Gigabit Ethernet Comes of Age

Scaling Network Performance to 1000 Mbps

Gigabit Ethernet Comes of Age
Scaling Network Performance to 1000 Mbps

Contents

Gigabit Ethernet Comes of Age

Scaling Network Performance to 1000 Mbps

Is there a cost-effective way to accommodate exploding application and user bandwidth demands on network backbones while limiting network complexity? This white paper presents 3Com's strategy for IT managers faced with bandwidth capacity constraints and considering migrating to a Gigabit Ethernet (GbE) technology solution. The paper briefly examines the forces driving the adoption of Gigabit Ethernet, the concepts behind the technology, and the key advantages of implementing Gigabit Ethernet over other technologies. The paper also explains the recently ratified Institute of Electrical and Electronics Engineers (IEEE) 802.3z Gigabit Ethernet standard and accompanying cabling specifications and their significance in network implementation. Finally, the paper covers some important deployment issues and presents an incremental migration path to a Gigabit Ethernet infrastructure.

Factors Driving Network Bandwidth

Progress is relentless; new technologies spur new application development, and new applications, in turn, fuel the need for further technology advancements. In a relatively short time span, network connections have evolved from shared or switched Ethernet to shared or switched Fast Ethernet to accommodate rising bandwidth demand. Within the enterprise networks, business applications are now advancing to embrace high-resolution graphics, video, and other rich media types that exceed the capacity of even Fast Ethernet performance. This latest generation of bandwidth-intensive applications can be organized into four categories:

- Scientific modeling, publication, and medical imaging applications produce multimedia and graphics files that are ballooning in size from megabytes to gigabytes to terabytes.
- Internet and intranet applications create unpredictable any-to-any traffic patterns. Increasingly, with servers distributed across the enterprise and users accessing Web sites inside and outside the corporate network, there is no way to predict where traffic will go.
- Data warehousing and backup applications handle gigabytes or terabytes of data distributed among hundreds of servers and storage systems.
- Bandwidth-intensive groupware applications such as desktop video conferencing, interactive whiteboarding, and real-time video that support mission-critical business applications not only require more raw bandwidth, they also demand low latency and limited jitter to be effective.

Table 1 summarizes the applications and their impact on the network. As you can see, these applications are driving the demand for increased bandwidth, applying pressure at the desktop, the server, and the hub, and creating higher capacity demands at the network core.

Another critical factor spurring backbone bandwidth demand is the cost effectiveness of edge and workgroup switching. Switching is fast becoming the predominant form of Ethernet connection; switched Ethernet shipments are predicted to outnumber shared Ethernet shipments by the year 2000 (Dell'Oro Group, Portola Valley, California). Switching at the network edge puts a tremendous load on the network backbone where traffic aggregates.

What Is Gigabit Ethernet?

Gigabit Ethernet is Ethernet that provides speeds of 1000 Mbps—one billion bits per second. It uses the same Ethernet frame format and media access control technology as all other 802.3 Ethernet technologies. It also uses the same 802.3 full-duplex Ethernet technology and 802.3 flow control. In terms of data rate, Gigabit Ethernet simply moves the decimal point, increasing the Fast Ethernet transmission rate tenfold, from 100 Mbps to 1000 Mbps. Table 2 compares Fast Ethernet and Gigabit Ethernet technologies.

Like its Ethernet and Fast Ethernet precursors, Gigabit Ethernet is a physical (PHY) and media access control (MAC) layer technology, specifying the Layer 2 data link layer of the OSI protocol model. It complements upper-layer protocols TCP and IP, which

Acronyms and Abbreviations

ASIC
application-specific integrated circuit

ATM
Asynchronous Transfer Mode

CoS
Class of Service

CPU
central processing unit

CSMA/CD
carrier sense multiple access/collision detection

FDDI
Fiber Distributed Data Interface

GbE
Gigabit Ethernet

GBIC
gigabit interface connector

GMII
Gigabit Media Independent Interface

IEEE
Institute of Electrical and Electronics Engineers

IP
Internet Protocol

MAC
media access control

MAN
metropolitan area network

MM
multimode

MTU
maximum transmission unit

Table 1. *Application Bandwidth Drivers*

	Data Types/Size	Traffic Implication	Network Requirement
Scientific Modeling, Publications, Medical	• Data Files • MB to GB	• Increased bandwidth needed for large files	• Higher bandwidth for desktops, servers, backbone
Internet/Intranet	• Data files now • Audio now • Video future • High transaction rate	• Increased bandwidth needed for large files • Low latency • CoS/QoS	• Higher bandwidth for desktops, servers, backbone • Low latency
Data Warehousing, Backup	• Large files • Data files • GB to TB	• Many data streams • Increased bandwidth needed for large files • Search and access require low latency • Backup performed during fixed period	• Higher bandwidth for servers and backbone • Low latency
DVC, Interactive Whiteboard	• Constant data stream • Up to 3+ Mbps to the desktop	• CoS/QoS • Many data streams	• Higher bandwidth for servers and backbone • Low latency • Predictable latency

Source: Gigabit Ethernet Alliance

specify the Layer 4 transport and Layer 3 network portions and enable reliable communication services between applications.

The network topology for Gigabit Ethernet follows the traditional rules of Ethernet. At Layer 2, the Spanning Tree Protocol is used to ensure that there are no logical loops in the network, creating a hierarchical tree topology. More complex LAN configurations, including parallel data paths, are created through the use of routing technology. At any speed, Ethernet is designed for transporting data in the local area environment. Like Ethernet and Fast Ethernet, Gigabit Ethernet leverages other technologies and standards to provide higher-level services such as Class of Service (CoS) traffic prioritization and Quality of Service (QoS) data delivery that limit jitter and latency.

Why Gigabit Ethernet?

There are three compelling reasons to deploy Gigabit Ethernet for greater bandwidth at the backbone:

- Gigabit Ethernet delivers the scalable performance of Ethernet technology.
- Gigabit Ethernet leverages the installed base of Ethernet equipment.
- Gigabit Ethernet leverages the historical cost drivers of Ethernet switching equipment.

