

Defining IT Infrastructure

Levels of IT Infrastructure

Evolution of IT Infrastructure: 1950–2005

Technology Drivers of Infrastructure Evolution

New World Telecommunications enhanced its efficiency and competitiveness by better selection and management of its hardware and software technology. The company consolidated servers, updated its storage technology, and implemented software that could integrate its key business systems. These investments gave the company a more flexible IT infrastructure for meeting current and future business needs.

As a manager, you'll face many decisions about investments in hardware, software, and other elements of IT infrastructure that are necessary for your firm to conduct its business—and which may even give your firm new capabilities to surge ahead of competitors. This chapter provides an overview of the technology and service components of IT infrastructure and the most important trends in hardware and software platforms. We look at the capabilities of major technology vendors as well as the technologies themselves to help managers make intelligent choices about the appropriate technology platform for their firm. If you would like a more introductory discussion of hardware and software or a review of basic hardware and software concepts, you can find this information at the Laudon Web site for Chapter 6.

IT INFRASTRUCTURE

In Chapter 1, we defined information technology (IT) infrastructure as the shared technology resources that provide the platform for the firm's specific information system applications. IT infrastructure includes investment in hardware, software, and services—such as consulting, education, and training—that are shared across the entire firm or across entire business units in the firm. For instance, a firm's investment in thousands of new desktop computers networked together and linked to the Internet is an infrastructure investment because it serves many groups, goals, and business initiatives.

Supplying U.S. firms with IT infrastructure is a \$1.8 trillion dollar industry when telecommunications, networking equipment and telecommunications services (Internet, telephone, and data transmission) are included. Investments in infrastructure account for between 25 and 35 percent of information technology expenditures in large firms (Weill et al., 2002).

A firm's IT infrastructure provides the foundation for serving customers, working with vendors, and managing internal firm business processes (see Figure 6-1). In this sense, IT infrastructure defines the capabilities of the firm today and in the near term of, say, three to five years (the length of time it takes to make a significant change in the firm's IT infrastructure). For instance, if you want to provide customers with a daily balance of cell phone charges on a Web site, and if you plan to use this capability as a strategic differentiating feature of your firm, this implies certain infrastructure capabilities. If you are a bank and want to sell banking services anywhere in the United States or the world to all your customers, whether they have car loans, home loans, brokerage

accounts, or just checking accounts, you will need an infrastructure that crosses these different lines of business and operates on a global enterprise-wide basis.




FIGURE 6-1 The connection between the firm, IT infrastructure, and business capabilities

The services a firm is capable of providing to its customers, suppliers, and employees are a direct function of its IT infrastructure. Ideally, this infrastructure should support the firm's business and information systems strategy. New information technologies have a powerful impact on business and IT strategies, as well as the services that can be provided to customers.

Defining IT Infrastructure

There are two ways of defining IT infrastructure: as technology or as services clusters. In one sense IT infrastructure is like the plumbing or electrical systems in a building: a set of physical devices and software applications that are required to operate the entire enterprise. But an even more useful service-based definition focuses on the services provided by all this hardware and software. In this definition, IT infrastructure is a set of firmwide services budgeted by management and comprising both human and technical capabilities (Weill et al., 2002). These services include the following:

- Computing platforms used to provide computing services that connect employees,

customers, and suppliers into a coherent digital environment, including large mainframes, desktop and laptop computers, and personal digital assistants (PDAs) and Internet appliances.

- Telecommunications services that provide data, voice, and video connectivity to employees, customers, and suppliers.
- Data management services that store and manage corporate data and provide capabilities for analyzing the data.
- Application software services that provide enterprise-wide capabilities such as enterprise resource planning, customer relationship management, supply chain management, and knowledge management systems that are shared by all business units.
- Physical facilities management services that develop and manage the physical installations required for computing, telecommunications, and data management services.
- IT management services that plan and develop the infrastructure, coordinate with the business units for IT services, manage accounting for the IT expenditure, and provide project management services.
- IT standards services that provide the firm and its business units with policies that determine which information technology will be used, when, and how.
- IT education services that provide training in system use to employees and offer managers training in how to plan for and manage IT investments.
- IT research and development services that provide the firm with research on potential future IT projects and investments that could help the firm differentiate itself in the marketplace.

This “service platform” perspective makes it easier to understand the business value provided by infrastructure investments. For instance, the real business value of a fully loaded Pentium 4 PC operating at 2 gigahertz that costs about \$1,500 with a high-speed Internet connection is hard to understand without knowing who will use it and how it will be used. When we look at the services provided by these tools, however, their value becomes more apparent: The new PC makes it possible for a high-cost employee making \$100,000 a year to connect to all the company’s major systems and the public Internet. The high-speed Internet service saves this employee about one hour per day in reduced wait time for corporate and Internet information. Without this PC and Internet connection, the value of this one employee to the firm might be cut in half.

