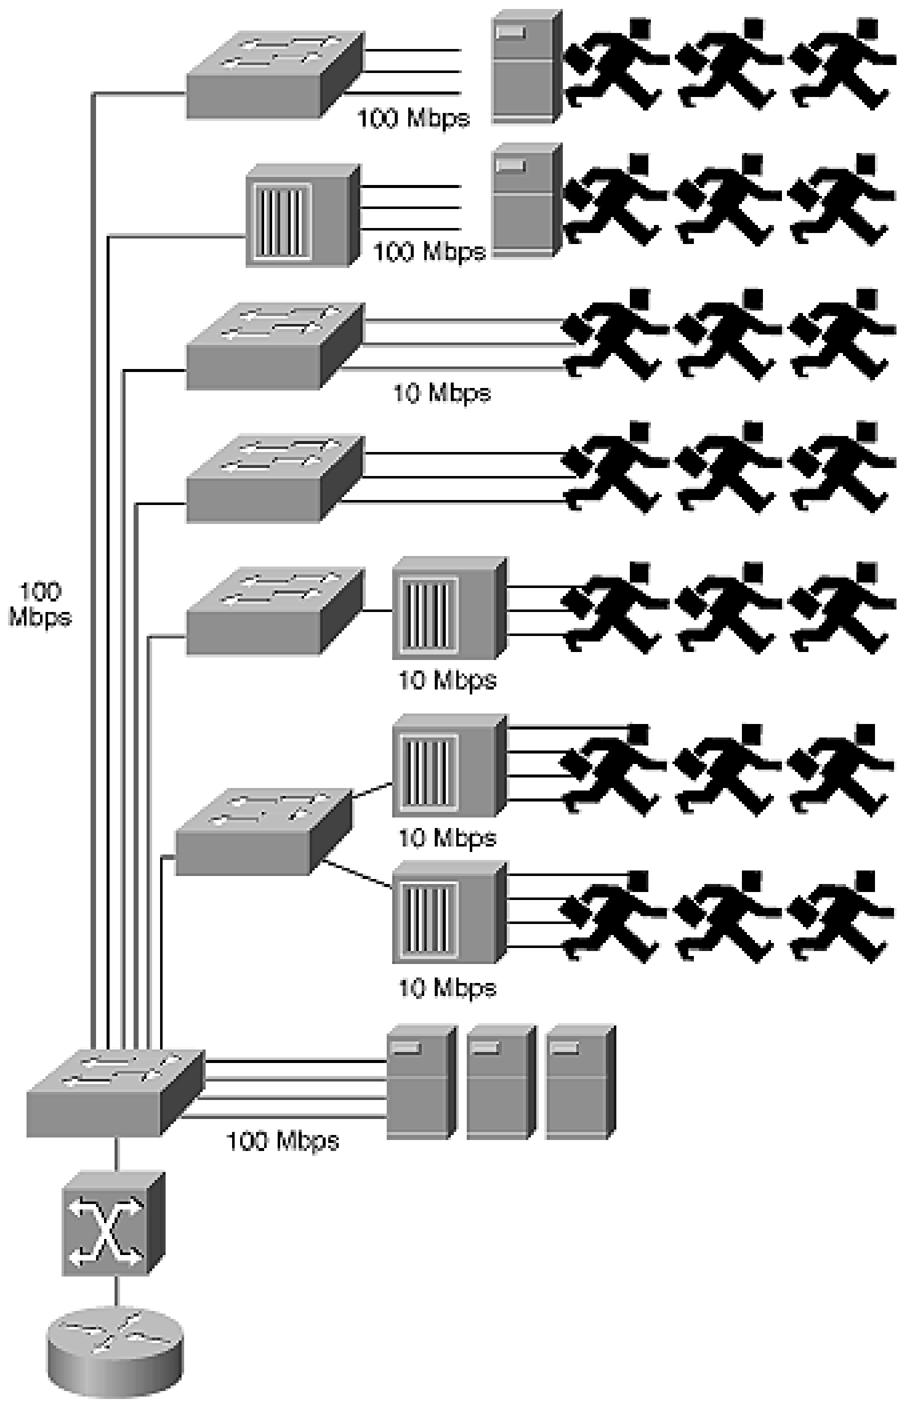

Routers

A Fast Ethernet Interface Processor (FEIP) card for the Cisco 7000 family of high-end routers offers one or two 100BaseTX ports that support full- and half-duplex operation. In conjunction with the Cisco Catalyst® 5000 switch, the FEIP can aggregate a large number of switched 10-Mbps LANs and give them high-speed access to routing services. Each FEIP 100BaseT interface provides an RJ-45 port for connection to 100BaseTX and an alternate DB-40 MII port that, with external transceivers, can connect to 100BaseFX and 100BaseT4 physical sublayers. All FEIP interfaces will support virtual LANs (VLANs), enabling network managers to use switches and routers to logically organize enterprise-wide workgroups.