Table 2. *Gigabit Ethernet Is Ethernet*

	Fast Ethernet	Gigabit Ethernet
Speed	100 Mbps	1000 Mbps
Frame Format	802.3 Ethernet	802.3 Ethernet
MAC Layer	802.3 Ethernet	802.3 Ethernet
Flow Control	802.3x Ethernet	802.3x Ethernet
Primary Mode	Full duplex	Full duplex
Signaling	FDDI	Fibre Channel

Acronyms and Abbreviations

NDIS
network driver interface specification

NIC
network interface card

NOS
network operating system

OSI
Open Systems Interconnection

PCI
Peripheral Component Interconnect

PCS
physical coding sublayer

PHY
physical

PMA
physical medium attachment

QoS
Quality of Service

RMON
Remote Monitoring

RSVP
Resource Reservation Protocol

SM
single-mode

TCP
Transmission Control Protocol

UTP
unshielded twisted pair

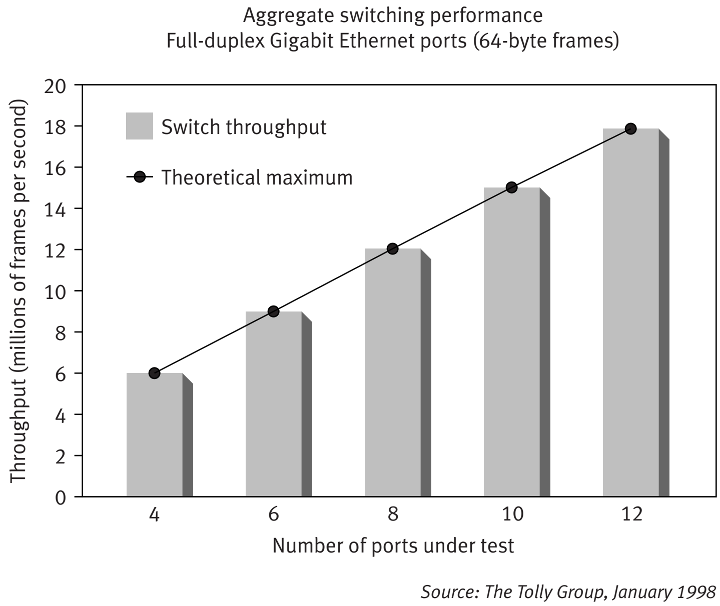

Figure 1. *Gigabit Ethernet Wire-Speed Performance*

Gigabit Ethernet Delivers Cost-Effective 1000 Mbps Ethernet Performance

Ethernet is the dominant and most ubiquitous LAN technology today, and it is clearly the most widely understood LAN technology. According to industry analyst International Data Corporation (IDC, Framingham, MA), more than 85 percent of all installed network connections were Ethernet by the end of 1997, representing more than 118 million interconnected PCs, workstations, and servers. Ethernet is popular because it offers the best combination of price, simplicity, scalability, and management ease of use. Because Gigabit Ethernet is Ethernet, network managers can leverage their investment in staff training and skills, reducing their total cost of ownership.

Gigabit Ethernet extends Ethernet's scalable 10/100 Mbps performance to 1000 Mbps. Figure 1 shows third-party tests results for the 3Com SuperStack® II Switch 9300, a twelve-port Gigabit Ethernet switch. The switch achieved an aggregate throughput of more than 17.8 million frames per second in a mesh configuration using all 12 full-duplex switch ports, or 1.48 million frames per second for each port. Once the bit overhead of the 12-byte interframe gap and the 8-byte preamble are included in the calculation, the test shows that the Gigabit Ethernet link performs at full wire speed of 1.0 billion bits per second with 64-byte frames.

Gigabit Ethernet Leverages Ethernet Equipment Costs

The goal of the IEEE 802.3z Task Force, which developed the Gigabit Ethernet standard, was to specify connections that delivered 10 times the performance of Fast Ethernet at very affordable prices. Because Gigabit Ethernet leverages existing Ethernet technologies, it also leverages Ethernet's fiercely competitive industry cost curve.

Gigabit Ethernet products are a good value today. Currently, switched 1000 Mbps Gigabit Ethernet provides the best price performance of all the leading high-speed LAN technologies. Figure 3 illustrates data from the Dell'Oro Group showing that switched Gigabit Ethernet delivers the lowest cost per data rate—less than $2 U.S. per Mbps—compared to switched FDDI, switched 155 Mbps ATM, switched 622 Mbps ATM, and switched Fast Ethernet.

The Gigabit Ethernet Standard

The 1000BASE-X (IEEE 802.3z) Gigabit Ethernet standard was ratified in June 1998, after more than two years of intense effort within

Leveraging Experience and Expertise: Lowering the Cost of Ownership

The total cost of ownership is an important factor in evaluating any new networking technology. Total cost of ownership includes not only the purchase price of equipment, but also the cost of network administration, application support, training, maintenance, and troubleshooting. Figure 2 shows the results of research by the Gartner Group that quantifies the total annual cost of ownership for a network-attached desktop. Note that 73 percent of the total is "people costs": 33 percent in managing the NOS and supporting the applications that run on the network, 13 percent in administering the network, and 27 percent in "hidden" costs incurred when network problems cause productivity losses.

With familiar technologies like Ethernet, people costs are lower because network managers are able to leverage the skill sets and training of their existing network administration staff and leverage existing network analysis and network management tools. And because increasing the number of protocols and technologies always increases network complexity, deploying an Ethernet-based backbone solution is simpler than deploying FDDI or ATM. A less complex network, in turn, lowers the cost of ownership.

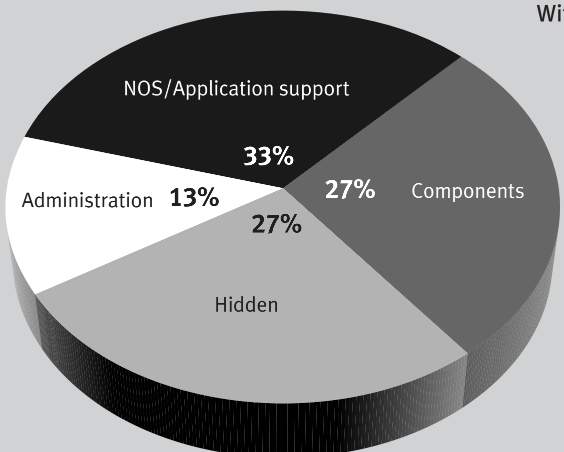

Source: The Gartner Group

Figure 2. *Total Cost of Ownership for a Networked Desktop*

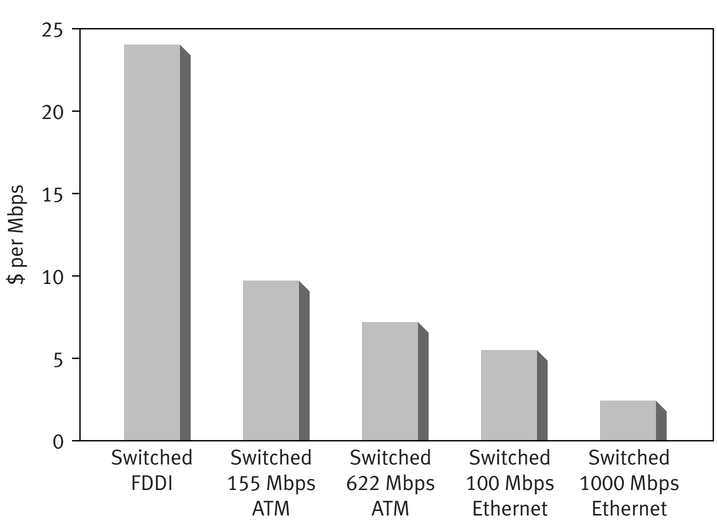

Source: Dell'Oro Group, CY1997

Figure 3. *The Price/Performance of Bandwidth*

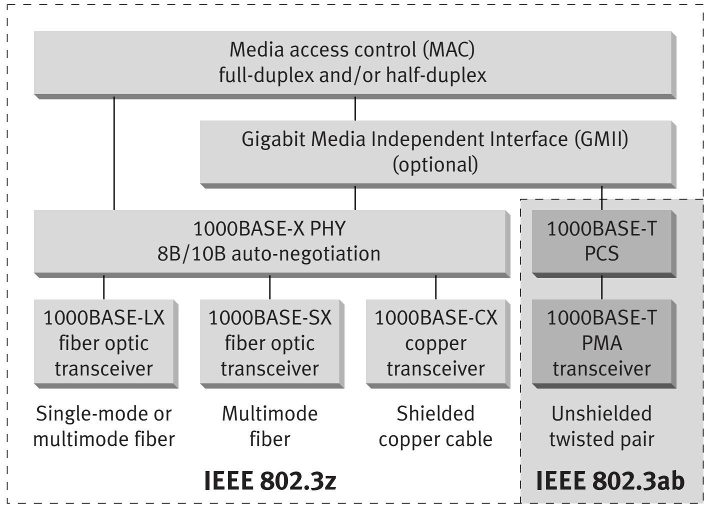

Source: IEEE

Figure 4. *Gigabit Ethernet Architecture Standard*

the IEEE 802.3 Ethernet committee. The key objective of the 802.3z Gigabit Ethernet Task Force was to develop a Gigabit Ethernet standard that encompassed the following:

- Allowed half- and full-duplex operation at speeds of 1000 Mbps
- Used the 802.3 Ethernet frame formats
- Used the CSMA/CD access method with support for one repeater per collision domain
- Addressed backward compatibility with 10BASE-T and 100BASE-T technologies

Because the fundamental features of the 802.3z specification have been stable during the last stages of the standardization process, network vendors have been able to build and deliver quality, mature products to the marketplace for many months. In addition, numerous interoperability demonstrations have been sponsored by the Gigabit Ethernet Alliance and other independent organizations, giving customers confidence in using Gigabit Ethernet products in their production networks.

Gigabit Ethernet Cabling and Distance Specifications

As shown in Figure 4, the physical (PHY) layer is a crucial part of the Gigabit Ethernet specification. It provides the interface between the media access control (MAC) layer and the transceivers in Gigabit Ethernet hardware. The PHY layer performs encoding, decoding, carrier sense, and link monitor functions.

The IEEE 802.3z Gigabit Ethernet standard ratified in June includes three physical layer specifications, two for fiber optic media—1000BASE-SX and 1000BASE-LX—and one for shielded copper media—1000BASE-CX. Another group, the IEEE 802.3ab Task Force, is defining the physical layer to run Gigabit Ethernet over the installed base of Category 5 unshielded twisted pair (UTP) cabling. This 1000BASE-T effort is scheduled for completion in March 1999. Figure 4 shows the various layers of Gigabit Ethernet and the distinction between the 802.3z specification and the 802.3ab specification. A fourth transceiver type, 1000BASE-LH, is a multivendor specification that defines optical transceivers that support distances greater than the 1000BASE-LX specification.

Fiber cabling specifications. There are three specifications for fiber cabling:

- **1000BASE-SX** ("S" for short wavelength) defines optical transceivers or physical layer devices for laser fiber cabling. The wavelength of SX lasers is 770 to 860 nanometers (nm) (commonly referred to as 850 nm). 1000BASE-SX is based on the Fibre Channel signaling specification and uses the 8B/10B

Table 3. *Gigabit Ethernet Distances for Fiber Optic Media*

Transceiver	Fiber Diameter (microns)	Bandwidth (MHz*km)	Minimum Range (meters)
1000BASE-SX	MM 62.5	160	2–220
	MM 62.5	200	2–275
	MM 50	400	2–500
	MM 50	500	2–550
1000BASE-LX	MM 62.5	500	2–550
	MM 50	400	2–550
	MM 50	500	2–550
	SM 9	NA	2–5000

MM = multimode fiber
SM = single-mode fiber

physical coding sublayer (PCS). Targeted for multimode fiber only, 1000BASE-SX transceivers are less costly than those found in products implementing the long wavelength specification.

- 1000BASE-LX ("L" for long wavelength) defines optical transceivers or physical layer devices for laser fiber cabling. The wavelength of LX lasers is 1270 to 1355 nm (commonly referred to as 1350 nm). 1000BASE-LX is also based on the Fibre Channel signaling specification and uses FC 8B/10B physical coding sublayer (PCS). 1000BASE-LX is specified for use on either multimode or single-mode fiber.

Table 3 shows the Gigabit Ethernet distance specifications for 1000BASE-SX and 1000BASE-LX fiber optic media. Note that the distance Gigabit Ethernet can reach depends on the bandwidth (measured in MHz*km)—the greater the bandwidth of the fiber, the further the distance supported. It's also important to note that IEEE specifies minimum rather than maximum ranges, and under average operating conditions, the minimum specified distance can be exceeded by a factor of three or four. However, most network managers are conservative when they design networks and use the IEEE specifications as the maximum distances.

- 1000BASE-LH ("LH" for long haul) is a multivendor specification; each vendor has a set of transceivers covering different distances. Although it is not an IEEE standard, many vendors are working to interoperate with IEEE 1000BASE-LX equipment and are using the gigabit interface connector (GBIC) multivendor specification in order to provide a common form factor and

Table 4. *Gigabit Ethernet Distances for 1000BASE-LH Fiber Optic Media*

Transceiver	Fiber Diameter (microns)	Wavelength (nm)	Minimum Range (meters)
1000BASE-LH (Extended distance)	SM 9	1310	1 km–49 km
1000BASE-LH (Extended distance)	SM 9	1550	50 km–100 km

Multimode vs. Single-Mode Fiber Cabling

Single-mode (SM) fiber cable reduces the radius of the fiber cable core to between 5 and 10 microns, which is very close to the order of a wavelength, so that one single angle or "mode" of light energy passes along the fiber. Because there is a single transmission path with single-mode fiber, the light can travel long distances with little loss of signal. Single-mode fiber is commonly used between buildings.

In contrast, multimode fiber has wider cores. Light is reflected along the core at multiple angles. The light is propagated along multiple paths, each path with a different length and hence a different time to traverse the fiber. These multiple angles or modes cause the signal elements to spread out in time. Consequently, distortions occur that limit the distance over which the integrity of the light signal can be maintained. Multimode fiber is the predominant type of LAN fiber installed within buildings and is less expensive than single-mode fiber.

greater flexibility. Table 4 shows specific wavelength distances recently qualified by 3Com; this list will grow as more of the GBIC distances undergoing qualification testing are approved. The appearance of these long-distance devices has sparked the interest of carriers in offering Gigabit Ethernet–based MAN (metropolitan area network) services. There are implementation caveats, however. Network planners must provide sufficient cabling attenuation to avoid a 40 km transceiver overdriving a 10 km transceiver on the receive end.

Copper cabling specifications. There are two copper cabling specifications:

- **1000BASE-CX** ("C" for copper) defines transceivers or PHY layer devices for shielded copper cabling. 1000BASE-CX is intended for short-haul copper connections (25 meters or less) within wiring closets. In 1998, no vendor had announced plans to develop 1000BASE-CX products. This physical interface uses the Fibre Channel 8B/10B physical coding sublayer (PCS), a different driver on shielded copper, and 150 ohm balanced cable, and it supports wiring distances up to 25 meters. (Note that IBM Type 1 cable is specifically mentioned in the IEEE 802.3z standard as not recommended.)
- **1000BASE-T** ("T" for twisted pair) is the developing specification for Gigabit Ethernet over four-pair Category 5 UTP copper cabling for distances up to 100 meters, enabling network managers to build networks with diameters of 200 meters. Extensive signal processing will be required on the 1000BASE-TX transceiver itself. The standards effort is scheduled for completion in June 1999.

Implementing Gigabit Ethernet in the Network

Assessing a new technology like Gigabit Ethernet involves many facets, including comparing it to other technologies and evaluating its impact on the existing network topology and network equipment. The topics discussed here are intended to help network managers consider both the advantages and limitations of Gigabit Ethernet.

Gigabit Ethernet and Link Aggregation

Link or port aggregation or trunking often comes up in discussions of Gigabit Ethernet migration. Link aggregation is the ability to support multiple, point-to-point, parallel active links between switches or between a switch and a server. The advantages of link aggregation are higher bandwidth, redundant links, and load sharing.

Link aggregation and Gigabit Ethernet are complementary technologies. If the business requirement is simply to add more band-

Table 5. *PCI Bus Bandwidth*

Bus Width	Bus Speed	Maximum Bus Bandwidth	When Available
32 bit	33 MHz	1 Gbps	Installed servers
64 bit	33 MHz	2 Gbps	Now shipping
64 bit	66 MHz	4 Gbps	1999 servers*

** At least one x86 non-Intel chipset server system supports 66 MHz PCI already.*

width between devices, and the devices can be upgraded to Gigabit Ethernet, Gigabit Ethernet is the right choice. If the business requirement includes the need for resilient and redundant links with load balancing, then link aggregation (aggregating multiple links into one logical connection) is the right choice.

Gigabit Ethernet links can of course themselves be aggregated with link aggregation technology. The primary application of link aggregation technology in the near future will be to build resilient, redundant links between Layer 2 and Layer 3 Gigabit Ethernet switches.

Many vendor-specific link aggregation implementations exist in the industry today. For example, 3Com supports link aggregation for FDDI, Fast Ethernet, and Gigabit Ethernet. And in keeping with its standards-based solution strategy, 3Com helped initiate the IEEE 802.3ad standards effort for link aggregation in 1998.

Optimizing Servers to Handle Gigabit Ethernet

While Gigabit Ethernet relieves the bottleneck at the server, server environments are not yet fully optimized to handle the entire available gigabit bandwidth. The good news is that server systems and their software and hardware components are rapidly evolving to ensure that the server system bandwidth capacity will handle the Gigabit Ethernet data rate. In addition, Gigabit Ethernet network interface cards (NICs) currently under design will overcome some of the server system bottlenecks and be optimized for the upcoming evolution in server architecture. Critical server system improvements are described below.

PCI Bus Bandwidth

The PCI bus is the predominant bus in x86 platforms, and is also available in some non-x86 systems. Although the original PCI bus present in earlier server systems had insufficient bandwidth to carry gigabit-speed I/O traffic, in newer server systems the bus is quickly becoming wider and faster and is not a bottleneck for the Gigabit Ethernet network connection.

While Table 5 shows PCI buses supporting gigabit rates today, the real bandwidth of the PCI bus is slightly lower than the maximum bus bandwidth due to the PCI bus overheads involved. At a minimum, the wider 64-bit bus at 33 MHz is required. In full-duplex mode, Gigabit Ethernet provides 2 Gbps bandwidth. Therefore, with the faster 66 MHz bus speed due out in the next few months, the PCI bus will no longer be the critical bottleneck in the system.

In a 64-bit 33 MHz PCI server, the PCI bus with its overhead will be slower than the full-duplex Gigabit Ethernet line rate. To solve the slowdown, the Gigabit Ethernet NIC uses onboard NIC memory to buffer packets while the PCI bus is busy or catching up. In future 64-bit 66 MHz PCI servers, the PCI bus will be faster than the line rate and the faster PCI bus will need to be offloaded before the packet is put on the wire for transmission. The NIC's onboard memory would then be used for storing the transmit packets. NIC memory can thus have an important impact on performance. The memory is typically carved into receive and transmit buffers. The NIC's onboard memory also serves to improve performance by storing and buffering packets on a heavily loaded bus with devices contending for access to the bus.

Flow control capability available on some Gigabit Ethernet NICs also helps to prevent buffer overflows: flow control signals the sender to pause if the receive buffer is about to overflow. This controlled slowdown in transmission provides better performance overall by eliminating costly timeouts and subsequent retransmissions.

Better techniques for efficiently utilizing the PCI bus are also being engineered. These techniques and innovations allow the PCI NIC to better interact with the bus, and improve performance and PCI bus utilization in next-generation servers.

Network Operating System

The history of networking has always been about moving bottlenecks to some other part of the network. With the development of Gigabit Ethernet, Ethernet packets travel at speeds of one billion bits per second. With the PCI bus developments described above, an I/O bus on the end station can support speeds of 1 billion bits per second or faster. The bottleneck has now moved from the network to inside the host itself and the challenges to increase performance are not Gigabit Ethernet issues or even network interface card issues, but fundamental network operating system (NOS) and host issues.

The performance of the server connection depends heavily on the network operating system and underlying protocols. UNIX operating systems appear better adapted to handling Gigabit Ethernet speeds, while the TCP/IP protocol running under Microsoft NT 4.0 still has much room for improvement. TCP/IP is a connection-oriented and complex protocol that requires high CPU bandwidth to process packets at gigabit per second rates.

Because it is a leading networking vendor and the Ethernet pioneer, 3Com is working closely with CPU and NOS vendors to push the envelope of host performance. With single-processor UNIX implementations, performance has been measured in the range of 500 to 800 Mbps. Industry testing with Windows NT 4.0 and the latest Pentium CPU-based server systems running TCP/IP show network performance in the range of 600 Mbps. Under Windows NT 4.0, throughput will increase with newer and more powerful Intel processor-based servers. The measured throughput on Windows NT single-processor Pentium systems has increased significantly from earlier Pentium servers to reach the current rate of greater than 600 Mbps. This throughput continues to increase with advancements in server system and processor technology. With the advent of Windows NT 5.0, these results will improve even more dramatically.

Under Windows NT 4.0, when greater CPU power is available to process the TCP/IP stack, server connection performance improves. In general, connection performance also improves for multiprocessor systems. However, the network driver interface specification's (NDIS's) single-threading interface limits the benefits of multi-CPU processing and provides less improvement and processor scalability. Microsoft plans to provide better-performance NDIS-based implementations in the future.

CPU Utilization at Gigabit Speeds

At gigabit speeds, routine networking tasks such as TCP/IP checksum calculations can easily tie up the processor, resulting in 100 percent CPU utilization. This leaves no processing power for running other applications. Therefore, well-designed Gigabit Ethernet NICs offload such tasks from the host processor, performing them in the NIC hardware. These Gigabit Ethernet NICs also consolidate interrupts, interrupting the CPU less frequently and for multiple tasks each time. This reduces the number of times the processor has to save context to service interrupts, and results in lower CPU utilization for better application performance.

To overcome some of the bottlenecks caused by the host, a few vendors have supported a proprietary solution known as "jumbo frames" to provide a better data payload-to-packet overhead ratio. A jumbo frame has a non-Ethernet standard maximum transmission unit (MTU), exceeding 1518 bytes. For example, one jumbo frame size supported is that of 9 Kbytes. Although they provide

Partnering Gigabit Ethernet with ATM in Enterprise Networks

Gigabit Ethernet discussions invariably involve a comparison with Asynchronous Transfer Mode (ATM) technology. Originally developed by the telephony industry, ATM has been successfully deployed in campus network and WAN applications. It offers speeds ranging from 155 Mbps through 622 Mbps, and soon 2.4 Gbps will be available. ATM was designed for large, scalable, resilient network backbones. One of its key objectives is to integrate telephony, data, voice, and video traffic onto one consolidated network for LANs and WANs.

ATM's explicit Quality of Service (QoS) capability—from the transmitting end station to the receiving end station over dedicated virtual circuits—and its ability to control jitter and limit latency across the network enable it to support real-time traffic. End stations can request a specific QoS level from the network, and ATM will deliver that service across a campus or wide area. Indeed, the ability to deliver QoS and support real-time, interactive traffic is one of ATM's greatest strengths.

Within the overall enterprise network, ATM and 10, 100, or 1000 Mbps Ethernet technologies can coexist and complement each other. The enterprise desktop is likely to remain Ethernet-based while the enterprise WAN will increasingly be based on ATM transport. From the ATM perspective, the issue of ATM versus Gigabit Ethernet becomes a question of how far into the backbone the ATM "cloud" penetrates. From the Ethernet perspective, the question is how far toward the WAN Ethernet will extend. Where ATM and Ethernet meet within the backbone depends on cost and the applications and services required, as well as on network size, topology, and redundancy requirements. The box below summarizes general guidelines on where to use ATM and Gigabit Ethernet technologies in your enterprise network.

some performance benefits, the major drawback of jumbo frames is that the frame format itself is proprietary. Any switch or NIC that supports jumbo frames cannot interoperate directly with other standards-based Ethernet switches and NICs on the network.

Memory Subsystems

Memory subsystem performance affects overall server performance and impacts network connection throughput. Memory with zero-wait states is typically better suited for servers since frequent access to the large contiguous data stores is usually required. While the initial latency will typically be higher for such memory, with large contiguous data accesses, the server throughput could reach a high peak sustained level.

Cache memory holds data that is accessed often or that is expected to be accessed next. This data is available for the NIC much faster than data off a disk storage subsystem. Cache memory and system memory help improve performance by reducing the impact of the

Guidelines for Using Gigabit Ethernet and ATM

Use Gigabit Ethernet when:

- Simplicity and cost are high priorities
- Seamless connectivity with Ethernet desktops is critical
- 802.1p/Q and RSVP can be leveraged for traffic prioritization and multimedia

Use ATM when:

- One infrastructure for voice/video/data is desired
- Meshed topologies for Layer 2 redundancy and load sharing are required
- Seamless LAN/MAN/WAN switching is important

storage access latencies. While a large amount of cache memory is beneficial to performance, the amount of cache memory optimum for the server varies based on the frequency and type of cache accesses. Look for flexibility in the server's cache size configuration.

The system bus speed between the processor and external (L2) cache is also important. The system bus speed (not to be confused with the PCI I/O bus speed) has been at 66 MHz until recently. In newer servers with processors supporting a 100 MHz system bus, the server throughput has improved markedly.

Storage Subsystems

Storage is typically the slowest component of the server. It usually uses rotating media technology with mechanical limitations. Fast and Wide SCSI, Fibre Channel, and disk arrays bring high performance to this area, but storage still remains the slowest link in the server performance chain. Fortunately, system memory and cache can also be used to bring higher performance to storage.

Gigabit Ethernet and Layer 3 Switching

Managing high-performance networks today involves more than the deployment of bandwidth and capacity. Network managers must also build control into the network. Traditionally, CPU-based routers have been used in the network to manage and control traffic between subnets, isolate faults, control chatty protocols, build firewalls, and so on. With the development of Gigabit Ethernet, interface speeds have reached speeds of one billion bits per second or almost 1.5 million packets per second with 64 byte packets. Traditional routers with their general-purpose CPUs and router code residing in memory cannot keep up with these speeds.

Recent developments in integrated circuit technology have enabled Layer 3 switches to attain wire-speed packet forwarding through application-specific integrated circuits (ASICs) that embed the Layer 3 routing intelligence in the switch hardware itself. This firmware intelligence performs multiple, simultaneous analysis and routing operations on packets. The Layer 3 switch is optimized for LAN routing and LAN protocols such as IP, IPX, and AppleTalk.

The Layer 3 switch is able to forward and route traffic at wire speed and make intelligent decisions about the type of traffic flowing through it. Layer 3 switching is actually interface independent: it can support 10 or 100 Mbps Ethernet, Gigabit Ethernet, FDDI, and ATM. Layer 3 switching technology is the key to successful ultra-high-speed backbone connectivity. Like intelligent, high-performance Layer 2 switches, Layer 3 switches support flexible deployment options such as port aggregation and virtual LANs. They also deliver advanced levels of network management with extensive built-in Layer 2 Remote Monitoring (RMON) capabilities, as well as higher-level RMON2 monitoring capabilities.

In addition, Layer 3 switches can use their packet-by-packet filtering capability to support Class of Service (CoS) and Quality of Service (QoS) based on the IEEE 802.1p (now 802.1D) and the 802.1Q standards as well as the IETF Resource Reservation Protocol (RSVP) bandwidth reservation scheme. Enterprise networks today need CoS and QoS capabilities to successfully converge different types of application traffic—from traditional IBM 3270 terminal and network backup applications to newer transaction processing, streaming video, and LAN telephony applications—onto one backbone infrastructure.

Gigabit Ethernet Migration Scenario

Gigabit Ethernet is best suited for unclogging network bottlenecks that occur in three main areas:

- Switch-to-switch connections
- Connections to high-speed servers
- Backbone connections

The following scenario details a network migration to Gigabit Ethernet. As shown in Figure 5, the initial building backbone in the network scenario is 10 Mbps Ethernet. Several Ethernet segments or workgroups are collapsed into a 10/100 Mbps switch, which in turn has several 10 Mbps Ethernet server connections. In the scenario:

- Power users are experiencing bottlenecks from their 10 Mbps links

- Users on the shared segments are experiencing slow response times
- All new desktops are equipped with 10/100 PCI cards

Figure 6 on page 14 implements the first upgrade phase in three areas: (1) upgrading the network backbone to 100 Mbps Fast Ethernet, (2) upgrading the power workgroup to 100 Mbps Fast Ethernet connections between the end stations and the switch, and (3) implementing 10 Mbps switching in other workgroups that need dedicated bandwidth. The result? Power users immediately acquire 100 Mbps connection speeds, and other workgroups that need it enjoy dedicated 10 Mbps bandwidth while the investment in existing switches and NICs is preserved. The speed of the backbone also increases tenfold to accommodate the overall increase in network bandwidth demand.

Figure 7 on page 14 shows upgrading the riser downlinks to Gigabit Ethernet to increase backbone bandwidth even further. The switches in the wiring closets that support power users or large workgroups are upgraded with Gigabit Ethernet downlink modules, the basement switch is upgraded to 100/1000 capability, and key servers are upgraded with Gigabit Ethernet NICs. A key point to notice is that the upgrades of the backbone to Gigabit Ethernet support the extension of switching to the edges. Both office workgroups now support switched 10 or switched 10/100 Ethernet to the desktops.

Also in Figure 7, the Fast Ethernet backbone switch that aggregates multiple 10/100 switches is upgraded to a Gigabit Ethernet switch, supporting multiple 100/1000 switches. Gigabit full-duplex repeaters are installed as needed to aggregate servers or build Gigabit Ethernet support for workgroups that manipulate very large multimedia and graphics files. In addition, once the backbone is upgraded to a Gigabit Ethernet switch, high-performance server farms can be connected directly to the backbone with Gigabit Ethernet NICs, increasing throughput to the servers for users with high-bandwidth application files or for bandwidth-intensive data warehousing and backup operations. Overall, the network now supports a greater number of segments and more bandwidth per segment.

Summary

Gigabit Ethernet is an IEEE-standard technology that preserves Ethernet's simplicity, leverages its cost drivers, and takes advantage of the

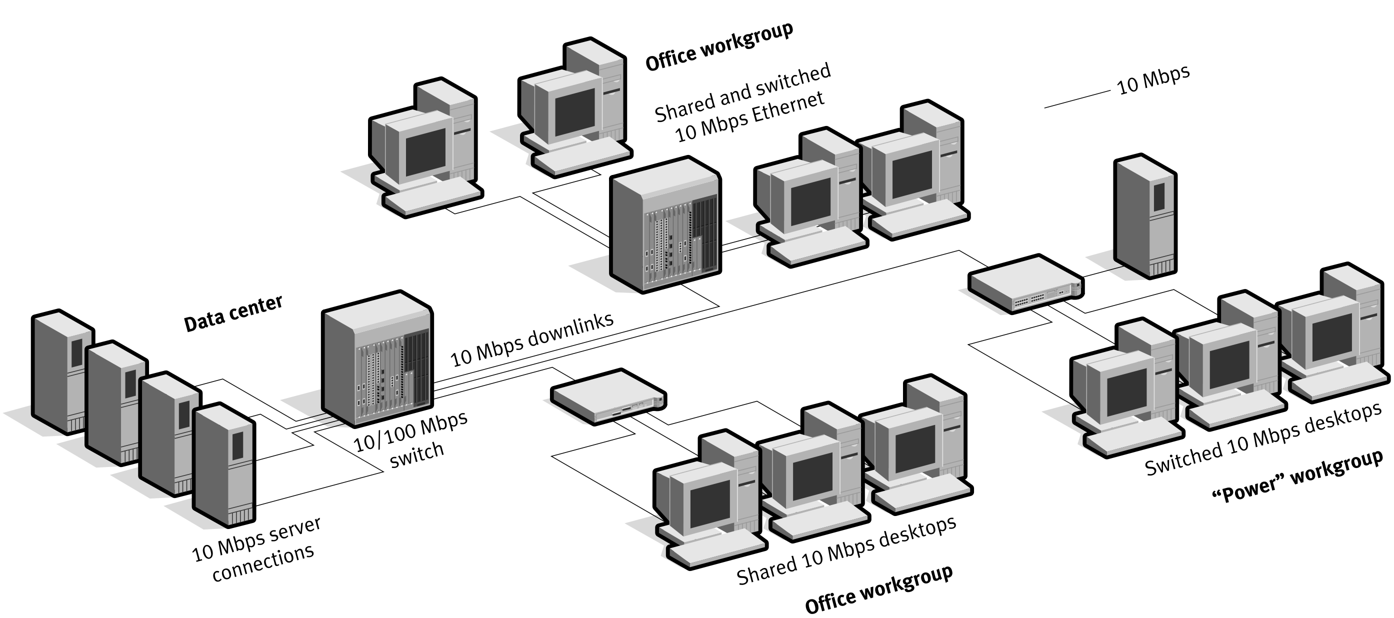

Figure 5. *Scaling Network Capacity: Starting Point*

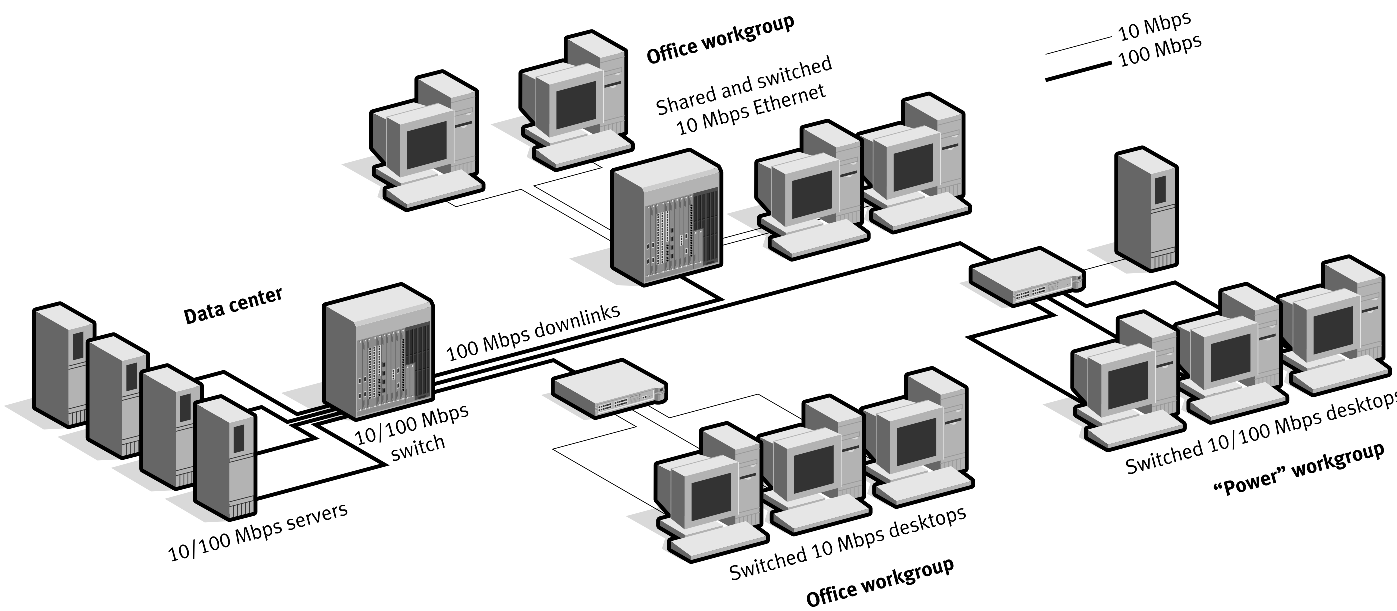

Figure 6. *Scaling Network Capacity: Phase 1*

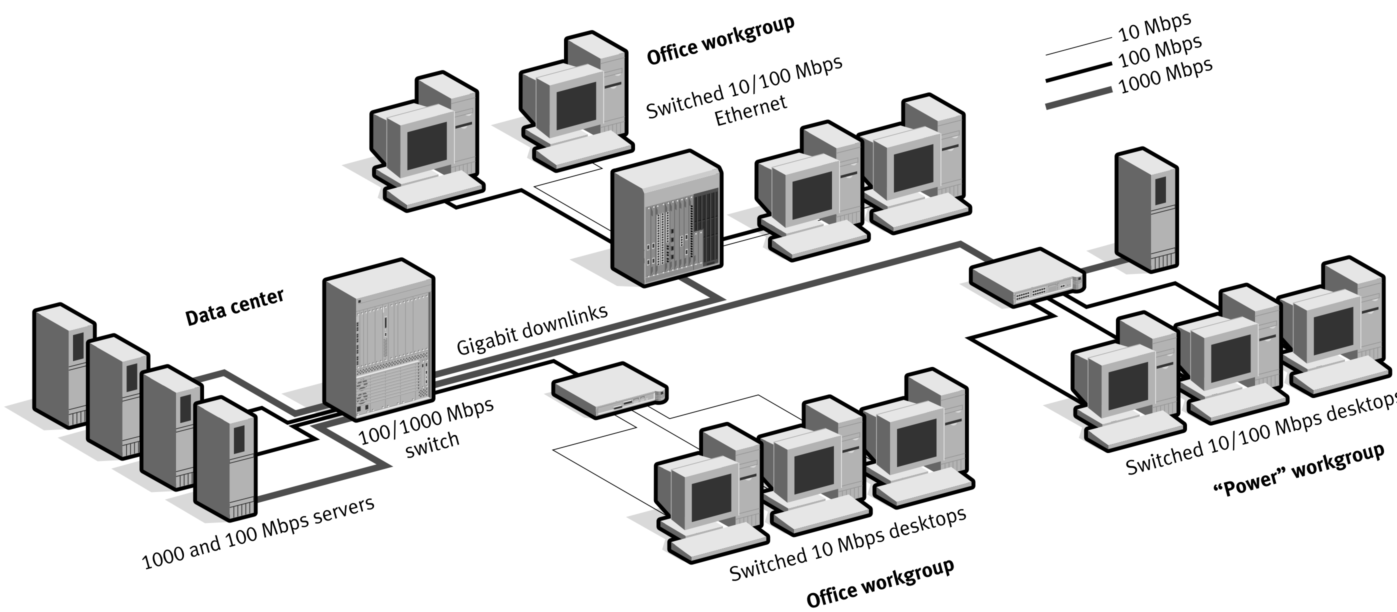

Figure 7. *Scaling Network Capacity: Phase 2*

installed base of Ethernet devices. It also takes advantage of the familiarity, knowledge, and expertise network managers have with both Ethernet- and IP-based traffic and services. The deployment of Gigabit Ethernet in the network increases the ability of network planners and managers to implement more switched Ethernet and switched Fast Ethernet to the horizontal wiring closets and servers and out to the desktops.

Gigabit Ethernet can be implemented incrementally into an enterprise network, replacing Fast Ethernet core switches and enabling the deployment of more switched 10 and switched 10/100 Mbps Ethernet to the edge of the network. 1000BASE-SX supports

For Further Reference

1 3Com Gigabit Ethernet solutions Web page, www.3com.com/technology/gigabit.html
2 *Layer 3 Switching: An Introduction* white paper, www.3com.com/technology/tech_net/white_papers/50066001.html
3 *3Com Link Aggregation* white paper, www.3com.com/technology/tech_net/white_papers/500666.html
4 Gigabit Ethernet Test Consortium Web site, www.iol.unh.edu/consortiums/ge/
5 IEEE 803.3z Task Force Web site, grouper.ieee.org/pub/802_main/802.3/gigabit/
6 Gigabit Ethernet Alliance Web site, www.gigabit-ethernet.org/technology/whitepapers

multimode fiber up to 550 meters, while 1000BASE-LX supports multimode up to 550 meters and single-mode fiber up to 5000 meters. Products based on 1000BASE-SX and 1000BASE-LX are shipping in volume today. The 1000BASE-T standard to support Gigabit Ethernet operation on the installed base of Category 5 UTP copper cabling is targeted for completion by March 1999. Hardware and software server system and server NIC development is well under way to ensure that the server system bandwidth capacity will handle Gigabit Ethernet data rates, an issue of real concern in Gigabit Ethernet implementation. Layer 3 switching technology adds the critical component of "brain" to Gigabit Ethenet's "brawn."

3Com Corporation is the leader in Gigabit Ethernet development and is committed to the adoption of standards-based Gigabit Ethernet in the marketplace. The company is a founding member of the Gigabit Ethernet Alliance, the main industry consortium committed to accelerating the Gigabit Ethernet standards process, end-user education, and multivendor interoperability. 3Com is also a founding member of the Gigabit Ethernet Consortium at the Interoperability Lab at the University of New Hampshire, which develops test suites and conducts group and individual product testing. 3Com offers the industry's most complete Gigabit Ethernet solution for medium to large enterprises—from Gigabit Ethernet server NICs and full-duplex repeaters to workgroup and core Layer 3 Gigabit Ethernet switches.

3Com®

More connected.™

3Com Corporation enables individuals and organizations worldwide to communicate and share information and resources anytime, anywhere. As one of the world's preeminent suppliers of data, voice, and video communications technology, 3Com has delivered networking solutions to more than 200 million customers worldwide. The company provides large enterprises, small and medium enterprises, carriers and network service providers, and consumers with comprehensive, innovative information access products and system solutions for building intelligent, reliable, and high-performance local and wide area networks.

3Com Corporation
P.O. Box 58145
5400 Bayfront Plaza
Santa Clara, CA
95052-8145
Phone: 1 800 NET 3Com
or 1 408 326 5000
Fax: 1 408 326 5001
World Wide Web:
www.3com.com

Asia Pacific Rim

Sydney, Australia
Phone: 61 2 9937 5000
Fax: 61 2 9956 6247

Melbourne, Australia
Phone: 61 3 9934 8888
Fax: 61 3 9934 8880

Beijing, China
Phone: 86 10 68492 568
Fax: 86 10 68492 789

Shanghai, China
Phone: 86 21 6350 1581
Fax: 86 21 6350 1531

Hong Kong
Phone: 852 2501 1111
Fax: 852 2537 1149

India
Phone: 91 11 644 3974
Fax: 91 11 623 3192

Indonesia
Phone: 62 21 572 2088
Fax: 62 21 572 2089

Osaka, Japan
Phone: 81 6 536 3303
Fax: 81 6 536 3304

Tokyo, Japan
Phone: 81 3 5977 3266
Phone: 0120 313266
(toll free from Japan)
Fax: 81 3 5977 3370

Korea
Phone: 82 2 3455 6300
Fax: 82 2 319 4710

Malaysia
Phone: 60 3 715 1333
Fax: 60 3 715 2333

New Zealand
Phone: 64 9 366 9138
Fax: 64 9 366 9139

Philippines
Phone: 632 892 4476
Fax: 632 811 5493

Singapore
Phone: 65 538 9368
Fax: 65 538 9369

Taiwan
Phone: 886 2 2 377 5850
Fax: 886 2 2 377 5860

Thailand
Phone: 662 231 8151 5
Fax: 662 231 8158

3Com Austria
Phone: 43 1 580 17 0
Fax: 43 1 580 17 20

3Com Benelux B.V.

Belgium
Phone: 32 2 725 0202
Fax: 32 2 720 1211

Netherlands
Phone: 31 346 58 62 11
Fax: 31 346 58 62 22

3Com Canada

Calgary
Phone: 1 403 265 3266
Fax: 1 403 265 3268

Edmonton
Phone: 1 403 423 3266
Fax: 1 403 423 2368

Montreal
Phone: 1 514 683 3266
Fax: 1 514 683 5122

Ottawa
Phone: 1 613 566 7055
Fax: 1 613 233 9527

Toronto
Phone: 1 416 498 3266
Fax: 1 416 498 1262

Vancouver
Phone: 1 604 434 3266
Fax: 1 604 434 3264

3Com Eastern Europe/CIS

Bulgaria
Phone: 359 2 962 5222
Fax: 359 2 962 4322

Czech/Slovak Republics
Phone: 420 2 21845 800
Fax: 420 2 21845 811

Hungary
Phone: 36 1 250 8341
Fax: 36 1 250 8347

Poland
Phone: 48 22 6451351
Fax: 48 22 6451352

Russia
Phone: 7 095 258 09 40
Fax: 7 095 258 09 41

3Com France
Phone: 33 1 69 86 68 00
Fax: 33 1 69 07 11 54
Carrier and Client Access
Phone: 33 1 41 97 46 00
Fax: 33 1 49 07 03 43

3Com GmbH

Munich, Germany
Phone: 49 89 627320
Fax: 49 89 62732233

3Com Iberia

Portugal
Phone: 351 1 3404505
Fax: 351 1 3404575

Spain
Phone: 34 1 509 69 00
Fax: 34 1 307 79 82

3Com Latin America

U.S. Headquarters
Phone: 1 408 326 2093
Fax: 1 408 764 5730

Miami, Florida
Phone: 1 305 461 8400
Fax: 1 305 461 8401

Argentina
Phone: 54 1 312 3266
Fax: 54 1 314 3329

Brazil
Phone: 55 11 246 5001
Fax: 55 11 246 3444

Chile (serving Chile, Bolivia, and Peru)
Phone: 562 240 6200
Fax: 562 6240 6231

Colombia
Phone: 57 1 629 4847
Fax: 57 1 629 4503

Mexico
Phone: 52 5 520 7841/7847
Fax: 52 5 520 7837

Peru
Phone: 51 1 221 5399
Fax: 51 1 221 5499

Venezuela
Phone: 58 2 953 8122
Fax: 58 2 953 9686

3Com Mediterraneo

Milan, Italy
Phone: 39 02 253011
Fax: 39 02 27304244

Rome, Italy
Phone: 39 6 5279941
Fax: 39 6 52799423

3Com Middle East
Phone: 971 4 319533
Fax: 971 4 316766

3Com Nordic AB

Denmark
Phone: 45 48 10 50 00
Fax: 45 48 10 50 50

Finland
Phone: 358 9 435 420 67
Fax: 358 9 455 51 66

Norway
Phone: 47 22 58 47 00
Fax: 47 22 58 47 01

Sweden
Phone: 46 8 587 05 600
Fax: 46 8 587 05 601

3Com Southern Africa
Phone: 27 11 807 4397
Fax: 27 11 803 7405

3Com Switzerland
Phone: 41 844 833 933
Fax: 41 844 833 934

3Com UK Ltd.

Edinburgh
Phone: 44 131 240 2900
Fax: 44 131 240 2903

Ireland
Phone: 353 1 820 7077
Fax: 353 1 820 7101

Manchester
Phone: 44 161 873 7717
Fax: 44 161 873 8053

Winersh
Phone: 44 1189 27 8200
Fax: 44 118 969 5555

To learn more about 3Com products and services, visit our World Wide Web site at www.3com.com. 3Com is a publicly traded corporation (Nasdaq: COMS).

© 1998 3Com Corporation. All rights reserved. 3Com, the 3Com logo, and SuperStack are registered trademarks of 3Com Corporation. *More connected.* is a trademark of 3Com Corporation. AppleTalk is a trademark of Apple Computer. Pentium is a trademark of Intel. Windows and Windows NT are trademarks of Microsoft. IPX is a trademark of Novell. Other brand or product names may be trademarks or registered trademarks of their respective owners. All specifications are subject to change without notice.

Printed in U.S.A. on recycled paper

503003-002 6/99