Levels of IT Infrastructure

Firm infrastructure is organized at three major levels: public, enterprise, and business unit (Figure 6-2). There may be other lower levels, such as departments or individual employees, but for now we focus on these three.

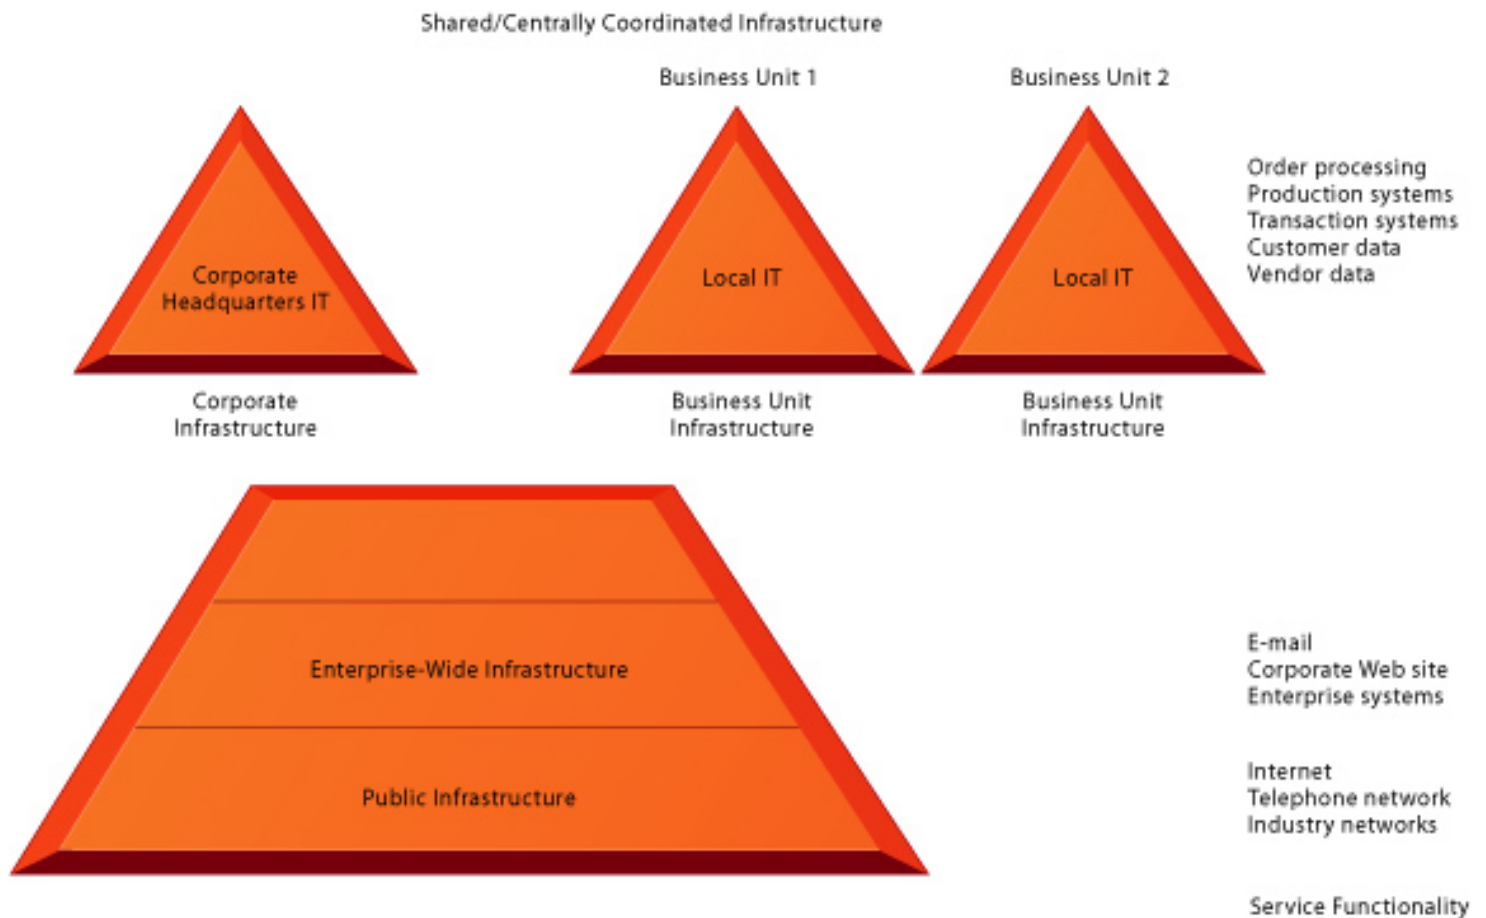


FIGURE 6-2 Levels of IT infrastructure

IT infrastructure exists at three different levels: public, enterprise, and business unit. Each level of infrastructure provides a set of IT services and capabilities.

All firms are dependent on public IT infrastructure, which includes the Internet, the public switched telephone network, industry-operated networks, and other IT support facilities such as cable systems and cellular networks. Enterprise-wide infrastructure includes services such as e-mail, a central corporate Web site, corporate-wide intranets, and an increasing array of enterprise-wide software applications (see Chapter 11). Business units also have their own infrastructure that is uniquely suited to their line of business such as specialized production software and systems, customer and vendor systems, and local order entry and other transaction systems.

In multiunit businesses, typical of most large firms, a central corporate

infrastructure is also used to manage the entire enterprise, receive reports from business units, and exercise central oversight.

Figure 6-2 raises some interesting questions: What, for instance, should be handled by enterprise infrastructure and what by local business units? How much should local business units be charged (and on what basis) for use of enterprise architecture? Local business units may have unique demands that an enterprise-wide infrastructure may not address.

[Return to Top](#)

Evolution of IT Infrastructure: 1950–2005


The IT infrastructure in organizations today is an outgrowth of over 50 years of evolution in computing platforms. We have identified five stages in this evolution, each representing a different configuration of computing power and infrastructure elements (see Figure 6-3). The five eras are automated special-purpose machines, general-purpose mainframe and minicomputer computing, personal computers, client/server networks, and enterprise and Internet computing.

Stages in IT Infrastructure Evolution

Electronic
Accounting
Machine
(1930–1950)



Mainframe/
Minicomputer
(1959–present)



Personal
Computer
(1981–present)



Client Server






FIGURE 6-3 Eras in IT infrastructure evolution

Illustrated here are the typical computing configurations characterizing each of the five eras of IT infrastructure evolution.

These eras do not necessarily end for all organizations at the same time, and the technologies that characterize one era may also be used in another time period for other purposes. For example, some companies still run traditional mainframe or minicomputer systems. Mainframe computers today are used as massive servers supporting large Web sites and corporate enterprise applications. More detail on infrastructure history can be found on the Laudon Web site for Chapter 6.

ELECTRONIC ACCOUNTING MACHINE ERA: 1930–1950

The first era of business computing used specialized machines that could sort computer cards into bins, accumulate totals, and print reports (DaCruz, 2004). Although the electronic accounting machine was an efficient processor of accounting tasks, the machines were large and cumbersome. Software programs were hardwired into circuit boards, and they could be changed by altering the wired connections on a patch board. There were no programmers, and a human machine operator was the operating system, controlling all

system resources.

GENERAL-PURPOSE MAINFRAME AND MINICOMPUTER ERA: 1959 TO PRESENT

The first commercial all-electronic vacuum tube computers appeared in the early 1950s with the introduction of the UNIVAC computers and the IBM 700 Series. Not until 1959 with the introduction of the IBM 1401 and 7090 transistorized machines did widespread commercial use of mainframe computers begin in earnest. In 1965, the general-purpose commercial mainframe computer truly came into its own with the introduction of the IBM 360 series. The 360 was the first commercial computer with a powerful operating system that could provide time sharing, multitasking, and virtual memory in more advanced models.

Mainframe computers eventually became powerful enough to support thousands of online remote terminals connected to a centralized mainframe using proprietary communication protocols and proprietary data lines. The first airline reservation systems appeared in 1959 and became the prototypical online, real-time interactive computing system that could scale to the size of an entire nation.

IBM dominated mainframe computing from 1965 onward and still dominates this \$27 billion global market in 2004. Today IBM mainframe systems can work with a wide variety of different manufacturers' computers and multiple operating systems on client/server networks and networks based on Internet technology standards.

The mainframe era was a period of highly centralized computing under the control of professional programmers and systems operators (usually in a corporate data center), with most elements of infrastructure provided by a single vendor, the manufacturer of the hardware and the software. This pattern began to change with the introduction of minicomputers produced by Digital Equipment Corporation (DEC) in 1965. DEC minicomputers (PDP-11 and later the VAX machines) offered powerful machines at far lower prices than IBM mainframes, making possible decentralized computing, customized to the specific needs of individual departments or business units rather than time sharing on a single huge mainframe.

PERSONAL COMPUTER ERA: (1981 TO PRESENT)

Although the first truly personal computers (PCs) appeared in the 1970s (the Xerox Alto, MIT's Altair, and the Apple I and II, to name a few), these machines had only limited distribution to computer enthusiasts. The appearance of the IBM PC in 1981 is usually credited as the beginning of the PC era because this machine was the first to become widely adopted in American businesses. At first using the DOS operating system, a text-based command language, and later the Microsoft Windows operating system, the Wintel PC computer (Windows operating system software on a computer with an Intel microprocessor) became the standard desktop personal computer. Today, 95 percent of the world's estimated 1 billion computers use the Wintel standard.

Proliferation of PCs in the 1980s and early 1990s launched a spate of personal desktop productivity software tools—word processors, spreadsheets, electronic presentation software, and small data management programs—that were very valuable to both home and corporate users. These PCs were standalone systems until PC operating system software in the 1990s made it possible to link them into networks.

CLIENT/SERVER ERA (1983 TO PRESENT)

In client/server computing, desktop or laptop computers called clients are networked to server computers that provide the client computers with a variety of services and capabilities. Computer processing work is split between these two types of machines. The client is the user point of entry, whereas the server provides communication among the clients, processes and stores shared data, serves up Web pages, or manages network activities. The term server refers to both the software application and the physical computer on which the network software runs. The server could be a mainframe, but today server computers typically are more powerful versions of personal computers, based on inexpensive Intel chips and often using multiple processors in a single computer box.

The simplest client/server network consists of a client computer networked to a server computer, with processing split between the two types of machines. This is called a two-tiered client/server architecture. Whereas simple client/server networks can be found in small businesses, most corporations have more complex, multitiered (often called N-tier) client/server architectures in which the work of the entire network is balanced over several different levels of servers, depending on the kind of service being requested (see Figure 6-4).




FIGURE 6-4 A multitiered client/server network (N-tier)

In a multitiered client/server network, client requests for service are handled by different levels of servers.

For instance, at the first level a Web server will serve a Web page to a client in response for a request for service. Web server software is responsible for locating and managing stored Web pages. If the client requests access to a corporate system (a product list or price information, for instance), the request is passed along to an application server. Application server software handles all application operations between a user and an organization's back-end business systems. The application server may reside on the same computer as the Web server or on its own dedicated computer. Chapters 7 and 8 provide

more detail on other pieces of software that are used in multitiered client/server architectures for e-commerce and e-business.

Client/server computing enables businesses to distribute computing work across a number of smaller, inexpensive machines that cost much less than minicomputers or centralized mainframe systems. The result is an explosion in computing power and applications throughout the firm.

Novell Netware was the leading technology for client/server networking at the beginning of the client/server era. Today Microsoft is the market leader, with its Windows operating systems (Windows Server, Windows XP, Windows 2000), controlling 78 percent of the local area network market.

ENTERPRISE INTERNET COMPUTING ERA (1992 TO PRESENT)

The success of the client/server model posed a new set of problems for corporations. Many large firms found it difficult to integrate all of their local area networks (LANs) into a single, coherent corporate computing environment. Applications developed by local departments and divisions in a firm, or in different geographic areas, could not communicate easily with one another and share data.

In the early 1990s, firms turned to networking standards and software tools that could integrate disparate networks and applications throughout the firm into an enterprise-wide infrastructure. As the Internet developed into a trusted communications environment after 1995, business firms began using the Transmission Control Protocol/Internet Protocol (TCP/IP) networking standard to tie their disparate networks together. We discuss TCP/IP in detail in Chapter 8.

The resulting IT infrastructure links different types and brands of computer hardware and smaller networks into an enterprise-wide network so that information can flow freely across the organization and between the firm and other organizations. Enterprise networks link mainframes, servers, PCs, mobile phones, and other handheld devices, and connect to public infrastructures such as the telephone system, the Internet, and public network services.

The enterprise infrastructure employs software that can link disparate applications and enable data to flow freely among different parts of the business. Chapter 2 describes how enterprise applications perform this function. Other solutions for enterprise integration include enterprise application integration software, Web services, and outsourcing to external vendors that provide hardware and software for a comprehensive enterprise infrastructure. We discuss these solutions in detail in section 6.4 and in Chapter 11.

The enterprise era promises to bring about a truly integrated computing and IT services platform for the management of global enterprises. The hope is to deliver critical business information painlessly and seamlessly to decision makers when and where they need it to create customer value. This could be everything from getting inventory data to the mobile salesperson in the customer's office, to helping a customer at a call center with a problem customer, or providing managers with precise up-to-the-minute information on company performance.