LAN Switches

The Catalyst 5000 family of LAN switches is composed of five chassis: the Catalyst 5500, 5509, 5505, 5000, and 5002. These modular chassis range in size from the 13-slot Catalyst 5500, 9-slot Catalyst 5509, 5-slot Catalyst 5505 and 5000, and the 2-slot Catalyst 5002. All five chassis share the same set of interface modules and software features, providing scalability while maintaining interoperability and investment protection across all chassis. The switches employ a flexible switching architecture based upon a series of Supervisor Engine modules—Supervisors Engines I, II, and III—that provide switching, network management, and uplink ports. The media-independent architecture of the Catalyst 5000 family support all LAN and ATM switching technologies through a wide array of 10/100/1000 Ethernet, FDDI, Token Ring, and ATM switch modules shown in Table 4.

Table 4 Catalyst 5000 Family Interface Support

Catalyst 5000 Family Switching Modules	Number of Interfaces Supported per Slot	Maximum Interfaces per Chassis
Group Switched 10BaseT Ethernet	48 interfaces	528
Switched 10BaseT Ethernet (RJ-21)	48 interfaces	528
Switched 10BaseT Ethernet (RJ-45)	24 interfaces	264
Switched 10BaseFL Ethernet	12 interfaces	132
Group Switched 100BaseTX Ethernet	24 interfaces	264
Switched 10/100BaseTX Fast Ethernet	24 interfaces	264
Switched 100BaseFX Fast Ethernet	12 interfaces	134
ATM LAN Emulation (OC-3 or OC-12)	1 (dual PHY)	7 (dual PHY)
CDDI / FDDI	1 interface	11
Gigabit Ethernet	9/3/2 interfaces	38

Management Tools

Cisco offers the following tools for managing its routers and switches:

- The CiscoWorks suite of network management applications that offers broad functionality for device, fault, and performance management of routers and switches
- The CiscoView application that operates as part of the CiscoWorks suite and presents a graphical display of the router or switch chassis and components; the CiscoView application supports configuration, performance monitoring, and basic troubleshooting
- An intuitive, drag-and-drop application that enables managers to configure their networks based on logical user groups and display an enterprise-wide logical view of these groups
- A traffic monitoring and analysis application that leverages embedded RMON agents in the Cisco LAN switches

Conclusion

Fast Ethernet technology is now available from major networking product vendors. The 100BaseT standard leverages existing 10-Mbps Ethernet and FDDI technologies and, in many cases, existing cabling. Fast Ethernet can be deployed for high-speed links to routers, to build high-performance backbone switches and centralized server farms, to operate as local workgroup servers, and to provide both shared and switched bandwidth to the desktop.

APPENDIX: Interfaces and Connectors

Table A-1 sets forth MII connector Contact Assignments.

Table A-1 MII Connector Contact Assignments

Pin	Signal Name	Pin	Signal Name
1	+5V	21	+5V
2	MDIO	22	Common
3	MDC	23	Common
4	RxD<3>	24	Common
5	RxD<2>	25	Common
6	RxD<1>	26	Common
7	RxD<0>	27	Common
8	RxDV	28	Common
9	Rx_Clk	29	Common
10	Rx_Er	30	Common
11	Tx_Er	31	Common
12	TxClk	32	Common
13	Tx_En	33	Common
14	TxD<0>	34	Common
15	TxD<1>	35	Common
16	TxD<2>	36	Common
17	TxD<3>	37	Common
18	Col	38	Common
19	CRS	39	Common
20	+5V	40	+5V

Figure A-1 shows the RJ-45 connection pinout for 10BaseT and 100 Mbps. Figure A-1 shows the MII connector.

Figure A-1 RJ-45 Connection Pinout for 10BaseT and 100 Mbps Connector

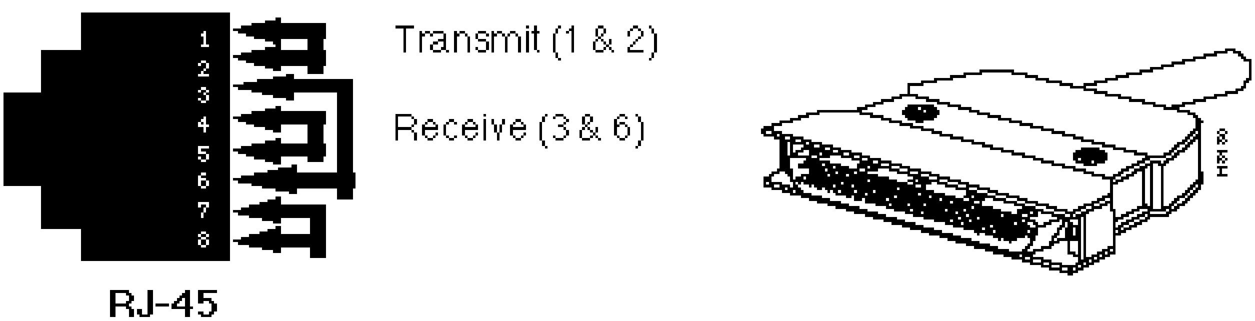

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
http://www-europe.cisco.com
Tel: 33 1 69 18 61 00
Fax: 33 1 69 28 83 26

Americas
Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
http://www.cisco.com
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

Cisco Connection Online Web site at http://www.cisco.com/offices.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic • Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore
Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela