PC Server 500

User's Reference

IBM

PC Server 500

User's Reference

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page 100.

Second Edition (September 1994)

The following paragraph does not apply to the United Kingdom or any country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

It is possible that this publication may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Requests for technical information about IBM products should be made to your IBM authorized reseller or IBM marketing representative.

Contents

About This Book

This User's Reference is an overview of the technologies and features that make up your server. It is intended for anyone who has an interest in learning more about the IBM PC Server 500.

This reference does not provide specific instructions for performing procedures. That information is in your *User's Handbook.*

How This Book is Organized

Chapter 1 is an overview of the features and expansion capabilities of your server. This chapter also includes the PC Server 500 specifications.

Chapter 2 provides information about server performance and the advantages of using a Local Area Network (LAN).

Chapter 3 contains basic information about microprocessors, and the types of microprocessors and memory in the PC Server 500 models.

Chapter 4 contains information about the Micro Channel design, including the Micro Channel bus, direct memory access (DMA), arbitration, burst data transfers, and the streaming-data procedure in your server.

Chapter 5 explains the PC Server 500 data storage capabilities and includes information about the direct access storage devices (DASD) that you can install, such as diskette drives, hard disk drives, CD-ROM drives, and tape drives.

Chapter 6 contains information about the disk array subsystem shipped in some models of PC Server 500. It describes the different RAID (redundant array of inexpensive disks) levels as well as the advantages of each.

Chapter 7 provides information about the configuration process. It also contains a description of all of the configuration programs provided with your PC Server 500.

Chapter 8 provides basic information about the serial, parallel, keyboard, mouse, video, and SCSI-2 input/output (I/O) connectors on the PC Server 500. Detailed pin assignments for each I/O connector are provided to help you determine if a device is compatible with your server.

Chapter 9 contains information about the super video graphics array (SVGA) adapter and video device drivers. This chapter also includes detailed information about some of the displays that you can attach to your server.

Chapter 10 explains about securing network hardware, securing data on the network, and the selectable-drive startup feature.

Appendix A contains your product warranty and special notices, such as a list of trademarks used in this book.

If you find a term you're not familiar with, refer to the glossary located in the back of this book. An index is also provided.

Related Publications

The *User's Handbook* provided with this manual contains detailed information about setting up, configuring, using, installing options in, and troubleshooting your server. See the “About This Book” section in the front of the User's Handbook for a description of what that book contains.

The following PC Server 500 publications are available for purchase:

The *PC Server 500 Hardware Maintenance Manual Supplement* contains a parts catalog, error codes, and advanced diagnostic procedures. This manual is intended for trained service technicians. (Diagnostic diskettes are not included.)

The *PC Server 500 Technical Reference* provides an in-depth register-level description of the server. This information is intended for programmers as well as application and adapter developers.

For a complete listing of publications available in the U.S. and Puerto Rico, call 1-800-426-7282. In Canada, call Customer Assistance at 1-800-465-1234. In all other countries, contact your IBM authorized reseller or marketing representative.

Chapter 1. Introducing Your IBM PC Server 500

The IBM* PC Server 500 is designed to meet your needs today, and provide plenty of expansion capabilities for the future. Your system is designed to be a high-performance server, where microprocessor performance, memory, system expansion, flexibility, and large amounts of data storage are high priorities.

Your PC Server 500 has been designed with performance, ease of use, and reliability in mind. Its advantages include *hot swapping* (the ability to install and remove hard disk drives without server downtime) and File Failure Prevention with NetFinity* support from the ServerGuide* package. Your server also takes advantage of the newest Pentium** microprocessor. As always, this IBM server meets stringent world wide certifications for power, EMC (electromagnetic compatibility), and safety.

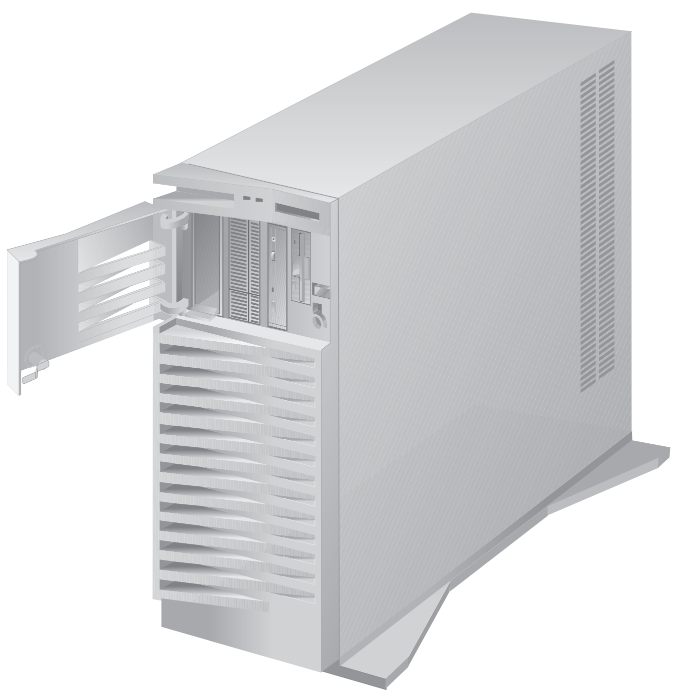

* Trademark of the IBM Corporation.

** Trademark of the Intel Corporation.

Features at a Glance

Microprocessor

Varies by model.

Memory

Standard: 32MB ECC expandable to 256MB
70 ns ECC
8 single inline memory-module sockets

Diskette Drives

Standard: One 3.5-inch, 2.88MB
Optional (internal):
- 3.5-inch, 2.88MB
- 3.5-inch, 1.44MB
- 5.25-inch, 1.2MB

CD-ROM Drive:

Standard: Double Speed Slim CD-ROM SCSI-2

Keyboard

Standard: 101-key keyboard

Mouse

Standard: IBM Mouse

Expansion

Eight 32-bit Micro Channel* slots
22 drive bays

Video

SVGA adapter
Compatibility:
- Video graphics array (VGA)
- Enhanced graphics adapter (EGA)
- Color graphics adapter (CGA)
- Multicolor graphics array (MCGA)

Hard Disk Drives

Number of drives and capacity vary by model
Hot-swap and plug-in capabilities
Twin-tailing
Easy front access
SCSI-2 interface

Security Features

Door lock
LogicLock with active security
U-bolt enabled

Upgradable POST and BIOS

Two, 256KB flash ROMs on the processor board
POST/BIOS upgrades (when available)

Information Panel

Displays diagnostic and status information

Integrated Functions

Vital product data (VPD)
LED usability support
Video port on an adapter
Two serial ports
Two parallel ports
Mouse port
Keyboard port
Battery-backed clock and time/date calendar

Power Supply

434 watt with auto-range voltage selection (115–230 V ac)
Built-in overload and surge protection
Power supply upgrade expansion option
- 220 watt automatic range voltage selection add-on
- Built-in overload and surge protection

SCSI-2 Controller (varies by model)

IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A
IBM SCSI-2 Fast/Wide Adapter/A

* Trademark of the IBM Corporation.

What Your IBM PC Server 500 Offers

The IBM PC Server 500 combines impressive performance using the Pentium microprocessor, large data-storage capabilities, enhanced system expandability, and SVGA video graphics. This system is designed to be used as a powerful, cost-effective server. Using the Micro Channel architecture, the PC Server 500 provides compatibility with a wide range of existing hardware devices and software applications.

Micro Channel Bus

Your PC Server 500 uses the Micro Channel architecture, which defines the size, physical arrangement, and logical relationships of the connectors. Micro Channel architecture makes the server versatile and expandable, allowing for a wide range of design and application alternatives. This server takes advantage of advancements in microprocessor speed, memory capacity, application programs, communication devices, networks, and operating environments.

Micro Channel architecture supports burst data transfers, in which data is sent in multiple bytes with no intervention by the microprocessor. This improves system performance and allows faster data transfers between devices than in non-Micro Channel computers.

Controls and Status Indicators

The most commonly used indicators on the front of the server are shown here.

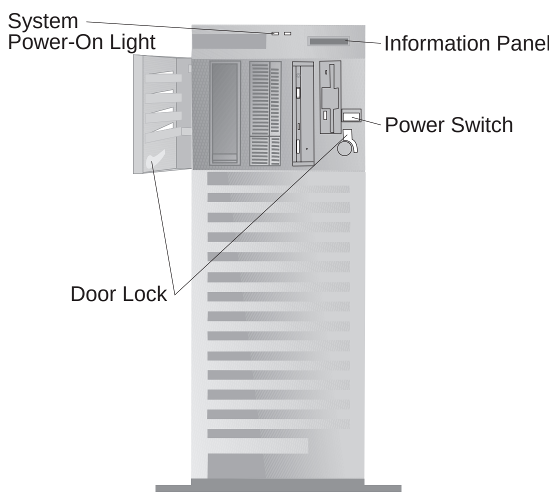

Power Switch: Pressing this switch turns your server on or off.

Information Panel: Diagnostic data and other system information appears here.

Power-On Light: This light is on when your server is turned on.

Door Lock: You can lock the door on your server to deter tampering with the internal components.

The most commonly used drive controls and indicators on the front of the server are shown here.

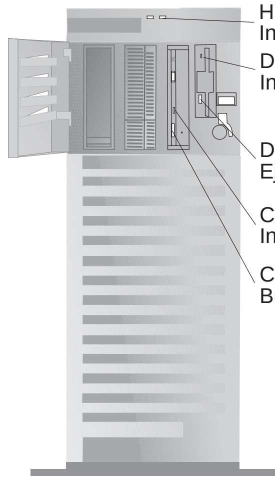

Hard Disk Drive In-Use Light: Normally, this light is on when the server is writing to or reading from the hard disk. However, on servers with a disk array subsystem, this light is operating-system dependent.

Diskette Drive In-Use Light: This light is on whenever the diskette drive is accessed.

Diskette Eject Button: Pushing this button releases the diskette from the drive.

CD-ROM Drive In-Use Light: This light is on whenever the CD-ROM drive is accessed.

CD-ROM Eject Button: This button pushes the CD-ROM tray out from the server so that you can insert or remove a CD.

Expansion Bays

Your server comes with one 3.5-inch 2.88MB diskette drive, one CD-ROM drive, and at least one hard disk drive.

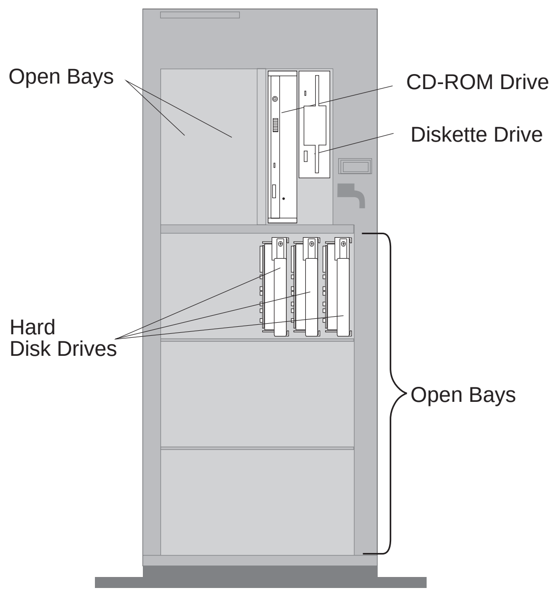

Diskette Drive: The 3.5-inch, 2.88MB diskette drive uses 1MB, 2MB, and 4MB diskettes.

CD-ROM Drive: The 5.25-inch CD-ROM drive is located next to the diskette drive. For information about inserting CDs and using the CD-ROM drive, see the *User's Handbook.*

Open Bays: Open drive bays make it possible for you to have as many as 18 internal hard disk drives in your server. Another open bay next to your CD-ROM drive allows you to install a tape drive, rewritable optical drive, or another hard disk drive. For installation instructions and information on the types of drives that you can install in each bay, see the “Installing Options” chapter in your *User's Handbook.*

Hard Disk Drives: The number of drives installed in your server and their capacity depends upon your model.

Expansion Slots

Your server has eight Micro Channel expansion slots. The video adapter and a SCSI-2 adapter are installed in two of the slots.

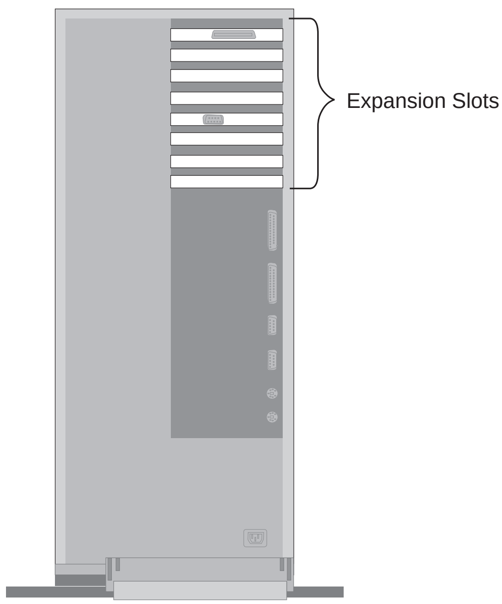

The remaining slots are available for future expansion and growth. For example, you can add adapters to provide communications, specialized graphics, and sound. Many adapters provide bus-master capabilities, which enable the adapters to perform operations without interrupting the system microprocessor.

These Micro Channel slots have Plug and Play capabilities. When operating systems that support this feature become available, your server automatically will configure itself when plug and play peripherals are installed.

Input/Output Ports

The input/output connectors (ports) are shown here.