That is the promise, but the reality is wrenchingly difficult and awesomely

expensive. Most large firms have a huge, tangled web of hardware systems and software applications inherited from the past. This makes achieving this level of enterprise integration a difficult, long-term process that can last perhaps as long as a decade and cost large companies hundreds of millions of dollars. Table 6-1 compares each era on the infrastructure dimensions discussed above.

TABLE 6-1 Stages in the IT Infrastructure Evolution

Infrastructure Dimension	Electronic Accounting Machine Era (1930–1950)	Mainframe Era (1959 to Present)	PC Era (1981 to Present)	Client/Server Era (1983 to Present)	Enterprise Era (1992 to Present)
Signature Firm(s)	IBM Burroughs NCR	IBM	Microsoft/Intel Dell HP IBM	Novell Microsoft	SAP Oracle PeopleSoft
Hardware Platform	Programmable card sorters	Centralized mainframe	Wintel computers	Wintel computers	Multiple: <ul style="list-style-type: none"> • Mainframe • Server • Client
Operating System	Human operators	IBM 360 IBM 370 Unix	DOS/Windows Linux IBM 390	Windows 3.1 Windows Server Linux	Multiple: <ul style="list-style-type: none"> • Unix/Linux • OS 390 • Windows Server
Application and Enterprise Software	None; application software created by technicians	Few enterprise-wide applications; departmental applications created by in-house programmers	No enterprise connectivity; boxed software	Few enterprise-wide applications; boxed software applications for workgroups and departments	Enterprise-wide applications linked to desktop and departmental applications: <ul style="list-style-type: none"> • mySAP • Oracle E-Business Suite • PeopleSoft Enterprise One
Networking/ Telecommunications	None	Vendor-provided: <ul style="list-style-type: none"> • Systems Network Architecture (IBM) • DECNET (Digital) • AT&T voice 	None or limited	Novell NetWare Windows 2003 Linux AT&T voice	LAN Enterprise-wide area network (WAN) TCP/IP Internet standards-enabled
System Integration	Vendor-provided	Vendor-provided	None	Accounting and consulting firms Service firms	Software manufacturer Accounting and consulting firms System integration firms Service firms
Data Storage and Database Management	Physical card management	Magnetic storage Flat files Relational databases	DBase II and III Access	Multiple database servers with optical and magnetic storage	Enterprise database servers
Internet Platforms	None	Poor to none	None at first Later browser-enabled clients	None at first Later: <ul style="list-style-type: none"> • Apache server • Microsoft IIS 	None in the early years Later: <ul style="list-style-type: none"> • Intranet- and Internet-enabled

ENHANCED CLIENTS

- Apache server
- Microsoft IIS

SERVERS

- Intranet- and Internet-delivered enterprise services
- Large server farms

[Return to Top](#)

Technology Drivers of Infrastructure Evolution

The changes in IT infrastructure we have just described have resulted from developments in computer processing, memory chips, storage devices, telecommunications and networking hardware and software, and software design that have exponentially increased computing power while exponentially reducing costs. Let's look at the most important developments.

MOORE'S LAW AND MICROPROCESSING POWER

In 1965, Gordon Moore, the director of Fairchild Semiconductor's Research and Development Laboratories, an early manufacturer of integrated circuits, wrote in *Electronics* magazine that since the first microprocessor chip was introduced in 1959, the number of components on a chip with the smallest manufacturer costs per component (generally transistors) had doubled each year. This assertion became the foundation of Moore's Law. Moore later reduced the rate of growth to a doubling every two years (Tuomi, 2002).

This law would later be interpreted in multiple ways. There are at least three variations of Moore's Law, none of which Moore ever stated: (1) the power of microprocessors doubles every 18 months (Tuomi, 2002); (2) computing power doubles every 18 months; and (3) the price of computing falls by half every 18 months.

Figure 6-5 illustrates the relationship between number of transistors on a microprocessor and millions of instructions per second (MIPS), a common measure of processor power. Figure 6-6 shows the exponential decline in the cost of transistors and rise in computing power.




FIGURE 6-5 Moore's Law and microprocessor performance

Packing more transistors into a tiny microprocessor has exponentially increased processing power.




FIGURE 6-6 Falling cost of chips

Packing more transistors into less space has driven down transistor cost dramatically as well as the cost of the products in which they are used.

There is reason to believe the exponential growth in the number of transistors and the power of processors coupled with an exponential decline in computing costs will continue into the future. Chip manufacturers continue to miniaturize components. Intel has recently changed its manufacturing process from 0.13-micron component size (a micron is a millionth of a meter), introduced in 2002, to a newer 90-nanometer process in 2004 (a nanometer is a billionth of a meter). With a size of about 50 nanometers, today's transistors should no longer be compared to the size of a human hair but rather to the size of a virus, the smallest form of organic life.

By using nanotechnology, Intel believes it can shrink the size of transistors down to the width of several atoms. Nanotechnology uses individual atoms and molecules to create computer chips and other devices that are thousands of times smaller than current technologies permit. IBM and other research labs have created transistors from nanotubes and other electrical devices (IBM, 2004) and have developed a manufacturing process that could produce nanotube processors economically (Figure 6-7). Other new technologies include strained silicon, 300-millimeter production wafers (which decrease the costs of production), and denser interconnections among components.




FIGURE 6-7 Examples of nanotubes

Nanotubes are tiny tubes about 10,000 times thinner than a human hair. They consist of rolled up sheets of carbon hexagons. Discovered in 1991 by researchers at NEC, they have the potential uses as minuscule wires or in ultrasmall electronic devices and are very powerful conductors of electrical current.

Whereas the first Pentium microprocessors operated at 75 megahertz, today's Pentiums are available with 3-gigahertz speeds. However, increasing processor speeds at the same exponential rate as in the past may no longer be possible. As processor speeds increase, heat is generated that cannot be dissipated with air fans.

Another brake on future increases in microprocessor speed is more market-oriented: Most consumers may not need vast increases in microprocessor speed but instead are more interested in low power consumption for longer battery life and low weight to increase laptop and handheld computer portability. For this reason, Intel and other firms are designing the next generation of chips to be less power hungry and lower in weight even if they are the same or even slower speeds. Other options include putting

multiple processors on a single chip.

THE LAW OF MASS DIGITAL STORAGE

A second technology driver of IT infrastructure change is the Law of Mass Digital Storage. The world produces as much as 5 exabytes of unique information per year (an exabyte is a billion gigabytes, or 10^{18} bytes). The amount of digital information is roughly doubling every year (Lyman and Varian, 2003). Almost all of this information growth involves magnetic storage of digital data, and printed documents account for only 0.003 percent of the annual growth.

Fortunately, the cost of storing digital information is falling at an exponential rate. Figure 6-8 shows that PC hard drive capacity—beginning with a Seagate 506 in 1980 that had 5 megabytes of memory—has grown at a compound annual growth rate of 25 percent in the early years to over 60 percent a year since 1990. Today's PC hard drives have storage densities approaching 1 gigabyte per square inch and total capacities of over 200 gigabytes (IBM, Seagate).

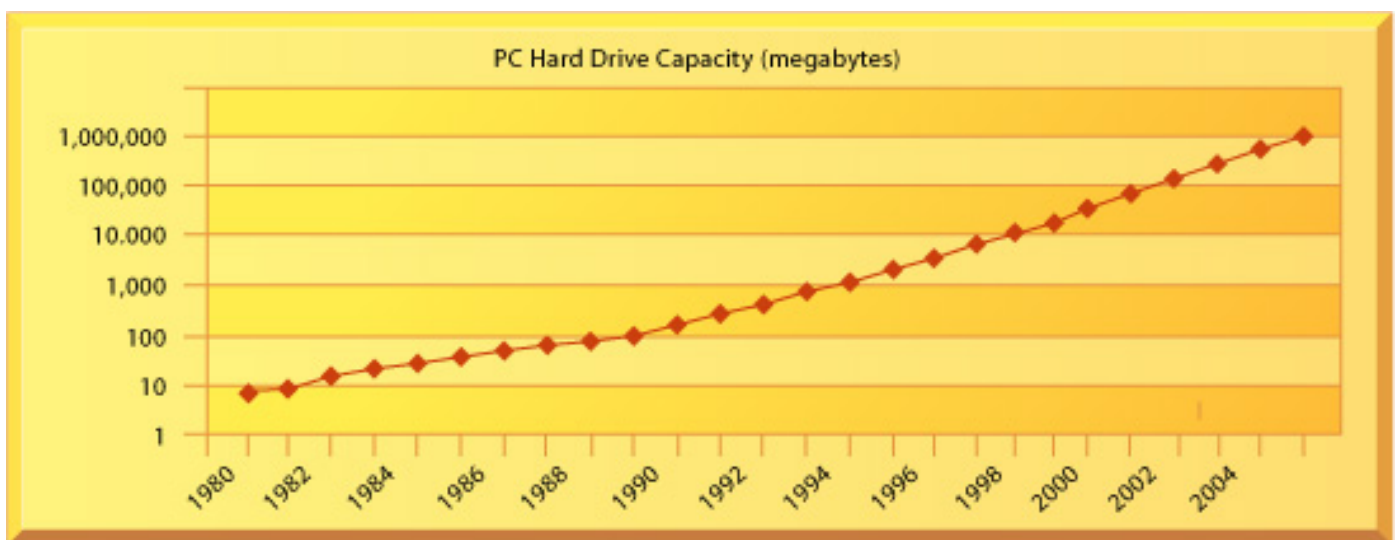


FIGURE 6-8 The capacity of hard disk drives grows exponentially, 1980–2004

From 1980 to 1990, hard disk drive capacities for PCs grew at the rate of 25 percent annual compound growth, but after 1990 growth accelerated to more than 65 percent each year.