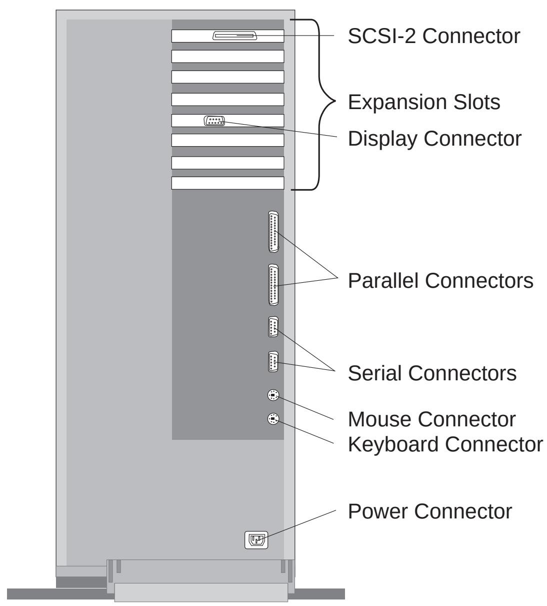

SCSI-2 Connector: Attach external SCSI devices here.

Display Connector: The display signal cable attaches to the connector on this adapter.

Parallel Connectors: Two 25-pin parallel ports are provided. This is where the signal cable for a parallel printer, or other parallel device, connects to your server. (Parallel Port A is a high speed port.)

Serial Connectors: Two 9-pin serial connectors (A and B) are provided. The serial signal cable for a modem or other serial device usually connects here. If you are using a 25-pin signal cable, you need a 9-pin-to-25-pin adapter cable.

Mouse Connector: This is where the mouse cable connects to the server. This port sometimes is called an auxiliary-device or pointing-device port.

Keyboard Connector: The cable from your keyboard connects here.

International Capabilities

The IBM PC Server 500 is suitable for use worldwide. Electric power can vary from country to country; however, the voltage-sensing feature allows the server to automatically set itself to match differing power sources.

Language is another factor to consider. Your system can be used with special editions of operating systems that let the system and keyboard operate with other languages. With the appropriate adapter, those same operating systems also let you view Asian characters on your display. Now you can use one system to communicate in several languages.

Specifications

Size

Depth: 755 mm (29.7 in.)
Depth with cable cover: 835 mm (32.8 in.)
Height: 622 mm (24.5 in.)
Width: 353 mm (13.9 in.)

Weight

Configured with one hard disk drive: 31.29 kg (69 lb)

Environment

Air temperature:
- System on: 10° to 35° C (50° to 95° F)
  Altitude: 0 to 914 m (3000 ft.)
- System on: 10° to 32° C (50° to 90° F)
  Altitude: 914 m (3000 ft.) to 2133 m (7000 ft.)
- System off: 10° to 43° C (50° to 110° F)
  Maximum Altitude: 2133 m (7000 ft.)
- Storage: 1° to 60° C (39° to 140° F)

Humidity:
- System on: 8% to 80%
- System off: 8% to 80%

Maximum altitude: 2133 m (7000 ft)

Heat Output

Approximate heat output in British Thermal Units (BTU) per hour:
- Minimum configuration: 150 BTU (44 watts)
- Maximum configuration: 3,600 BTU (1055 watts)

Electrical Input

Sine-wave input (50 to 60 Hz) is required
Input voltage:
- Low range:
  - Minimum: 90 V ac
  - Maximum: 137 V ac
- High range:
  - Minimum: 180 V ac
  - Maximum: 265 V ac
- Input kilovolt-amperes (kVA) approximately:
  - Minimum configuration as shipped: 0.16 kVA
  - Maximum configuration: 1.0 kVA

Total Power Available for Drives

Nominal Operating Current allowed: (Base with expansion upgrade)
- +5 V dc line: 10 A base/23 A with upgrade
- +12 V dc line: 10 A base/23 A with upgrade

Acoustical Noise Emission Values

Average sound pressure levels at operator position:
- 1 hard disk drive (low fan speed):
  - 40 dB operating
  - 38 dB idle
- 1 hard disk drive (high fan speed):
  - 42 dB operating
  - 41 dB idle
- 3 hard disk drives (low fan speed):
  - 40 dB operating
  - 38 dB idle
- 3 hard disk drives (high fan speed):
  - 42 dB operating
  - 42 dB idle

Average sound pressure at bystander position (1 meter):
- 1 hard disk drive (low fan speed):
  - 36 dB operating
  - 35 dB idle
- 1 hard disk drive (high fan speed):
  - 38 dB operating
  - 38 dB idle
- 3 hard disk drives (low fan speed):
  - 36 dB operating
  - 35 dB idle
- 3 hard disk drives (high fan speed):
  - 38 dB operating
  - 38 dB idle

Declared (upper limit) sound power levels:
- 1 hard disk drive (low fan speed):
  - 5.5 bels operating
  - 5.4 bels idle
- 1 hard disk drive (high fan speed):
  - 5.7 bels operating
  - 5.6 bels idle
- 3 hard disk drives (low fan speed):
  - 5.5 bels operating
  - 5.4 bels idle
- 3 hard disk drives (high fan speed):
  - 5.7 bels operating
  - 5.7 bels idle

Note: One or more of the cooling fans in this system will operate at either high or low speed, depending on the hardware configuration of the server. Noise levels for both are included in the table.

Acoustical levels are measured in controlled acoustical environments according to the American National Standards Institute (ANSI) procedure S12.10 and ISO 7779, and are reported in accordance with ISO 9296. Sound pressure levels in your location might exceed the average values stated because of room reflections and other nearby noise. The declared sound power levels indicate an upper limit, below which a large proportion of machines operate.

RAS Features

Three of the most important factors in server design include *reliability*, *availability*, and *serviceability* (RAS). Your server has a number of features to ensure it is reliable, available, and serviceable.

Reliability

ECC memory corrects all the single-bit memory errors and detects all of the dual-bit memory errors.

Availability

Quick and easy access from the processor complex to the information panel has been achieved through the design and use of a serial link. This enables the processor complex to communicate error code information to the information panel even if the video adapter, Micro Channel bus, and other system board logic is nonfunctional.

Serviceability

Vital product data is stored in nonvolatile memory. This data includes serial number information and replacement part numbers. This information makes remote maintenance of your server easier and more efficient.

Chapter 2. Understanding Server Performance

Your productivity and satisfaction, as well as everyone else's on your network, are dependent upon server performance. A number of factors can influence the performance of your server, such as the configuration of the server hardware, application workloads, server operating-system configuration parameters, and the server operating environment.

This chapter provides information to help you avoid potential performance bottlenecks, and suggestions for how you can improve your server's performance.

One of the main capabilities of a local area network is sharing of resources, such as data and expensive peripheral devices. The ability to share resources can mean a decrease in the cost of an individual workstation, because every workstation might not need its own software, printer, or hard disk. Sharing resources also can result in improved work flow, quality of work, and productivity. Some reasons for this are:

- Centralized databases can be shared easily.
- Security of information is more easily ensured.
- Meeting notices, notes, messages, letters, and text files can be sent to multiple users simultaneously, using electronic mail.

SynchroStream Feature

Your PC Server 500 includes an IBM SynchroStream controller, which is an advanced processor complex technology. The SynchroStream controller synchronizes large amounts of data flowing among the Micro Channel adapters, the microprocessor cache controller, and the system memory. The processor complex processes data by allowing it to stream in parallel at full bandwidth to each subsystem. This means that data never has to stream through an intermediate subsystem. The SynchroStream technology provides a higher level of server throughput and performance than servers that don't use this technology.

Integrated Systems Management

The ServerGuide package that came with your server gives you the option of installing IBM NetFinity. NetFinity is a LAN management program that enhances the capabilities of your LAN. To find out more about NetFinity and how it can help you, view the demo and read the information in the ServerGuide package.

Performance Tuning

There is a performance tuning feature available on the ServerGuide CD (provided with your server). You can easily customize the programmable registers of your system so that your applications run faster than before. To improve the performance of your server, select Performance Tuning from the ServerGuide main menu and follow the instructions on the screen.

Server Performance

You can optimize the performance of your server by balancing the performance of its main components. These components include the:

- Network program
- Application programs
- Network adapter
- System memory
- System microprocessor
- Hard disk subsystem (disk controller and drives)

The Network Program

Always check the documentation that comes with the network program to determine which parameters affect performance and which, if any, parameters you can change to increase performance. The original parameter settings that are provided by most network programs usually support a network with 12 to 18 concurrent users. When you have a larger network or increase the size of a network, you will probably need to adjust some of these settings to maintain adequate performance.

It is a good practice to record the original settings before you make any changes. In the unlikely event that performance decreases, you can change the settings back to their original values. It is also a good practice to document changes by making remarks, such as the date and reason for the changes, above or below the values that you change. This type of information can save you valuable time if you need to make additional adjustments.

Application Programs

Most complex programs allow you to change specific parameters to improve their performance in a network environment. Use the instructions that come with the application programs to determine which parameters to change and when and how to change them to affect performance.

For the same reasons as mentioned previously in "The Network Program" on page 14, it is a good practice to record the original settings before you make any changes and, if possible, document the changes by making remarks above or below the values that you change.

The Network Adapter

Before you can use your server on a network, your server must have a network adapter installed. The primary value of using a bus-master network adapter is added capacity or, in the case of a heavily loaded server, improved performance. When you use a bus-master network adapter, the network throughput is less dependent on the type of system microprocessor and more dependent on the network adapter design and the network software.

The appropriate network adapter for your server depends on the type of network you have (broadband, baseband, token-ring, or Ethernet).

IBM and several other companies offer baseband, broadband, token-ring, and Ethernet network adapters for Micro Channel computers. These adapters operate in all Micro Channel computers, but some of them, such as the bus-master network adapters, are more suited for use in network servers.

Bus-master network adapters provide a higher throughput rate, because they handle most of the network processing and leave the system microprocessor free for other tasks.

Adding a faster network adapter provides additional throughput. Therefore, the total number of transactions that the server processes per second increases and the application response time decreases. But this reduction in response time is a limited benefit. As you add more users, the effects of the faster network adapter are offset by the increase in data space access and write-data processing (which reduces the disk-cache-hit rate and overall server performance).

For a current list of the network adapters that are available for your model, contact your authorized reseller or marketing representative.

System Memory

When you have a large network or you increase the size of a small network, you might need to install additional memory to maintain adequate performance. This additional memory might be needed to support additional application programs, buffers, or control blocks.

The System Microprocessor

One way to improve system microprocessor performance is to lighten its load by using bus masters. Bus masters can handle tasks, such as the network processing or data transfers, and provide relief for an overloaded system microprocessor. (For additional information about bus masters, see “Bus Masters” on page 32.)

The Hard Disk Subsystem

An IBM SCSI-2 bus-master controller is a standard feature in your server. This type of controller significantly increases performance by handling the processing of SCSI commands and data transfers, leaving the system microprocessor free to handle other tasks.

You can attach up to 15 SCSI devices to the existing SCSI controller to increase disk storage and performance. If those are not enough, you can install additional SCSI controllers for further expansion. You also can use SCSI storage enclosures to house additional storage devices.

Disk arrays use multiple hard disk drives that provide faster input and output rates than a single large-capacity drive, because the files are stored on individually addressable disks and can be found more quickly during read requests. The amount of performance improvement depends on the application programs that are being run on the server and the type of array that you use. (See Chapter 6, “Disk-Array Subsystems” on page 49 for more information about disk-array technology.)

Estimating Server Performance and Capacity

There are no simple methods of evaluating your server performance for all environments. The hardware, software, applications, workloads, and number of users all affect server performance.

When you are ready to do a performance evaluation, you must always use a systematic approach. A good example of this type of approach is:

- Define the functional requirements of your PC Server 500 for your working environment.
- Define what you consider to be adequate performance, such as response times, transaction rates, and number of users.
- Understand what factors affect server performance.
- Decide what an approximate configuration of applications for your work environment would be.
- Measure the current performance of your server using a variety of application configurations, then use the performance results to locate performance bottlenecks and to determine which configuration provides sufficient performance today, and allows room for growth.
- Ensure that the performance of the server will still be acceptable if one of the components fails.

Improving Server Performance

Before you begin to experiment with applying upgrades or enhancements to improve performance, remember that the subsystem that causes a bottleneck depends largely on the applications that you are using.

Try to determine which applications most closely resemble your environment, then concentrate on obtaining improvements for that application. It is possible that some changes might improve the performance of your application, but degrade the performance for other applications. If this occurs, you will need to assess the trade-offs.

When you understand the factors that affect server performance, you can then use a methodical approach to improving it. A good example of this type of approach is:

1. Measure the current server performance.
2. Identify the performance bottlenecks.
3. Upgrade components that cause the bottlenecks.
4. Measure the new performance of the server.

For example, if the bottleneck is memory, you could install additional memory; if it's the hard disk subsystem, then you could install an additional drive or create a disk array. You can replace slow network adapters with faster ones, or you can install additional adapters, and so on.

Chapter 3. Processor Complex

The *processor complex* consists of the devices and features in the server that perform logical operations and calculations, control access to memory, and manage data-transfer operations. The following devices and features make up the processor complex:

- The microprocessor
- The memory subsystem
- Cache memory
- The direct memory access (DMA) controller

Some of the components of the processor complex are on a *processor board,* which you can remove and replace with a different one to modify the performance characteristics of the server. The processor board connects to two 164-pin connectors on the system board. These connectors provide:

- The Micro Channel interface, which allows data to be transferred between the processor complex and the adapters in the Micro Channel expansion slots.
- The system board interface, which allows the transfer of data between the processor complex and devices on the system board, such as the parallel, serial, keyboard, and auxiliary-device ports.
- Two memory interfaces, which the processor complex uses to read from and write to system memory (for an explanation of system memory, see "Physical Memory and Virtual Memory" on page 23). All access to system memory is through the memory controller in the processor complex.

Microprocessor

The *microprocessor* is an integrated circuit that performs most of the control and computing functions of your server. A Pentium is installed in your PC Server 500.

Clock speed is the rate at which the microprocessor performs instructions. The design of the microprocessor determines the maximum clock speed at which it can operate reliably, and the design of the server determines the optimal clock speed for the server. A quartz crystal on the system board or processor board generates a pulse to the microprocessor, making the microprocessor operate at a specific speed. Depending on the design of the server, the optimal clock speed might not be the maximum speed for which the microprocessor is rated.

Clock speed is measured in megahertz (MHz). A *clock cycle* is the time unit that is used to measure how long the microprocessor takes to perform instructions at a given clock speed; clock cycles are measured in nanoseconds (ns) and are the inverse of the clock speed. For example, at a clock speed of 16 MHz, one clock cycle takes 62.5 ns; at a clock speed of 90 MHz, one clock cycle takes 11 ns. The microprocessor in your server operates at two different clock speeds: one speed for transferring data into and out of the microprocessor and a faster speed for operations within the microprocessor.

The microprocessor needs data and instructions for each processing operation that it performs. Data and instructions are copied from memory into data-storage locations, known as *registers,* in the microprocessor. Registers are also used to store the data that results from each processing operation, until the data is transferred to memory.

The microprocessor *data interface* is the connection between the microprocessor and the data bus. Your server uses a microprocessor whose data-interface width matches the width of the data bus in the server, and the width of the data bus determines which type of adapters (16-bit or 32-bit) the server supports. Your server has a 32-bit data bus, and it uses a microprocessor that has a 64-bit data interface, which supports both 16-bit and 32-bit adapters.

The set of instructions that the microprocessor can perform determines whether a particular program will run in the server. Programs that are written for a 32-bit interface require only that the microprocessor be capable of decoding and performing 32-bit instructions. The width of the data bus does not affect software compatibility.

The size of the microprocessor *address interface,* which is the connection between the microprocessor and the address bus, determines the width of the address bus and the amount of physical memory the microprocessor can address. A microprocessor with a 24-bit address interface can address a maximum of 16MB (MB equals 1 048 576 bytes) of physical memory. However, the microprocessor in your server has a 32-bit address interface and can address up to 4GB (GB equals 1 073 741 824) of physical memory. (Physical memory is described on page 23.)

Real Mode and Protected Mode

The microprocessor has a *real mode,* in which the microprocessor operates exactly like the 8088 microprocessor (including the 1MB addressing limitation). The real mode does not support virtual memory (see "Physical Memory and Virtual Memory" on page 23). When the microprocessor is in the real mode, your server can run programs that were written for the IBM personal computer. The real mode is the default condition for the microprocessor; that is, the microprocessor is always in the real mode after the server is turned on and before any operating system is started.

The microprocessor also has a *protected mode,* which supports virtual memory and allows the microprocessor to address more than 1MB of memory. The protected mode is so named because, in a multitasking environment, the microprocessor manages the memory so that each program is protected from interference from other programs running at the same time.

The operating system that you install will determine whether the microprocessor stays in the real mode or is switched to the protected mode. For example, when the server runs DOS, the microprocessor stays in the real mode. DOS makes only limited use of the capabilities that are provided by the protected mode. When the IBM Operating System/2 (OS/2) operating system is started, it switches the microprocessor to the protected mode. The OS/2 operating system takes advantage of the virtual memory and multitasking capabilities provided by the protected mode of the microprocessor.

Note: OS/2 2.0 has a real mode and a protect mode. These operating-system modes are not related to the real mode and protected mode of the microprocessor.

Math Coprocessor

A *math coprocessor* is a specialized type of microprocessor that performs only mathematical operations. The microprocessor passes data and instructions to the math coprocessor, which performs the mathematical operations and returns the results to the microprocessor. The microprocessor in your server has a built-in math coprocessor. Programs that require extensive mathematical calculations take much less processing time when the math coprocessor performs the calculations, because it is designed specifically for that task, and the system microprocessor can perform other functions at the same time. A program that is intended to use the math coprocessor must contain specific instructions that are recognized by the math coprocessor. Otherwise, the program will not activate the math coprocessor, and the microprocessor will handle the calculations.

Memory Subsystem

The *memory subsystem* consists of random access memory and read-only memory, physical memory and virtual memory, the memory controller, and cache memory.

Random Access Memory and Read-Only Memory

Random access memory (RAM) is used for temporary storage of data and instructions during processing. RAM is contained in memory components known as *single inline memory modules* (sometimes referred to as SIMMS). It is *volatile* memory, which means that to retain its contents, it must be constantly refreshed by an electrical current. While the server is turned off, no current is supplied to the single inline memory modules, so no data is retained in RAM.

The amount of RAM that a server supports is determined by the server model and by the operating system that you use. Your server supports 256MB, (70 ns memory module kits), of error correcting code (ECC) memory. You can add memory to the standard configuration of your server by installing them in matched pairs.

Read-only memory (ROM) is used for storage of programs that the server uses for startup procedures and other internal operations. These programs are permanently encoded in an integrated circuit, known as a *ROM module*, on the processor complex. ROM is *nonvolatile* memory, which means that it retains its contents when the server is turned off. Generally, the contents of a ROM module cannot be modified. However, the type of ROM modules in your server can be reprogrammed. *Electrically erasable programmable ROM (EEPROM)* modules (sometimes referred to as *flash memory*) on the processor complex can be reprogrammed while they are in the server. For example, the POST and BIOS code in your server is stored in *flash memory* so that it can be updated whenever enhancements are made. If you have a disk-array model, there is also an EEPROM on the SCSI-2 RAID adapter that stores the disk array configuration information.

In addition to RAM and ROM, the system board has a small area of *nonvolatile RAM*. Using a small current from the system battery, nonvolatile RAM retains its contents while the server is turned off. Nonvolatile RAM is used to store the power-on password, the time and date, and system configuration information.

Physical Memory and Virtual Memory

Physical memory consists of all the writable memory locations in the server, most of which are on single inline memory modules. Portions of the physical memory are reserved for the operating system, for the power-on self-test, and for the video data that is being displayed. *System memory* is the part of physical memory that is available for instructions and data that the server uses to run programs.

Virtual memory is memory that appears to be allocated to application programs. The operating system uses a portion of the hard disk as virtual memory, swapping data and instructions between the hard disk and physical memory. Virtual memory makes multitasking possible. In a multitasking environment, the memory requirements of all the programs that might be running in the server at the same time can occasionally far exceed the amount of physical memory that is available. The operating system allocates virtual memory to meet the total memory requirements of each program and then manages the available physical memory to meet the actual requirements at each point in time. Thus, the amount of virtual memory that is allocated can be much greater than the amount of physical memory that is installed in the server.

The maximum amount of physical memory that you can install in your server, and the maximum amount of virtual memory that the operating system can allocate, are determined by the type of microprocessor in your server. Your server ships with 32MB of memory and you can install a maximum of 256MB. You can use 4-, 8-, 16-, or 32MB, 70ns, ECC memory-module kits when adding memory to your server.

The width of the microprocessor address interface limits the amount of physical and virtual memory that the microprocessor can address. The microprocessor in your server has a 32-bit address interface and can address a maximum of 4GB (2^{32} bytes) of physical memory and 64TB (TB equals approximately 1 000 000 000 000 bytes) of virtual memory, in the protected mode.

Memory Controller

The memory controller is a device on the processor board that controls access to system memory by the microprocessor and I/O devices. Registers in the memory controller contain information about the amount and type of memory that is installed in the server. During a system reset, the POST routine writes this information into the registers.

The functions of the memory controller vary and can include:

- Dual bus capability, which allows the microprocessor to read from and write to system memory while a bus master is controlling the Micro Channel bus. (For more information, see “Dual Bus” on page 33.)
- Memory timing control, which coordinates data-transfer operations involving single inline memory modules that operate at different speeds.
- Cache control, which ensures the validity of the contents of the cache. (For information about cache memory, see “Cache Memory” on page 26.)

- SynchroStream controller, which synchronizes large amounts of data flowing among the Micro Channel adapters, processor cache controller, and system memory.
- Bus-width allocation, which supports 8-, 16-, 32-, and 64-bit data-transfer operations.
- Memory interleaving, which is a method of reducing the time the microprocessor has to wait for system memory to respond during memory I/O operations.

Memory Interleaving and Noninterleaving

Some microprocessors are so fast that system memory cannot respond to read and write requests as quickly as the microprocessor can send them. In some models, system memory imposes one or more *wait states* on the microprocessor when it reads data from or writes data to system memory. A wait state is a period of time (one microprocessor clock cycle) during which the microprocessor suspends processing and waits for system memory to respond to a read or write operation (a *memory I/O operation*). Wait states cause the server to operate less efficiently than it would if the microprocessor were able to continue processing data during memory I/O operations.

Memory interleaving is a method of reducing the number of wait states that are imposed on the microprocessor during memory I/O operations.

In servers that use *noninterleaved memory*, the speed of memory I/O operations is limited by the speed at which system memory can respond. The speed of a memory I/O operation is measured in microprocessor clock cycles, so the microprocessor clock speed determines the minimum time that is required for a memory I/O operation.

To reduce the time that the microprocessor spends waiting for system memory to respond, some servers have *interleaved memory*. In these servers, system memory is divided into two banks: bank A and bank B. Data is stored sequentially in system memory so that the first 32 bits are stored in bank A, the second 32 bits in bank B, the third in bank A, and so forth. Each bank has a *data buffer*, which is a temporary holding area for data. When the microprocessor reads the first 64 bits of data from memory, 32 bits from bank A are sent to the microprocessor while bank B loads 32 bits into the bank B buffer. Then the data in the bank B buffer is sent to the microprocessor while bank A loads 32 bits into the bank A buffer. The microprocessor can be reading from one buffer while the other bank is completing its memory cycles for the next 32 bits of data. A similar process occurs when the microprocessor writes data to memory.

Cache Memory

Another method of reducing the need for wait states is the use of *cache memory,* which improves server performance by temporarily storing frequently used data and instructions in a *cache.* A cache is a small buffer between the microprocessor and system memory. The amount of internal cache memory is determined by the type of microprocessor installed. Your Pentium microprocessor has a 16KB (8KB of data and 8KB of instruction) internal level-1 cache. Cache contains high-speed memory, known as static random access memory (SRAM), that can respond to memory I/O operations without imposing wait states on the microprocessor.

Microprocessor cache memory is used to store the information most often used by the microprocessor. Internal cache allows a microprocessor to process information faster than if it had to use the system memory each time it needed new information. During processing, as requirements change, the cache controller copies other data and instructions into the cache, replacing data and instructions that are no longer needed in the cache. Server performance is improved each time the microprocessor finds what it needs in the cache (a *cache hit*). If it does not find what it needs in the cache (a *cache miss*), the cache controller must locate the data or instruction in system memory and copy it into the cache, while one or more wait states are imposed on the microprocessor. The cache controller manages the use of the cache so that the number of cache hits far exceeds the number of cache misses.

For maximum performance, PC Server 500 cache is set to the write-back mode; however, you can use the system programs to change the cache setting to the write-through mode.

Write-Back and Write-Through Cache

The PC Server 500 default cache setting is the write-back mode. The write-back mode provides the best performance, because the microprocessor updates the cache, then the cache controller updates system memory. In the write-back mode, the microprocessor updates the cache in 15 to 20 ns, then goes on to perform other functions. In the write-through mode, the microprocessor updates system memory directly. This mode is much slower, because the microprocessor must deal directly with system memory, which runs at 70 ns.

Level-2 Cache

Your server also has 256KB of parity protected *level-2 cache.* This 256KB cache feature increases the amount of cache memory, which increases the probability of cache hits. If the microprocessor does not find what it needs in the level-2 cache (a *second-level cache miss*), the cache controller locates the data or instruction in system memory and copies it into one of the caches.

Direct Memory Access Controller

Direct memory access (DMA) is a method of transferring data between system memory and I/O devices without requiring intervention by the microprocessor. The *DMA controller* is a component of the processor complex that manages all the DMA data transfers.

For an explanation of DMA and the DMA controller, see “Direct Memory Access” on page 30.

Chapter 4. Micro Channel Architecture

The Micro Channel *bus* is an electrical pathway for transferring information among the microprocessor, memory, adapters, and I/O devices. It has 32 signal paths for data, which means that 32 bits of data can be sent through the bus at the same time. Only one device at a time can use the Micro Channel bus to send or receive data. The bus has several functions. It can be used as an expansion bus, an address bus, or a data bus.

Expansion Bus

The expansion bus supports the addition of adapters to customize your server and add new function. The expansion bus has eight connectors for adapters. These connectors provide paths for transferring data to and from the adapters.

Address Bus and the Data Bus

When a device sends data, it identifies where the data is stored and where it is to be sent. Each location is identified by a unique number, known as an *address.* An address is assigned to each memory location and I/O device that is attached to the Micro Channel bus. The sending device uses the Micro Channel bus as an address bus to send the target-address information to the receiving device. Then the sending device uses the Micro Channel bus as a data bus to send the data.

The width of the address bus determines how many memory locations your server can support. Your server has a 32-bit address bus which allows up to 4GB of memory to be addressed.

The width of the data bus determines how much data can be transferred at one time. Micro Channel architecture provides data-bus widths of 8, 16, and 32 bits, depending on the capacity of the sending and receiving devices. For example, when a 32-bit device transfers data to an 8-bit device, the 32-bit device must send the data in four 8-bit transfers. Only when both the sending device and the receiving device are 32 bits wide is a 32-bit data bus required.

Direct Memory Access

Direct memory access (DMA) is a method of transferring data between system memory and I/O devices without requiring intervention by the microprocessor. DMA is more efficient than *programmed I/O,* in which the microprocessor reads the data from the sending device and then writes it to the receiving device. In DMA data transfers, data can bypass the system microprocessor as it moves between system memory and I/O devices. DMA improves server performance because the microprocessor does not have to interrupt its processing activities to manage data transfers.

The *DMA controller* is integrated into the processor board and manages all DMA data transfers. Transferring data between system memory and an I/O device requires two steps. Data goes from the sending device to the DMA controller and then to the receiving device. The microprocessor gives the DMA controller the location, destination, and amount of data that is to be transferred. Then the DMA controller transfers the data, allowing the microprocessor to continue with other processing tasks.

When a device needs to use the Micro Channel bus to send or receive data, it competes with all the other devices that are trying to gain control of the bus. This process is known as *arbitration*. The DMA controller does not arbitrate for control of the bus; instead, the I/O device that is sending or receiving data (the *DMA slave*) participates in arbitration. It is the DMA controller, however, that takes control of the bus when the central arbitration control point grants the DMA slave's request.

Arbitration

When multiple devices need to use the Micro Channel bus at the same time, these devices participate in *arbitration.* Every device that can arbitrate for control of the bus is assigned a priority level, known as an *arbitration level,* that is used to determine which device should control the bus next. The arbitration level for each device is contained in a software file, known as an *adapter description file (ADF).*

Note: Devices that contend for control of the Micro Channel bus at the same time must not have the same arbitration level. Occasionally, arbitration-level conflicts occur between devices. When this occurs, you can use the system programs to change the arbitration level of one of the conflicting devices.

Micro Channel architecture has a *fairness feature,* which ensures that each device gets a turn to control the bus, even if it has a low priority level. The fairness feature guarantees that none of the devices are locked out of the bus and that each device can gain control of the bus within a given amount of time. When you configure your server, you can disable the fairness feature for a device so that it controls the bus more than other devices. A device for which fairness has been disabled can monopolize the bus. Disabling the fairness feature for more than one device is risky. You could cause some devices for which the fairness feature is enabled to be completely locked out of the bus, including the microprocessor. Therefore, it is best to leave the fairness feature enabled for all devices.

The *central arbitration control point* is a location in the system master where contending devices send their arbitration signals. It does not actually decide which device should control the Micro Channel bus; the contending devices make that determination among themselves, using the arbitration logic that is programmed into the devices. However, it is the central arbitration control point that actually grants control after the decision is made.

Masters and Slaves

Two types of devices, *masters* and *slaves,* take part in the data-transfer process. One master and one slave participate in a data-transfer operation.

Masters

A master is a device that can own the Micro Channel bus. When a master owns the bus, it can send data to or receive data from a slave (a device, an adapter, or system memory) without interrupting the microprocessor. There are three types of masters: the system master, bus masters, and the DMA controller.

System Master

The *system master* assigns system resources, manages the system configuration, issues the commands of the primary operating system, and can grant control of the Micro Channel bus to a *bus master*.

Bus Masters

Your server supports up to 15 bus masters. *Bus masters* take control of the Micro Channel bus and transfer data directly to and from I/O devices and memory without requiring intervention by the system microprocessor or DMA controller.

A bus master can have its own microprocessor, instruction cache, and memory. By taking over some of the work of the system microprocessor, bus masters create a multiprocessing environment and increase overall system performance.

DMA Controller

The *DMA controller* manages data transfers between DMA slaves and memory slaves. This type of transfer is often called a third-party DMA operation. See “Direct Memory Access Controller” on page 27 and “Direct Memory Access” on page 30 for more information.

Slaves

A *slave* is a device that is selected by a controlling master as either the source or the target for a transfer. A slave can also begin a service request, such as an interrupt. There are three types of slaves: memory, I/O, and DMA.

Memory Slaves

A *memory slave* is a device that provides a block of system memory. Memory slaves respond to read and write operations by placing the requested data on the Micro Channel bus or by writing data from the bus to random access memory (RAM). A memory slave can be selected by any of the three types of Micro Channel masters.

I/O Slaves

An *I/O slave* is a device that communicates with or controls a separate piece of equipment, such as a printer. An I/O slave can be selected by the system master or by a bus master.

DMA Slaves

A *DMA slave* is the only type of slave that can initiate arbitration. DMA slaves require the DMA controller to manage data transfers. A DMA slave can be selected by the DMA controller or by a bus master.

Dual Bus

Your server has a *dual bus,* meaning that it has one data bus from the microprocessor to the memory controller and another data bus from the Micro Channel devices to the memory controller. This allows the microprocessor to read from and write to system memory while a bus master is controlling the Micro Channel bus.

The following list summarizes dual-bus operation:

- When the microprocessor is reading from or writing to its internal cache or to the optional 256KB (KB equals approximately 1000 bytes) cache, the bus master that is controlling the Micro Channel bus has exclusive access to system memory.
- The microprocessor and the bus master that is controlling the Micro Channel bus can use the system memory at the same time, provided that they do not try to use the same memory locations.
- When a bus master is reading from or writing to an I/O device or an adapter in a Micro Channel expansion slot, the microprocessor has exclusive access to system memory.

In servers that do not have a dual bus, the microprocessor is the default master, which means that it has to wait until no other masters are controlling the Micro Channel bus before it can have access to system memory.

Burst Data Transfers

In non-Micro Channel servers, transferring each byte of data is a two-step process. First, the microprocessor signals that it is going to send a byte of data. Then it sends the byte and signals that it is going to send the next byte. The microprocessor cannot perform any other tasks while it is managing a data-transfer operation in this way. Micro Channel architecture supports burst data transfers, in which data is sent in multiple bytes without intervention by the microprocessor. This improves system performance and allows faster data transfers between devices.

In some servers, data from the hard disk is moved into a buffer on the hard-disk controller before it is transferred across the bus. (A *buffer* is a temporary storage space that compensates for a difference in the rate of data flow when data is transferred from one device to another.) Micro Channel architecture allows burst data transfers from the hard disk to memory, without placing the data in a buffer on the hard-disk controller.

Streaming-Data Procedure

The *streaming-data procedure* allows high-speed transfer of data between bus masters and slaves. This procedure supports high-speed transfers of large blocks of data for devices such as hard disk drives and network adapters.

The streaming-data procedure transfers blocks of sequentially stored data. In basic data-transfer operations, a target address is assigned for every byte of data that is transferred. The streaming-data procedure assigns a target address only to the first byte of data, and the rest of the data in the block follows in sequence and is assigned to sequential addresses. A streaming-data transfer operation takes 100 nanoseconds to send 4 bytes of data. This is a data-transfer rate of 40 million bytes per second, which is twice as fast as basic data-transfer operations.

The Micro Channel bus has 32 data lines and 32 address lines. During streaming-data transfer operations, the 32 address lines are used only at the beginning of a transfer cycle. Then they remain idle for the rest of the cycle. The multiplexed streaming-data procedure uses the address bus as another 32-bit data bus, allowing data to be transferred 64 bits at a time.

Bus-Parity Checking

Bus parity checking is a method of verifying that data has not been changed during a data-transfer operation. Bus parity checking uses an extra bit, known as a *parity bit*, that is sent with each byte of data as it is transferred across the bus. The parity bit is set to 1 or 0 so that each byte has an odd number of 1's (if the server uses odd parity) or an even number of 1's (if the server uses even parity). If the parity (odd or even) of the received byte does not match the parity of the byte as it was sent, an error occurred during transmission and the receiving device can request that the data be sent again.

Bus parity checking has become a common feature in most servers. It is not 100% accurate, but it greatly reduces the chance for errors. It is essential for most operations because of fast I/O devices, complex I/O configurations, and large memory subsystems.

Channel-Check Reporting and Error Logging

Errors can occur not only during data-transfer operations, but also while data is stored in system memory. The contents of a memory location can be changed accidentally, a memory module can be defective, or other hardware failures can occur in the server.

Your server uses *channel-check reporting* to detect hardware errors and *error logging* to record the errors. These records can be used to diagnose and correct problems in the server.

The channel-check reporting facility automatically locates random and intermittent errors while your server is operating. Information about any failing component is saved in the error log so that you can identify and replace the failing component.

Chapter 5. Data Storage Devices

In a *direct access storage device (DASD),* the time that is required to access data is not dependent on the location of the data on the storage medium. Examples of DASD are diskette drives, hard disk drives, CD-ROM drives, and rewritable optical drives. All of these devices, except hard disk drives, are removable-media drives.

SCSI-2 Fast/Wide Subsystem

PC Server 500 comes with either a disk array or non-disk-array SCSI-2 Fast/Wide adapter preinstalled as a standard feature. (See Chapter 6, "Disk-Array Subsystems" on page 49 for more information about array sub-system information.) These adapters support the following:

- A variety of SCSI devices that conform to ANSI SCSI Standards X3.131-1986 or X3.131-1994.
- The array adapter supports two channels, with up to seven SCSI physical devices per channel. One of these channels is a dedicated internal channel and you can set the second as either an internal or external channel.
- The non-array adapter supports a total of fifteen SCSI devices. The internal channel supports up to seven SCSI physical devices. The external channel supports from one to fourteen SCSI devices. However, the total number of SCSI devices attached to a single SCSI adapter cannot exceed fifteen. (Keep in mind that wide devices support SCSI IDs 0 to 15, but narrow devices support only IDs 0 to 7.)
- Devices that are compatible with the SCSI Common Command Set.
- Functions that are unique to a particular SCSI device, when they are supported by a specific device driver that is provided by the device manufacturer.
- Replacing drives without turning off the server, which is called drive *hot swapping.*

Hard Disk Subsystem

Some models of the PC Server 500 are shipped with a built-in disk-array subsystem. If your server is a non-disk-array model, you can install an IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A and hard disk drives to create a disk-array subsystem.

The IBM RAID Adapter provides your server with new technology in hard disk management. The hard disk grouping techniques greatly enhance logical drive capacity and performance, and the *hot swap* capability for hard disk drive replacement (non-disk array models also) allows many failures to be corrected without turning off the server.

Internal Bays

Your server comes with one or more hard disk drives (depending upon the model), one 3.5-inch 2.88MB diskette drive, and one CD-ROM drive.

Hard Disk Drives: The drives installed in your server include one or more (depending upon the model) 1GB or 2GB, 3.5-inch SCSI-2 Fast/Wide Disk Drive.

Diskette Drive: The 3.5-inch, 2.88MB diskette drive uses 1MB, 2MB, and 4MB diskettes.

CD-ROM Drive: The 5.25-inch CD-ROM drive is installed vertically in bay A.

Open Bays: Additional DASD bays provide the capability of extensive storage expansion. You have flexibility in configuring data storage. For installation instructions and information on the types of drives that you can install in each bay, see the *User's Handbook.*

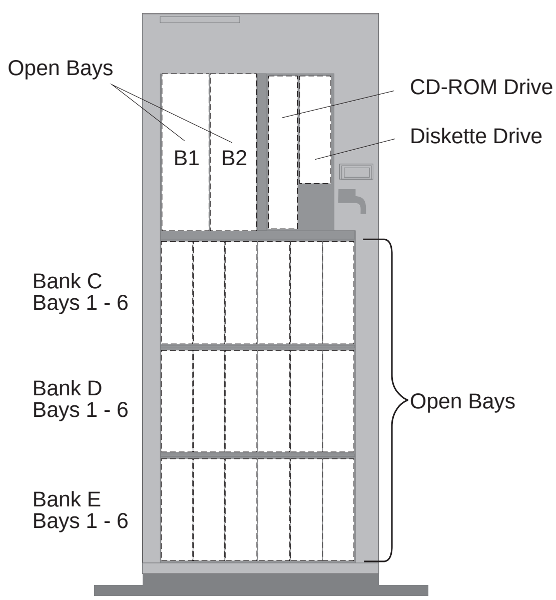

Diskette Drives

The following 3.5-inch and 5.25-inch diskette drives are available for purchase as optional drives that you can install in your server.

3.5-Inch, 2.88MB Diskette Drive

This drive uses 3.5-inch diskettes with a formatted capacity of 720KB (KB equals approximately 1000 bytes), 1.44MB (MB equals approximately 1 000 000 bytes), or 2.88MB. To prevent the loss of data, you must always format 1MB diskettes to store 720KB of data, 2MB diskettes to store 1.44MB of data, and 4MB diskettes to store 2.88MB of data. This drive has a *media-sense* feature, which uses the small window at the bottom-right corner of the diskette to detect whether the diskette capacity is 1MB, 2MB, or 4MB.

The 2.88MB diskette drive is available in two models, one of which has a *removable media security* feature. This feature prevents the removal or insertion of diskettes by an unauthorized person.

3.5-Inch, 1.44MB Diskette Drive

This drive uses 3.5-inch diskettes with a formatted capacity of 720KB or 1.44MB. To prevent the loss of data, always format 1MB diskettes to store 720KB of data and 2MB diskettes to store 1.44MB of data. The slim-high model of this drive includes the media-sense feature.

5.25-Inch, 360KB Diskette Drive

This drive uses 5.25-inch diskettes with a formatted capacity of 360KB. It is available only as an external drive.

5.25-Inch, 1.2MB Diskette Drives

This drive uses 5.25-inch diskettes with a formatted capacity of 360KB or 1.2MB. Diskettes that are formatted on this drive to store 360KB of data might not be usable in 360KB diskette drives. An external model and an internal model are available.

The internal model attaches to the diskette drive cable. The external model attaches to an adapter that you install in an expansion slot.

SCSI Hard Disk Drives

A SCSI controller and at least one (depending upon your model) SCSI hard disk drive are included with your server. You can connect up to six additional devices of various types, including hard disk drives, CD-ROM drives, rewritable optical drives, and tape drives, to each channel of the SCSI controller. There are two internal channels on the SCSI controller in your server.

The following figure lists the IBM SCSI hard disk drives that are available. These drives can be installed inside your server or in an external storage enclosure. All these drives comply with ANSI SCSI Standards X3.131-1986 and X3.131-1994. Seek time, expressed in milliseconds (ms), is the time that is required for the access arm of a DASD to be positioned on the appropriate track.

IBM SCSI Hard Disk Drive	**Average Seek Time Read**	**Write**
540MB (narrow)	8.5 ms	9.5 ms
1GB[1] (narrow)	8.5 ms	10.1 ms
1GB (wide)	7.8 ms	9.3 ms
2GB (wide)	7.8 ms	9.3 ms
2GB[2](narrow)	9.2 ms	10.7 ms
2GB[3] (wide)	9.2 ms	10.7 ms
4GB (wide)	8.4 ms	9.9 ms

Table 1. Average Seek Times for SCSI Hard Disk Drives

CD-ROM Drive

CD-ROM is an abbreviation for *compact disc-read only memory*. The removable medium that is used with a CD-ROM drive is a compact disc (CD) that the server can read from but not write to. Data can be viewed but not updated or changed.

---

1 This drive is 25.4 millimeters (1 inch) in height and is part number 70G8492.

2 This drive is 41.3 millimeters (1.63 inches) in height and is part number 70G8493.

3 This drive is 41.3 millimeters (1.63 inch) in height and is part number 70G8494.

Diskette drives, hard disk drives, and tape drives use a magnetic recording technique; CD-ROM drives use an optical technique that records and reads data at a much higher density than magnetic recording can achieve. All IBM CD-ROM drives use industry-standard 5.25-inch CDs with a capacity of approximately 650MB of data.

The CD-ROM drive that is installed in your server has an average access time of 20 ms (milliseconds); which is the interval between the instant the call for data is initiated and the instant the delivery of data is completed; a data transfer rate of 300 KBps (thousand bytes per second); and a burst data transfer rate of 4.2 MBps (million bytes per second).

This CD-ROM drive supports CD-ROM Extended Architecture (XA), which is a standard for interleaving different types of data such as video, audio, and text. This enhancement allows you to work with advanced multimedia applications.

This CD-ROM drive also supports the Kodak** Photo CD standard for recording pictures on CDs and is equipped with line output jacks and a volume-control dial. This CD-ROM drive complies with ANSI SCSI Standard X3.131-1986.

Startable CDs

The ServerGuide package shipped with your server is a startable (sometimes referred to as bootable) CD. This new feature allows you to start your server using a CD. A menu displays the options from which you can start.

Handling Compact Discs

Handle a CD by the edges only. You should periodically clean your CDs. Wipe them with a soft, lint-free cloth, from the center of the CD to the outer edge. Do not clean CDs in a circular pattern. This might damage the CD, causing loss of data.

** Trademark of Kodak Corporation.

3.5-Inch Rewritable Optical Drive

This enhanced rewritable optical drive is a SCSI direct access storage device that uses three types of removable 3.5-inch optical cartridges:

- The IBM PS/2 128MB Rewritable Optical Cartridge has a formatted capacity of 127MB and can be written to and read from many times.

- The optical read-only memory (O-ROM) cartridge has a formatted capacity of 122MB and can be read from but not written to. Using an O-ROM cartridge is similar to using a compact disc. An O-ROM cartridge can be used to distribute almost any type of data and software, such as operating systems and multimedia application programs.

- The partial read-only memory (P-ROM) cartridge has a formatted capacity of 122MB. Any part of the cartridge (up to 122MB) can contain read-only data. The remaining part of the cartridge can be formatted and then written to and read from many times.

  A P-ROM cartridge, for example, could be used by a testing or survey service. The test or survey questions would be entered by the service in the read-only part of the cartridge. The rest of the cartridge would provide sufficient space for read and write operations, so that the person taking the test or survey could enter the test answers or survey responses.

The IBM 3.5-Inch Enhanced Rewritable Optical Drive complies with ANSI SCSI Standards X3.131-1986 and X3.131-1994. You can install this drive inside servers that have a SCSI controller, or you can install it in an external storage enclosure and then connect the storage enclosure to an external SCSI connector. This drive reads data at a rate of 640 KBps and writes data at a rate of 213 KBps. Its burst data-transfer rate is 7.25 MBps. The average seek time (the time that is required for the access arm to be positioned on the appropriate track) for this drive is 40 ms.

Tape Drives

Backing up the entire contents of a hard disk to diskettes can be very time-consuming and, in the case of a network server that has multiple hard disks, might require hundreds of diskettes. A faster and more efficient way to back up the data is to use a tape-backup drive. Using a tape-backup drive, you can copy several billion bytes of data from the hard disks to a single tape.

Tape drives are sequential-access devices that provide a cost-effective way to back up your files. In a sequential-access (or serial-access) device, data is stored in the storage medium in the same sequence as it was entered.

There are three IBM tape drives available for use with your server. Two of these drives are internal drives that use 4 mm tape cartridges. Both of these drives feature data compression, which allows you to store more data on the cartridges. The main difference between these SCSI Digital Audio Tape drives is the external size of the drive. They are:

4/10GB 4 mm 3.5-Inch DAT Tape Option Kit
4/10GB 4 mm 5.25-Inch DAT Tape Option Kit

The third available tape backup drive is the *3445 External 5/10GB 8 mm Tape Drive*. This drive attaches to the external SCSI connector on the back of your server. It uses 8 mm tape cartridges and also offers the data compression feature.

IBM External SCSI Storage Enclosures

A SCSI storage enclosure is used to increase the data-storage capacity of any Micro Channel server that has a SCSI controller. IBM offers three types of SCSI storage enclosures. All types attach to the *external* SCSI connector of a SCSI controller. A SCSI Card to Option Cable is required to attach the storage enclosure to the server.

Model 3510 SCSI Storage Enclosure

The Model 3510 accommodates one 3.5-inch, half-high or slim-high SCSI hard disk drive, rewritable optical drive, or 5.25-inch CD-ROM drive. Multiple storage enclosures can be attached to the same SCSI controller. A SCSI Card-to-Option Cable connects the controller to the first Model 3510. An Option-to-Option Cable then connects the first Model 3510 to the second Model 3510. The last Model 3510 requires an Option Interface Terminator, which comes with the SCSI Card-to-Option Cable and terminates that end of the SCSI bus. The maximum length of cable that can be attached to each SCSI controller, including SCSI cable inside the server, is 6 meters (approximately 20 feet). For more information about maximum SCSI cable length, see your *User's Handbook*.

The Model 3510 has a universal 32-watt power supply, an external SCSI ID selector, and a push button power switch with power-on indicator.

Model 3511 External Storage Enclosure for SCSI Devices

The Model 3511 External Storage Enclosure for SCSI Devices is quite different from the Model 3510. From the outside, this floor-standing enclosure looks like a PS/2 Model 95 computer. This storage enclosure is available without the SCSI hard disk drive, and it can accommodate up to seven storage devices, such as hard disk drives.

The Model 3511 can hold up to five half-high or slim-high hard disk drives and one 2.3GB SCSI Tape Drive, providing excellent backup for your LAN server. You can also install up to two removable-media drives, such as CD-ROM drives and rewritable optical drives.

The manual provided with the Model 3511 includes information about power supply considerations. To avoid overloading the power supply, you must follow the guidelines that are provided. If you plan to install non-IBM SCSI devices, check with the drive manufacturer for the required power information.

Model 3516 Hot Swap Storage Expansion Enclosure

This DASD Hot Swap Storage Expansion Enclosure is a mini tower that holds up to seven hot swap hard disk drives. It ships without any drives, so you must install at least one drive before you attach the enclosure to the server. To connect the enclosure to the external SCSI-2 connector on the server, you need a 68-pin to 68-pin external SCSI-2 cable.

The SCSI controller in your server automatically sets the SCSI IDs in this storage enclosure, also. This means that you need not move jumpers to set SCSI IDs.

The manual provided with the enclosure includes information about setting it up, installing and removing drives, and attaching it to the server. There is also troubleshooting and getting service information in the manual.

External Disk-Array Subsystems

IBM offers two types of external SCSI disk-array subsystems: the Model 3514 and the Model 3515. Both models attach to the external connector of the SCSI adapter.

The Model 3514

This disk-array subsystem is a high availability, RAID-5 (see “Disk-Array Classifications” on page 56) disk array that reduces or eliminates data loss and system down time. It automatically performs data recovery and data reconstruction in the event of a single disk failure.

A Micro Channel array adapter and a signal cable come with the disk array for attaching it to the server. The floor-standing disk-array enclosure contains a controller, a redundant power supply, and three hard disk drives. There are two models of the storage enclosure: Model 3514-001, which contains three 398MB hard disk drives, and Model 3514-004, which contains three 968MB hard disk drives. Each model has space for up to five additional disk drives of the same size.

The Model 3515

This high-performance, RAID-1 (see “Disk-Array Classifications” on page 56) disk array protects against data loss and maintains high levels of performance. It automatically duplicates the data to a second hard disk and retrieves it in the event of a hard disk failure. If a hard disk fails, you can remove and replace the failing drive without turning off the server.

A Micro Channel array adapter, a connecting Serial Link cable, and an external disk subsystem come with the disk array for attaching it to the server. The Serial Link architecture allows concurrent disk operations. A single computer supports up to 48 drives with three adapters, for a maximum capacity of 51.3GB. There are three models of the disk subsystem: Model 3515-010, 220V is rack mounted, Model 3515-510, 110V and Model 3515-520, 220V are floor standing.

Twin-Tailing

One of the supported features of PC Server 500 is *twin-tailing*. Twin-tailing allows you to attach one or more drives to two servers. This gives you the opportunity to share data between the servers. One server is designated as the primary server (the only server that can write to the drive), but both servers can read from the drive.

Chapter 6. Disk-Array Subsystems

When you connect several hard disks together and configure the hard disk controller to access them in a predetermined pattern, you create what is known as a *disk array*. Disk arrays are used to increase security, performance, or reliability. They provide faster input and output rates than single large-capacity drives, because the files are stored on individually addressable disks and can be found more quickly during I/O requests.

The amount of security, performance, and improved reliability depends on the application programs that are running and the type of disk array being used. The type of disk array depends on the data-storage patterns supported by your hard disk controller, operating system, and application programs. (Your PC Server 500 supports RAID levels 0, 1, and 5.)

Internal Disk-Array Subsystems

IBM provides an integrated disk-array subsystem as a standard feature on some models of PC Server 500. Each disk-array model of PC Server 500 contains an IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A and two or more 1GB or 2GB hard disk drives. Your server has space for up to 18 *hot swappable* hard disk drives. (You can install one additional SCSI device in bay B, but this bay does not support hot swapping. The drive in bay B cannot be included in the array, unless no more than 5 drives are installed in bank C. See the *User's Handbook* for information about SCSI IDs.)

RAID Technology

RAID (redundant array of inexpensive disks) is the technology of grouping several hard disk drives in a server into an *array* that can be defined as a single logical drive. This *logical* drive then appears to the operating system as a single *physical* drive. This technology provides access and data transfer rates beyond the physical limitations of existing hard disk drives, greatly enhancing logical-drive capacity and performance. In addition, if one of the physical drives fails, the system continues to run, with no operator intervention required. The defunct drive can be replaced without turning off the server (*hot-swap*) and the new drive contents are rebuilt from the information on the other drives. This rebuilding process takes place in the background with the system online. Once the drive contents are rebuilt, full performance and fault-tolerant functions are restored.

When you install internal SCSI drives in your server, you can connect them to either the Channel 1 or the Channel 2 connector on the RAID adapter (see the *User's Handbook* for more information). The *external* connector on the adapter also is designated “Channel 2.” You cannot connect devices to both the internal channel 2 and external channel 2. This means that you can connect SCSI devices to both the internal channels and not the external channel; or, you can connect SCSI devices to one of the internal channels and the external channel, but never all three channels.

Hard Disk Drive Capacities

With a server, it is important to understand the implications of hard disk drive capacities and how they influence the way you create disk arrays.

Although the drives in the disk array can be of different capacities (for example 1GB or 2GB), they are treated in the disk array configuration as if they all have the capacity of the smallest disk drive. Therefore, if you have four drives of 1GB, 1GB, 1GB, and 2GB grouped in one disk array, the total capacity of the array will be 1GB times 4, or 4GB (instead of the 5GB physically available).

By the same logic, if you add a smaller drive to a group of drives, say a 1GB drive to a system containing three 2GB drives, the total capacity of an array that includes all four drives will be 4GB, instead of the 7GB physically available.

Hard Disk Drive Mapping

The drive locations shown on your RAID Configuration utility screen (described by channel and bay number) are shown mapped to the actual physical locations in your server in the following tables. The SCSI ID for each device is also shown in reference to the channel and bay to which it is attached. The first table shows the server configuration as it is shipped, with channel 2 not connected.

Displayed Channel Number	Displayed Bay Number	Bank and Bay Physical Location	SCSI ID
1	1	C1	0
1	2	C2	1
1	3	C3	2
1	4	C4	3
1	5	C5	3
1	6	C6	5
1	7	A1 (CD ROM)	6

Table 2. One Channel Mapping

The following table shows an example of using both internal channels of the RAID adapter. The first channel is connected to bank C, and the second channel is connected to the optional backplane on bank D.

Displayed Channel Number	Displayed Bay Number	Bank and Bay Physical Location	SCSI ID
1	1	C1	0
1	2	C2	1
1	3	C3	2
1	4	C4	3
1	5	C5	3
1	6	C6	5
1	7	A1 (CD ROM)	6
2	1	D1	0
2	2	D2	1
2	3	D3	2
2	4	D4	3
2	5	D5	4
2	6	D6	5
2	7	Not Used	6

Table 3. Two Internal Channels Mapping

The following table shows an example of using one internal and one external channel on the RAID adapter. The first channel is connected to bank C, and the second channel is attached to an external DASD storage enclosure. Refer to the documentation that came with the storage enclosure for physical locations.

Displayed Channel Number	**Displayed Bay Number**	**Bank and Bay Physical Location**	**SCSI ID**
1	1	C1	0
1	2	C2	1
1	3	C3	2
1	4	C4	3
1	5	C5	3
1	6	C6	5
1	7	A1 (CD ROM)	6
2	1	External	0
2	2	External	1
2	3	External	2
2	4	External	3
2	5	External	4
2	6	External	5
2	7	External	6

Table 4. One Internal and One External Channel Mapping

If you install another IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A in your server, you can use the same combinations shown here, depending upon your space availability.

Additional Storage Capacity

When you add hard disk drives to your server, you must configure a new disk array before you can use the drives. You can either reconfigure the existing disk array to include the capacity offered with the added drives, or group the added drives into their own array. You also can create an array with only one drive.

Logical Drives

It is helpful to understand how the system manages logical drives and how many you can define.

When you create an array, you are combining several hard disk drives into one storage area. The array then can be used as a single logical drive or can be subdivided into several logical drives. A logical drive of a disk array can be any size you choose within the size limitations of the array. The RAID Adapter supports up to four independent arrays and a total of eight logical drives. Each array can be formed from a maximum of eight drives. And an array can span both channels. (For information about *physical drives* supported, see Chapter 5, “Data Storage Devices” on page 37.)

For example, if you have only one array, it can be either a single logical drive or divided into as many as eight. If you have two or more arrays, you can have each one as one logical drive (a total of four), or you can divide them into multiple logical drives, as long as the total number of drives for the arrays is no more than eight.

The operating system considers each of these logical drives just as it does a physical hard disk drive. That is, the logical drives can be partitioned by the FDISK program (or its equivalent) in the same way that the operating system partitions a hard disk drive.

If you install an operating system, you either can allow the installation program to determine how the FDISK program (or its equivalent) allots the space within a logical drive, or manipulate the FDISK program yourself to partition the available space. The documentation you receive with your operating system explains how it handles *mapping*.

Improved System Performance

When hard disk drives are united into a single logical drive, data can be transferred in parallel from the multiple drives in the array. This parallel transfer yields data-transfer rates that are many times higher than with nonarrayed drives. This increased speed makes the server better able to meet the *throughput* (the amount of work in a given amount of time) or productivity needs of the multiple-user network environment.

The ability to respond to multiple data requests provides not only an impressive increase in throughput, but a decrease in response time. The combination of parallel transfers and simultaneous responses to multiple requests allows disk arrays to provide the highest level of performance in network environments.

Disk Array Adapter Features

The IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A supports a 32-bit Micro Channel data bus (automatically configurable); streaming-data transfer rates up to 40MBps; a 4MB adapter cache, which significantly improves server performance by keeping frequently used SCSI-device data available for immediate transfer; SCSI-2 Tagged Command Queuing; RAID levels 0, 1, and 5; *hot-spare drive*, which enables the disk-array subsystem to automatically perform data recovery and data reconstruction in the event of a single disk failure; and two internal 16-bit data-bus interfaces, which can support four arrays. (For information about external disk-array subsystems, see "IBM External SCSI Storage Enclosures" on page 44.)

With RAID technology, data is *striped* across an array of hard disk drives. *Striping* is the process of storing data across all the disk drives that are grouped in an array. This data-distribution scheme complements the way the operating system requests data.

Overlapped Input/Output Operation

Because the RAID Adapter provides multiple data paths to and from arrayed drives, your server can respond to requests from several users simultaneously. With its overlapped input/output operation, if one user requests data that resides on the first drive of the array, and a second user requests data that resides on the second drive, the adapter can simultaneously deliver both pieces of information.

Interleave Depth/Stripe Unit

The granularity at which data from one file is stored on one drive of the array before subsequent data is stored on the next drive of the array is called the *interleave depth.* The interleave depth can be set for optimum system performance to a stripe unit size from 8KB to 64KB bytes of data.

The collection, in logical order of these stripe units, from the first drive of the array to the last drive of the array, is called a *stripe.*

Queue Depth

Commands are queued in the adapter with a queue depth of 61. To obtain better performance, the commands in the queue will be *reordered* and *coalesced* on a hard disk drive basis. That is, the adapter organizes the commands according to which drive will be responding, and then orders and combines two or more commands, when possible, before sending them to the drives.

Cache

The RAID Adapter has 4MB of memory, of which more than 3MB is used as a data cache memory. This cache memory can be configured to operate in a write-through or write-back mode on a logical drive basis. Cache memory has parity to detect memory errors and retry algorithms to recover from errors that appear sporadically. For more information about cache, see “Cache Memory” on page 26.

Hot-Spare Drive Replacement

A *hot-spare drive* is a hard disk drive in your server that is defined for automatic use in the event of a drive failure. The hot-spare drive must be of equal or greater capacity than the drives in the array it is intended to replace. You can define as many hot-spare drives as you want.

If a drive fails, the system automatically switches to the hot-spare drive, and the data from the defunct drive is automatically recreated in the hot-spare drive. When the defunct drive is replaced using the operating system specific utility, the system automatically defines the replacement drive as a hot spare.

Data Protection

The RAID Adapter provides the server with a data redundancy technique that distributes data across all the drives in the array. In the event of a single drive failure, read and write requests are serviced by the remaining disks in the array, without loss of data. The defunct hard disk drive *appears* to be operating correctly, although at reduced performance, and despite its total failure.

When you replace the defunct drive, the data in logical drives assigned levels 1 and 5 is rebuilt on the new hard disk drive using the parity information stored on the other hard disk drives of the array. It is important that the defunct drive be replaced and rebuilt in a timely manner so that you avoid multiple-drive failures and reduced performance. Once the drive contents are rebuilt, full performance and fault-tolerant functions are restored.

Disk-Array Classifications

RAID is a term commonly used to define or classify the different types of data-storage patterns used by disk arrays. There are six basic RAID classifications, RAID-0 through RAID-5; three of which are supported by your server (levels 0, 1, and 5).

RAID Level 0

RAID level 0 stripes the data across all of the drives of the array and offers substantial speed enhancement, but provides for no data redundancy. Therefore, a hard disk failure within the array results in loss of data in the logical drive assigned level 0, but *only in that logical drive.*

Note: Although in this case, logical drives assigned level 1 or 5 in the same array do not suffer loss of data, the status of these drives is *Critical,* and they cannot sustain another hard disk drive failure. They operate at reduced performance. Replace and rebuild the defunct hard disk drive promptly to avoid a multiple hard disk drive failure.

When the defunct drive is replaced, all the logical drives assigned RAID levels 5 and 1 are rebuilt onto that hard disk drive; the level-0 logical drive is defined, but the data that was in the defunct level-0 logical drive is lost.

Even though the risk of data loss is present, you might want to consider assigning RAID level 0 to one of the logical drives so that you can take advantage of the speed offered with this level.

One way you might use this logical drive is to enter data that you back up each day and for which safety is not of primary importance; that is, you can re-create the data easily.

Another instance when you might want to use a level-0 logical drive is when the work you are doing requires maximum capacity—level 0 provides you the largest capacity of the three RAID levels offered, because no room is taken up for data parity storage.

Notes:

1. A hot-spare drive will not be activated in the event of a drive failure in an array that contains a logical drive assigned RAID level 0.
2. All logical drives in an array containing only one drive are automatically assigned RAID level 0.

RAID Level 1

RAID level 1 provides an enhanced feature for disk *mirroring* that stripes data and copies of the data across all the drives of the array. The first stripe is the data stripe; the second stripe is the *mirror* (copy) of the first data stripe, but written on another drive. Because the data is mirrored, the capacity of the logical drive when assigned level 1 is 50% of the physical capacity of the grouping of hard disk drives in the array. RAID level 1 uses two drives; if more than two drives are in the array, your RAID level is automatically switched to Enhanced RAID level 1.

If you have three or more drives in the array and want data redundancy, it's a good idea to choose level 5 before level 1. This illustration shows data arranged in a disk array with three hard disk drives, with the logical drive assigned RAID level 1. Notice that the mirror of the first data stripe (in this case, XXX YYY ZZZ) is shifted one drive. The other data stripes in level 1 follow the same pattern.

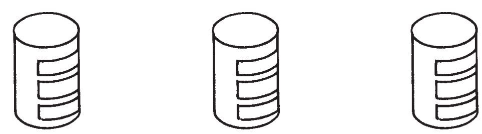

Stripe	XXX	YYY	ZZZ
Mirrored stripe	ZZZ	XXX	YYY
Stripe	AAA	BBB	CCC
Mirrored stripe	CCC	AAA	BBB

If you have only two drives available (for example, if you define one of the drives that came with your server as a hot spare), the second drive is a mirror copy of the first drive.

RAID Levels 2, 3, and 4

RAID levels 2, 3, and 4 are not supported for use on your PC Server 500. Level 2 uses a technique called Bit Interleave Data Striping and level 3 uses a technique called Bit Interleave Data Striping with Parity. Neither of these techniques can handle multiple, simultaneous small requests for data without causing a decrease in performance. RAID level 4 uses a technique similar to level 3.

RAID Level 5

RAID level 5 stripes data and parity across all drives of the array. When a disk array is assigned level 5 the capacity of the logical drive is reduced by one drive (for data parity storage).

Level 5 is generally the most desirable choice, because it offers both data protection and increased throughput. Level 5 gives you higher capacity than level 1, but level 1 offers better performance. If after using level 5 you are dissatisfied with the performance and can tolerate lower capacity, you can either redefine the level to level 1, or you can use a logical drive that you have assigned level 1.

Note: If you want to have a hot-spare drive and also assign level 5, you must have at least four hard disk drives, with no logical drives assigned level 0.

Chapter 7. Configuration

Configuring is the process of identifying the software and devices that are installed in or attached to your server, then deciding how the devices and software will use or share the available server resources.

Micro Channel servers use the Programmable Option Select (POS) feature, configuration programs, and configuration files to assign interrupt levels, input/output (I/O) addresses, memory addresses, arbitration levels, and other server resources.

The POS Feature

One of the most appealing features of the Micro Channel design is that it eliminates the use of jumpers and switches on the system board and on adapters. You no longer need to disassemble the server, locate the adapter publications, and manually set jumpers and switches to assign server resources.

Instead of jumpers and switches that you set manually, the system board and adapters have registers that act as electronic switches. These read and write registers are set by the POST routine each time you turn the server on. The registers are volatile storage areas, which means that they retain the configuration information only while the server is turned on. When you turn the server off, the configuration information is cleared from these registers.

A current copy of the configuration information is stored in a battery-backed memory area known as nonvolatile RAM. The same battery that supports the nonvolatile RAM module supports the real-time clock module, which provides a time-of-day clock, a 100-year calendar, and 64 bytes of nonvolatile RAM that the server uses to store setup information, such as the power-on password.

Each Micro Channel adapter is assigned a unique, 16-bit POS identification (ID) number, which is permanently encoded in the adapter hardware. The POST routine and the configuration programs use the POS ID to identify the type and location of each adapter in the server. The configuration programs also use the POS ID during the configuration process to locate the appropriate configuration files for each adapter.

In summary, the POS hardware consists of registers on the system board and on the adapters, nonvolatile RAM on the system board, and the hardware used to store the POS ID. This chapter describes how the POS hardware, together with the configuration programs and configuration files, simplifies the process of configuring the server hardware by:

- Eliminating the need for setting jumpers and switches on the system board and adapters
- Enabling you to configure the server automatically or manually using menu-driven programs
- Identifying the type and location of each device in the server
- Resolving server resource conflicts

The Configuration Programs

The *Reference Diskette* and the *Diagnostic Diskette* that came with your server are referred to as the System Diskettes. These diskettes contain the configuration programs (system programs) that you can use to configure the server. (These programs also are shipped in the ServerGuide package.)

These menu-driven programs guide you through the hardware configuration process. In addition to using the menus, you can press the F1 key to get helpful online information about each program (for example, when or why you might want to use it).

Using the Configuration Programs

The only time you must use the configuration programs is when you install a new device, remove or change the position of an existing device, or replace the battery that supports nonvolatile RAM.

Most people use the Automatic Configuration program, which allocates the server resources for optimum performance. However, if you are experienced and want to change the default values, or you have a device that requires special considerations, you have the option of using the Change Configuration program to customize the configuration.

IBM provides the configuration programs in three places:

On the hard disk

For non-disk array models that come with one hard disk drive, the configuration programs are preinstalled in a protected partition (called the System Partition) on the hard disk.

On the Reference Diskette

For both the disk-array and non-disk-array models the configuration programs are on the Reference Diskette. The Reference Diskette is included with the server.

On the CD

For all models of PC Server 500, if you misplace your System Diskettes, you can use the Diskette Factory on the CD that came in the ServerGuide package to create new diskettes.

Additional configuration programs are provided for disk-array models. These programs are located in the ServerGuide package. You can create a RAID Controller Diskette using the Diskette Factory or you can run the disk-array configuration programs from the CD. You can use the IBM RAID Controller programs to configure your hard disk drives into disk arrays and logical drives that best serve your needs. The RAID Controller diskette also contains device drivers and administration utility programs for the IBM RAID Controller. Refer to the *Configuring the Disk Array* section of your *User's Handbook* for information about using these device drivers and utility programs.

Configuration Programs Descriptions

The configuration programs are as follows:

Set configuration lets you view, change, back up, or restore the configuration information. It also runs the Automatic Configuration program.

During the configuration process, the Set Configuration program reads the POS ID of each device in the server to determine its type and location. Then it uses the information in the configuration files to assign the appropriate server resources.

View configuration lets you view the current configuration information that is stored in nonvolatile RAM.

Change configuration lets you customize the configuration or resolve a configuration conflict.

When a configuration conflict exists, an asterisk (*) appears next to the conflicting item on the Change Configuration screen. If the conflict is with a server resource that cannot be shared, the asterisk appears to the left of the slot number that contains the conflicting device. If other choices are available for the conflicting device, the asterisk appears to the right of the slot number.

Backup configuration lets you back up the configuration information to the hard disk (non-array models only) or to a diskette.

It is a good practice to run this program each time you change the configuration. That way, if the battery fails and the system configuration information is lost, you can use the Restore Configuration program to recover the server's configuration information.

Restore configuration lets you restore, from the hard disk or a diskette to nonvolatile RAM, the configuration information that was copied by the Backup Configuration program.

Run automatic configuration configures the system board and adapters automatically. This is the simplest way to configure the server hardware, because you do not have to make any decisions about resource assignments.

However, this program sets the configuration assignments, such as the memory address of an adapter, to the default values for the device. If you have a custom configuration or you are using an operating system or application that requires specific server resources, you should use the Change Configuration program.

Set and view SCSI device configuration lets you assign, view, or change the configuration of some SCSI devices.

Note: For disk array models, you must use the RAID Configuration Utilities to view the SCSI drive configuration information.

This Set and view SCSI device configuration program also allows you to enable or disable the Presence Error Reporting feature and the Accept SCSI Configuration feature.

The Presence Error Reporting feature allows you to keep the configuration information for a SCSI device in nonvolatile RAM, even after you remove the device from the server. When the POST routine checks for the presence of a SCSI device and finds the Presence Error Reporting feature set to Disabled, the POST routine does not report a configuration error. This feature is very useful if, for example, you share a single tape-backup drive among several servers.

The Accept SCSI Configuration feature allows you to add and remove SCSI devices without reconfiguring the server. When the Accept SCSI Configuration feature is set to Enabled, the SCSI configuration is not checked during POST. This feature is very useful if you frequently add and remove SCSI devices.

Display memory map lets you view the adapter memory assignments in the address range from hex 0C0000 through 0C7FFF. This program also shows the memory addresses that are available for additional hardware and software, such as an adapter or an expanded memory specification (EMS) driver.

Configuration Files

The Set Configuration program uses configuration files to assign interrupt levels, I/O addresses, memory addresses, arbitration levels, and other server resources.

All the files that are needed to configure the standard features of the server are preinstalled in the System Partition (non-array models only), on the Reference Diskette, or in the ServerGuide package.

Note: To make a copy of the disk-array configuration files, use the Diskette Factory in the ServerGuide package.

When you install an optional device, such as a network adapter, you must copy the configuration files from the Option Diskette that comes with the device to the System Partition or to the Reference Diskette.

The Set Configuration program uses three types of files to configure the server: adapter description files, adapter description programs, and initialization programs.

Some devices require only an adapter description file, but other, more complex devices might also require an adapter description program, an initialization program, or both.

Installing Option Configuration Files

The configuration files for an optional device normally are provided on an Option Diskette that comes with the device. This diskette contains files that the system uses to assign settings that define how the device will operate. Whenever you install a device that comes with an Option Diskette, you must copy the configuration files from the Option Diskette to the System Partition or to the backup copy of the Reference Diskette.

Most Option Diskettes contain three types of files:

Configuration files (file-name extension of ADF)
Diagnostic files (file-name extension of DGS)
Message files (file-name extension of PEP)

You should always use the instructions provided with the device to install configuration files. If no instructions are provided, you can use the Copy an Option Diskette program to install the configuration files.

Note: Some adapter manufacturers supply an adapter developer diskette. The copy-option procedure varies with an adapter developer diskette. Use the instructions provided along with the diskette to install these programs.

Before copying the configuration files, the Copy an Option Diskette program checks the System Partition or the Reference Diskette for duplicate file names. If it finds duplicate file names, it checks the dates of the files; then it copies only the files from the Option Diskette that have later dates than the existing files.

When these files are copied to the System Partition or the Reference Diskette, they become an integrated part of the system programs. Some Option Diskettes also contain device drivers and an automatic device-driver installation utility program. Keep the original Option Diskette in a safe place in case you need to rebuild the system programs or reinstall a device driver.

If your server has a hard disk with a System Partition, it is a good practice to make a backup copy of the partition after you install the configuration files. In the event of a hard-disk failure, you always would be able to recover the configuration information.

Configuration File Names

Every configuration file for a specific device uses the hexadecimal POS ID as part of its file name. For example, the POS ID for an adapter might be hex ABCD. If it is a very complex adapter, it requires an adapter description file, an adapter description program, and an initialization program. The names of these three files would be @ABCD.ADF, CABCD.ADF, and IABCD.ADF, respectively.

The following figure shows the types of files that are used to configure the server. You can identify each file by the first digit in the file name and by its file-name extension. The *xxxx* in the file name represents the POS ID for the device.

File Name	Description
@xxxx.ADF	Adapter description file
Cxxxx.ADF	Adapter description program
Ixxxx.ADF	Initialization program
Dxxxx.ADF	Adapter description file (system board)
Pxxxx.ADF	Adapter description file (built-in features)
Sxxxx.ADF	Adapter description program (system board)

Table 5. Identifying Configuration Files

Files that you can start from the operating-system prompt (known as *executable files*) normally have a file-name extension of EXE. Adapter description programs and initialization programs have an ADF extension to prevent you from accidentally starting these programs from the operating-system prompt, which would produce unpredictable results.

Adapter Description Files

The Set Configuration program uses the information in an adapter description file to identify the device and to allocate server resources.

Adapter description files list the POS ID and several choices for interrupt levels; I/O, RAM, and ROM addresses; arbitration levels; and other server resources that the device can use. Normally, these files also contain the informational text messages and interactive prompts that you use when you configure the server.

During the configuration process, the Set Configuration program locates and reads the adapter description file for each device. If the Set Configuration program cannot find the adapter description file, the server prompts you to insert the Option Diskette for the device in drive A. If you insert the diskette, and the Set Configuration program still cannot find the adapter description file, the Set Configuration program disables the device.

Adapter Description Programs

Some devices, such as the IBM SCSI Adapter with Cache, have a memory-address space that is not accessible before the operating system starts. These devices need an adapter description program to configure them during the configuration process.

The adapter description file for a device specifies the type of memory (system or nonsystem) that the device requires. System memory is controlled by the primary operating system, and only masters, such as bus masters, can use it. Nonsystem memory is not controlled by the operating system. When nonsystem memory is physically located on an adapter, it is known as memory-mapped I/O.

A device that requires an adapter description program has an entry in its corresponding adapter description file. This entry tells the Set Configuration program to run the adapter description program.

During the configuration process, the Set Configuration program reads the adapter description file for each device. If it finds an entry for an adapter description program, it locates and runs the adapter description program. If the Set Configuration program cannot find the adapter description program, the server prompts you to insert the Option Diskette for the device into drive A. If you insert the diskette, and the Set Configuration program still cannot find the appropriate adapter description program, the Set Configuration program disables the device.

Initialization Programs

Some devices require custom software, known as *device drivers*, to function properly. (Device drivers are programs that the operating system loads after the hardware is configured.) Some of these devices need an initialization program to activate them during the configuration process.

During the configuration process, the Set Configuration program reads the adapter description file for each device. If it finds an entry for an initialization program, it automatically installs the initialization program on track 0, head 0 of drive hex 80 (usually drive C).

The adapter description file also has an entry that assigns a priority to the initialization program. For example, if you have three devices in your server that require initialization programs, one device might require that its initialization program be run last.

All the initialization programs are run, according to their priority assignments, just before the operating system starts. If you start the server over a network or with an operating-system diskette, the initialization programs are bypassed, and the devices that require these programs are disabled.

Configuration Conflicts

The POS feature and the configuration programs usually assign server resources so that conflicts are avoided. But when a hardware device or a software program requires a resource that has already been assigned to another device or program, a conflict might occur.

Hardware Conflicts

The basic design of your Micro Channel server avoids conflicts among interrupt levels, arbitration levels, and DMA assignments. Conflicts among these resources are rare and usually occur when you install multiple adapters that require the same resources.

Most adapters are very flexible, and their adapter description files list many choices of interrupt levels; I/O, RAM, and ROM addresses; arbitration levels; and DMA channels that they can use. It is unlikely that these adapters will cause hardware configuration conflicts.

Unfortunately, there are other adapters whose adapter description files list only one choice. These adapters are most likely to cause hardware configuration conflicts.

You can find hardware conflicts by using the Change Configuration program. If a hardware conflict exists, an asterisk (*) appears next to the conflicting item on the Change Configuration screen. If the conflict is with a server resource that cannot be shared, the asterisk appears to the left of the slot number that contains the conflicting device. If other choices are available for the conflicting device, the asterisk appears to the right of the slot number.

The configuration programs assign server resources, starting with the adapter in the lowest-numbered expansion slot. If you have a hardware conflict that you cannot resolve using the Change Configuration program, you can install the more flexible adapters (those whose adapter description files list many choices for resources) in higher-numbered slots, and the less flexible adapters (those whose adapter description files list few choices or fixed resources) in the lower-numbered slots.

If you need to determine how flexible an adapter is, use a text editor to view its adapter description file.

Hardware and Software Conflicts

Some operating systems and application programs have special requirements for memory-address space that are not detected by the hardware configuration programs. Occasionally, these requirements are for memory-address space that is assigned already to a device.

When a memory-address conflict exists, one or more of the following might occur:

- The server cannot load the operating system.
- The server does not function.
- An application program does not operate or returns an error.
- Screen messages indicate a conflict exists.

There are two ways to resolve this type of conflict: change the software configuration or change the hardware configuration.

Changing the Software Configuration

The best way to resolve a memory-address conflict is to change the software configuration. This method allows you to continue using the Automatic Configuration program to configure the server hardware.

To resolve a memory-address conflict, use the Display Memory Map program to view the memory addresses that are available; then use a text editor to change the memory address in the configuration file (usually the configuration file for the software program or the CONFIG.SYS file).

Some programs do not allow you to modify the software configuration. To use these programs, you must change the hardware configuration so that it conforms with the requirements of the software program.

Changing the Hardware Configuration

A less desirable way to resolve a memory-address conflict is to change the hardware configuration. When you use the Change Configuration program, you create a custom configuration. You then must avoid using the Automatic Configuration program, because it sets the configuration assignments (such as the memory address of an adapter) to the default values for the device. So if you use the Change Configuration program to resolve a conflict, the conflict will recur if you run the Automatic Configuration program.

To change the hardware configuration, use the Display Memory Map program to view the addresses that are available; then use the Change Configuration program to assign a new memory address to the conflicting adapter.

Chapter 8. Input/Output Ports

The input/output (I/O) connectors (ports) on the back of your server are for attaching external devices, such as printers, keyboards, and displays, to your server. The I/O connectors include:

- Two serial ports
- Two parallel ports
- One video port
- One keyboard
- One auxiliary-device port
- One SCSI port

Serial Ports

Your server comes with two serial ports. Serial ports are used to communicate with printers, plotters, external modems, and auxiliary terminals and to transfer data between computers.

Serial ports transfer data one bit at a time, using direct memory access (DMA). DMA is a method of transferring data between I/O devices and system memory without intervention by the system microprocessor.

Serial ports can transfer data *asynchronously*, which means that they can transmit any number of characters at any time, with no restriction on the duration of the pauses between characters.

The serial ports can transmit and receive data and commands at a rate of 300 bits per second up to 345 6 000 bits per second. (To use a serial port at 345.6 thousand bits per second, you need a special shielded cable. For information about this cable, contact your IBM authorized reseller or marketing representative.)

Serial Port Connector

Each serial port uses a 9-pin, male D-shell connector on the back of the server. The pin-number assignments of this connector conform to the RS-232C industry standard. The following figure shows the pin-number assignments for the serial port connector.

Pin	Signal	Pin	Signal
1	Data carrier detect	6	Data set ready
2	Receive data	7	Request to send
3	Transmit data	8	Clear to send
4	Data terminal ready	9	Ring indicator
5	Signal ground		

Table 6. Serial Port Connector

Serial Port Numbering

When you turn on your server, the power-on self-test (POST) routine assigns the serial ports to specific communication port addresses, COM1 and COM2. You can install special adapters to increase the number of serial ports in your server. Servers can support a maximum of 16 serial ports (COM1 through COM16).

Some application programs use only COM1 and COM2, and some modems are designed to be selected only at certain communication port addresses. So in some situations, you might need to change communication port address assignments to resolve conflicts. To do this, use the Change Configuration program.

Parallel Ports

Your PC Server 500 server comes with two 25-pin parallel ports. Parallel ports are usually used to communicate with printers, and they transfer data one byte at a time using DMA.

Parallel port A (the second one from the top) is a high-speed port. If you share a printer with others on your network or the documents you print are large, you should connect your printer to port A to reduce your printing time.

Parallel Port Connectors

The parallel ports use a 25-pin, female D-shell connector on the back of the servers. The following figure shows the pin-number assignments for the parallel port connectors. (The high-speed port A, conforms to the Institute of Electrical and Electronics Engineers (IEEE) Standard 1284.)

Pin	Signal	Pin	Signal
1	STROBE	14	-AUTO FD XT
2	Data 0	15	-ERROR
3	Data 1	16	-INIT
4	Data 2	17	-SLCT IN
5	Data 3	18	Ground
6	Data 4	19	Ground
7	Data 5	20	Ground
8	Data 6	21	Ground
9	Data 7	22	Ground
10	-ACK	23	Ground
11	BUSY	24	Ground
12	PE	25	Ground
13	SLCT		

Table 7. Parallel Port Connector

Parallel Port Numbering

When you turn on your server, the POST routine assigns the high-speed port to LPT1 and the standard port to LPT2. PC Server 500 can support a maximum of four parallel ports. You can add additional parallel ports to your server by installing special adapters. The POST routine assigns the additional parallel ports to LPT2 through LPT4. You can change the port assignments by using the Change Configuration program.

Video Port

Your server comes with a 15-pin video port on the video adapter installed in expansion slot 5. The video port allows you to attach a video display to your server. The following figure shows the pin-number assignments for the video port.