Source: Authors.

Figure 6-9 shows that the number of kilobytes that can be stored on magnetic disks for one dollar from 1950 to 2004 roughly doubled every 15 months.

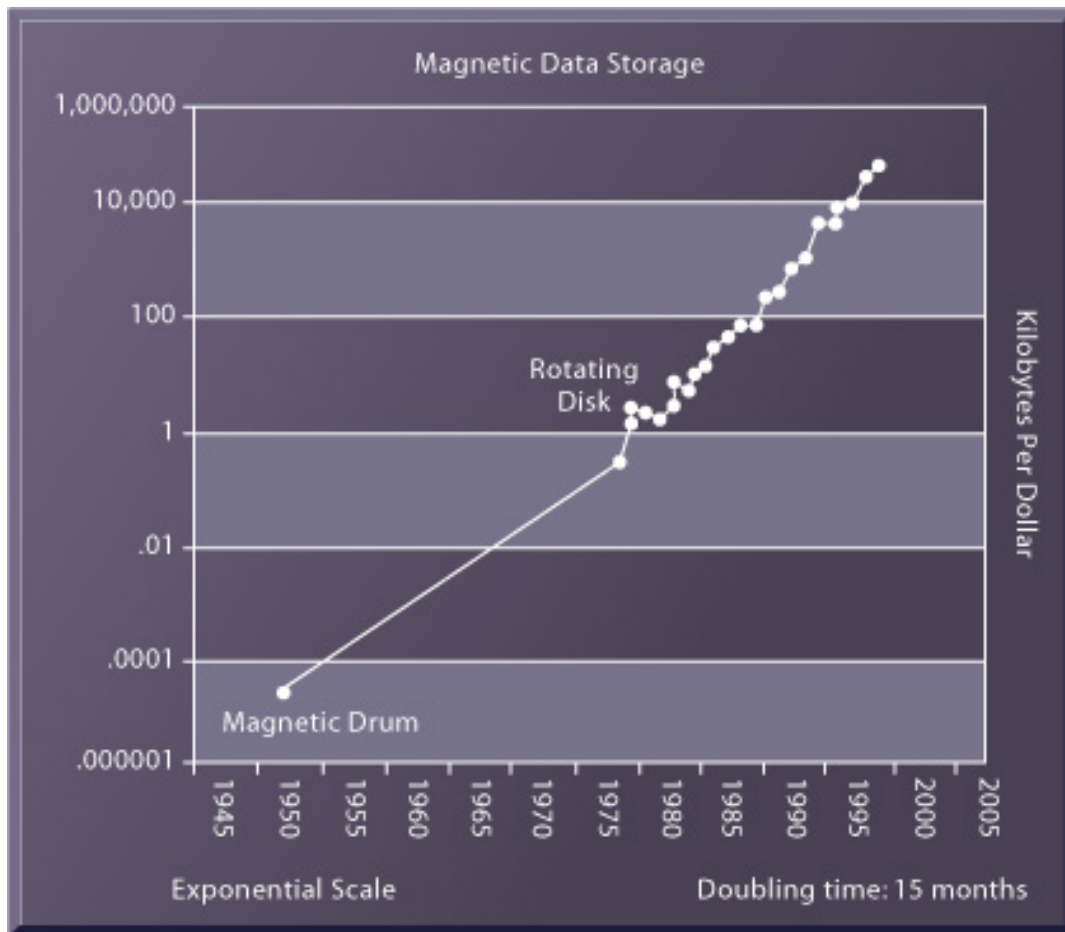


FIGURE 6-9 The cost of storing data declines exponentially, 1950–2004

Since the first magnetic storage device was used in 1955, the cost of storing a kilobyte of data has fallen exponentially, doubling the amount of digital storage for each dollar expended every 15 months on average.

Source: Kurzweill, 2003.

METCALFE'S LAW AND NETWORK ECONOMICS

Moore's Law and the Law of Mass Storage help us understand why computing resources are now so readily available. But why do people want more computing and storage power? The economics of networks and the growth of the Internet provide some answers.

Robert Metcalfe—inventor of Ethernet local area network technology—claimed in 1970 that the value or power of a network grows exponentially as a function of the number of network members. Metcalfe and others point to the increasing returns to scale that network members receive as more and more people join the network. As the number of members in a network grows linearly, the value of the entire system grows exponentially and theoretically continues to grow forever as members increase. Demand for information technology has been driven by the social and business value of digital networks, which rapidly multiply the number of actual and potential links among network members. We discuss Metcalfe's Law in greater detail in Chapter 8.

DECLINING COMMUNICATIONS COSTS AND THE INTERNET

A fourth technology driver transforming IT infrastructure is the rapid decline in the costs of communication and the exponential growth in the size of the Internet. An estimated 1 billion people worldwide now have Internet access, and over 250 million Web host computers exist in the United States. Figure 6-10 illustrates the exponentially declining cost of communication both over the Internet and over telephone networks (which increasingly are based on the Internet). As communication costs fall toward a very small number and approach 0, utilization of communication and computing facilities explodes.




FIGURE 6-10 Exponential declines in Internet communications costs

One reason for the growth in the Internet population is the rapid decline in Internet connection and overall communication costs. The cost per kilobit of Internet access has fallen exponentially since 1995. Digital Subscriber Line (DSL) and cable modems now deliver a kilobit of communication for a retail price of around 2 cents.

Source: Authors.

To take advantage of the business value associated with the Internet, firms must greatly expand their Internet connections, including wireless connectivity, and greatly expand the power of their client/server networks, desktop clients, and mobile computing devices. There is every reason to believe these trends will continue.

STANDARDS AND NETWORK EFFECTS

Today's enterprise infrastructure and Internet computing would be impossible—both now and in the future—without agreements among manufacturers and widespread consumer acceptance of technology standards. Technology standards are specifications that establish the compatibility of products and the ability to communicate in a network (Stango, 2004).

Technology standards unleash powerful economies of scale and result in price declines as manufacturers focus on the products built to a single standard. Without these economies of scale, computing of any sort would be far more expensive than is currently the case. Table 6-2 describes some important standards that have shaped IT infrastructure.

TABLE 6-2 Some Important Standards in Computing

Standard	Significance
American Standard Code for Information Interchange (ASCII) (1958)	Made it possible for computer machines from different manufacturers to exchange data; later used as the universal language linking input and output devices such as keyboards and mice to computers. Adopted by the American National Standards Institute in 1963.
Common Business Oriented Language (COBOL) (1959)	An easy-to-use software language that greatly expanded the ability of programmers to write business-related programs and reduced the cost of software. Sponsored by the Defense Department in 1959.
Unix (1969–1975)	A powerful multitasking, multiuser, portable operating system initially developed at Bell Labs (1969) and later released for use by others (1975). It operates on a wide variety of computers from different manufacturers. Adopted by Sun, IBM, HP, and others in the 1980s and became the most widely used enterprise-level operating system.
Transmission Control Protocol/Internet Protocol (TCP/IP) (1974)	Suite of communications protocols and a common addressing scheme that enables millions of computers to connect together in one giant global network (the Internet). Later, it was used as the default networking protocol suite for local area networks and intranets. Developed in the early 1970s for the U.S. Department of Defense.
Ethernet (1973)	A network standard for connecting desktop computers into local area networks that enabled the widespread adoption of client/server computing and local area networks and further stimulated the adoption of personal computers.
IBM/Microsoft/Intel Personal Computer (1981)	The standard Wintel design for personal desktop computing based on standard Intel processors and other standard devices, Microsoft DOS, and later Windows software. The emergence of this standard, low-cost product laid the foundation for a 25-year period of explosive growth in computing throughout all organizations around the globe. Today, more than 1 billion PCs power business and government activities every day.
World Wide Web (1989–1993)	Standards for storing, retrieving, formatting, and displaying information as a worldwide web of electronic pages incorporating text, graphics, audio, and video enables the creation of a global repository of billions of Web pages by 2004.

Beginning in the 1990s, corporations started moving toward standard computing and communications platforms. The Wintel PC with the Windows operating system and Microsoft Office desktop productivity applications became the standard desktop and mobile client computing platform. Widespread adoption of Unix as the enterprise server operating system of choice made possible the replacement of proprietary and expensive mainframe infrastructure. In telecommunications, the Ethernet standard enabled PCs to connect together in small local area networks (LANs; see Chapter 8), and the TCP/IP standard enabled these LANs to be connected into firmwide networks, and ultimately, to the Internet.