Pin	Signal
1	Red
2	Green or monochrome
3	Blue
4	Monitor ID 2 (not connected)
5	Ground
6	Red ground
7	Green ground or monochrome ground
8	Blue ground
9	Reserved
10	Ground
11	Monitor ID 0 (not connected)
12	Monitor ID 1 (not connected)
13	Horizontal synchronization (Hsync)
14	Vertical synchronization (Vsync)
15	Monitor ID 3 (not connected)

Table 8. Video Display Port

Keyboard and Auxiliary-Device Ports

The PC Server 500 system board has one keyboard port and one auxiliary-device port. An auxiliary device can be a mouse or other pointing device. The following figure shows the pin-number assignments for the keyboard and auxiliary-device ports.

Pin	Signal
1	Data
2	Not connected
3	Ground
4	+5 V dc
5	Clock
6	Not connected

Table 9. Keyboard and Auxiliary-Device Port

SCSI Connectors

The array model has one dedicated internal 2-by-34-pin port for attaching small computer system interface (SCSI) devices. There is a second internal 2-by-34-pin port that is common with an external 2-by-35-pin port. This means that when the external connector is used, the second internal connector cannot be used.

The non-array model has one internal 2-by-34-pin port and one external 2-by-34-pin port. You can attach a total of 15 SCSI devices to these ports. These SCSI ports allow you to expand the capabilities of your server by attaching different types of SCSI devices, such as drives or printers.

The IBM SCSI-2 Fast/Wide Adapter/A (non-array model) and the IBM SCSI-2 Fast/Wide Streaming-RAID Adapter/A (array model), which come as standard features on the PC Server 500 models, support both internal and external SCSI devices. Table 10 on page 76 shows the pin-number assignments for the 2-by-34-pin SCSI port.

Pin	Signal	Pin	Signal
1	Ground	35	Terminator power
2	Data 12	36	Terminator power
3	Ground	37	Reserved
4	Data 13	38	Reserved
5	Ground	39	Ground
6	Data 14	40	Ground
7	Ground	41	Ground
8	Data 15	42	-Attention
9	Ground	43	Ground
10	Data P1	44	Ground
11	Ground	45	Ground
12	Data 0	46	-Busy
13	Ground	47	Ground
14	Data 1	48	-Acknowledge
15	Ground	49	Ground
16	Data 2	50	-Reset
17	Ground	51	Ground
18	Data 3	52	-Message
19	Ground	53	Ground
20	Data 4	54	-Select
21	Ground	55	Ground
22	Data 5	56	-Control/data
23	Ground	57	Ground
24	Data 6	58	-Request
25	Ground	59	Ground
26	Data 7	60	-Input/output
27	Ground	61	Ground
28	Data P	62	Data 8
29	Ground	63	Ground
30	Ground	64	Data 9
31	Ground	65	Ground
32	Ground	66	Data 10
33	Terminator power	67	Ground
34	Terminator power	68	Data 11

Table 10. The 2-by-34-Pin SCSI Port

Chapter 9. Video Subsystem

Your PC Server 500 is equipped with a super video graphics array (SVGA) video adapter. This adapter uses a 15-pin, analog video connector. For information about this connector, see "Video Port" on page 74.

The SVGA adapter supports high resolution graphics and fast refresh rates that exceed the specifications of the International Organization for Standardization (ISO) 9241 Part 3, Visual Display Requirements. It also provides reduced-flicker operation when used with noninterlaced displays that meet ISO standards. For more information about noninterlaced displays, see "Displays" on page 78.

The SVGA subsystem includes a video display buffer that contains 1024KB (KB equals approximately 1 000 bytes) of *video dynamic random access memory (VDRAM)*. The amount of VDRAM in your server determines the resolution and the number of colors that are supported. Additional features of the SVGA subsystem include:

- Integrated VGA emulation
- Screen resolution of 1024 x 768 pels with up to 256 colors (noninterlaced memory), 640 x 480 pels with up to 16.8 million colors (noninterlaced memory), or 1280 x 1024 pels with up to 256 colors (interlaced memory)

  Note: Operating system drivers are required for 1280 x 1024 pels with up to 256 colors. See the documentation that came with your operating system for more information.
- Support for the 132-column text mode
- Reduced-flicker operation for noninterlaced displays that meet ISO standards

Displays come in many sizes and types. Your server's video controller supports a wide range of monochrome and color displays with various capabilities.

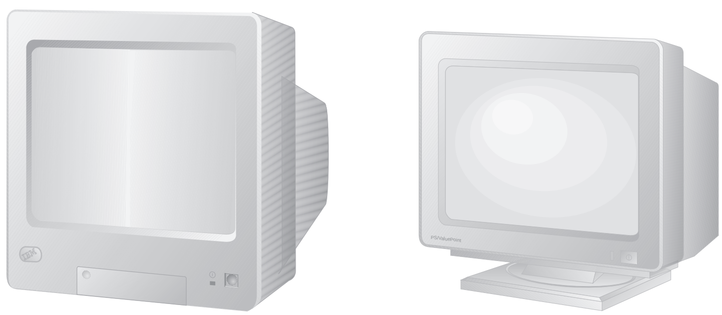

Displays

There are many types of displays. The most common distinguishing characteristics have to do with resolution, dot pitch, vertical-refresh rates, screen size, and color.

Resolution and Dot Pitch

If you look closely at the screen, you can see that the entire image is made up of many tiny dots. These dots are called *pels*. Resolution is measured by the number of pels needed to fill the screen (horizontal-by-vertical).

Dot pitch is more difficult to visualize. The inside of the display screen is coated with light-emitting materials called *phosphors*. Color displays use three phosphors: red, green, and blue. The phosphor coating is made up of dots arranged in a red-green-blue pattern. The distance from the center of one phosphor dot to the center of the next phosphor dot of the same color is called the *dot pitch*. Dot pitch is measured in millimeters, such as 0.28, 0.31, and 0.41.

In general, the display with the highest resolution and finest dot pitch produces the sharpest image. A display with a fine dot pitch, such as 0.28, produces sharper, better defined characters than a similar display that uses a coarser dot pitch, such as 0.41.

These characteristics can be very important. A system used extensively for detailed graphics might need a display that shows more colors, at a higher resolution and finer dot pitch, than one used exclusively for word processing. For some applications, a black-and-white (monochrome) display might do as well as a color display.

Vertical-Refresh Rate

The phosphors on the inside of the display screen hold their brightness and color for a very short time. Therefore, the image has to be redrawn many times each second to refresh the phosphors. The vertical-refresh rate is the speed at which the image on the screen is redrawn.

The image is redrawn from top to bottom. By the time the bottom rows are drawn, the top rows are starting to fade. There are two ways of redrawing the image; *interlaced* and *noninterlaced*. The interlaced method draws every other row from top to bottom, then starts at the top to fill in the rows that were skipped on the first pass. The noninterlaced method draws the complete image on each pass.

The vertical-refresh rate affects the stability of the image on the screen. Displays that have a fast, noninterlaced refresh rate, such as the IBM 952*x* displays, provide a very stable image with little or no flicker detectable. Displays with slower refresh rates tend to have more noticeable flicker. This is particularly noticeable with white and other light-colored backgrounds. You can minimize the flicker by changing the colors you work with. Many programs allow you to select colors.

Sensitivity to flicker varies from person to person, with young people being the most sensitive. Noninterlaced flicker occurs when the screen is redrawn at a rate that is slow enough to be detected by the human eye. Interlaced flicker occurs when the pattern of pels in adjoining lines is very different, or when an object on the screen is made up of alternating bright and dim lines.

Some application programs cannot reproduce all colors, and use a technique called *dithering* to combine two colors. For example, light red might be produced by alternating lines of white and red. As the image is continually refreshed using the interlaced method, the red lines are drawn on one pass and the white lines on the next pass. As a result, the light-red object might appear to flicker while the rest of the image on the screen remains stable.

Selecting a Display

In order to take full advantage of the video capabilities provided by your system, use a display with a fast noninterlaced refresh rate, such as the IBM 9524 or 9525 displays. However, as mentioned earlier, some applications are less demanding than others. IBM provides many displays that you can use with your system. The following is a partial list of displays and their characteristics.

Type	Dot Pitch	Resolution	Maximum Vertical Refresh Rate
6324	0.28	1024x768	72 Hz noninterlaced
		800x600	72 Hz noninterlaced
		640x480	72 Hz noninterlaced
6325	0.28	1024x768	72 Hz noninterlaced
		800x600	72 Hz interlaced
		640x480	72 Hz noninterlaced
9524	0.28	1024x768	72 Hz noninterlaced
		800x600	72 Hz noninterlaced
		640x480	72 Hz noninterlaced
9525	0.28	1024x768	72 Hz noninterlaced
		800x600	72 Hz noninterlaced
		640x480	72 Hz noninterlaced

Table 11. Display Types and Characteristics

For more details, refer to the information supplied with your display.

See your IBM authorized reseller or marketing representative for the latest list of displays.

Brightness and Contrast Controls

The Brightness and Contrast controls are used to adjust the image on the screen for maximum viewing comfort. On a new display, set the controls near their center position. Setting these controls to their maximum position can decrease the useful life of the display.

As the display gets older, the image tends to become dimmer. Adjust the Brightness and Contrast controls to compensate for this condition.

Screen-Saver Programs

Leaving the display turned on for prolonged periods without changing what appears on the screen can have an adverse affect on the display. Depending on what is on the screen, the image might burn into the phosphor coating (leaving a permanent image) or cause darkening of some areas of the screen.

You can avoid this problem by using a *screen-saver program*. A screen-saver program turns off the video signal when the keyboard has not been used for a defined period. This action results in a dark, blank screen. When you are ready to work with the system again, pressing a predefined key restores the image to the screen. Several companies produce screen-saver programs. Contact your IBM authorized reseller for details.

Avoiding Fatigue

Performing any visually demanding task for a prolonged period of time (for example, knitting, reading, or viewing a display screen), can cause eye fatigue. The American Academy of Ophthalmology considers displays safe for normal use and has declared they present no hazard to vision or do damage to the eye. To minimize eye fatigue, you should reduce any annoying glare and make sure your vision is properly corrected for the task. If eye discomfort persists, consult your eye-care professional or physician.

Similarly, sore muscles or joints can result from any job or activity during which you stay in a fixed position or repeat the same movement over a long period. Again, this is not unique to viewing a display screen. Eye fatigue and sore muscles and joints can be reduced by modifying work practices, such as taking breaks or rotating tasks, and by adjusting the display, lighting, and chair height.

Additional information is available by ordering the following publications through your nearest IBM branch office.

Comfort and You (G360-2050)
The VDT Workstation and Vision (SV04-0309)
Working with Displays (G325-0620)

Chapter 10. Security Features

Security features deter unauthorized use of your system and data. If you use your server in a public environment, such as an office, you might want to protect it and the data stored on it, by using one or more security features. Before implementing any security features, you might begin by evaluating your security needs. Where will the system be located? Does it need to be secured to permanent furniture or fixtures? Should use of the system be limited? When you have identified your security needs, you can activate or implement the appropriate security features.

LogicLock

LogicLock features, the advanced security features that come with your server include, tamper-evident switches. This active security feature uses microswitches on the covers to indicate if someone has tried to open the front cover. Keylocks and a system of passwords help deter problems of unauthorized use. Other features are operating-system dependent, which means if you want to use them, you must use an operating system that supports them. Other security features are optional features which can be purchased, such as cable covers. (To find out whether your operating system supports a particular security feature, see the documentation that came with your operating system.)

Securing Hardware

One of the best methods for protecting your hardware from theft is to keep it under lock and key. One method of deterring theft is to fasten the hardware to a stationary object, such as a table, a desk, or a wall. Another approach is to locate critical resources, such as servers, in a limited-access area behind locked doors.

The security features of your server include the U-bolt facility, which allows you to secure the system to a stationary object with a cable or chain. It also includes the door lock and keys.

U-Bolt Facility

If your server is in an area that is accessible to the public, you might need a method of securing it to a desk, table, or other stationary object. Two holes are provided at the rear of the server to accommodate a U-bolt. After the U-bolt is installed, you can use a cable or chain to lock the server in place. (This capability is referred to as a *cable-down* or a *bolt-down* feature.)

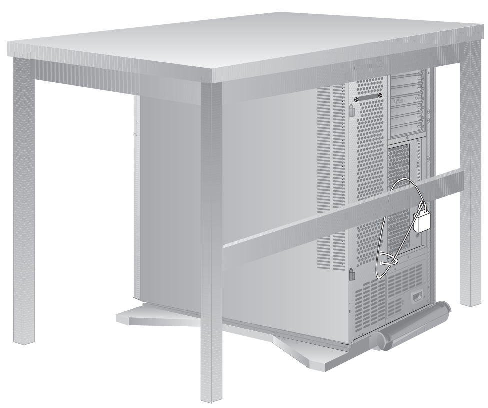

Step-by-step instructions for installing the U-bolt are in the "Installing Options" section of your *User's Handbook.*

Door Locks and Keys

There is one door lock which mechanically deters access to the drives. It controls access to the small door that covers the removable media drives as well as the entire front cover. If the covers are forced open without unlocking the lock, a microswitch detects the intrusion. The next time the system is started, the power-on self-test displays a message informing you that the covers have been tampered with, and that you must run the Automatic Configuration program before you can use the system.

Two keys are provided with your server. Always store the keys in a safe place. If you lose them, you must order a replacement lock mechanism and keys from IBM. Please note that anyone who has the key serial number and manufacturer's address can order duplicate keys, so store the tag in a safe place.

Cable Cover Option

At the rear of your server are several connectors that can be used to send and receive data. Anyone with the correct equipment and knowledge can use these connectors to gain access to the data stored on your hard disk and other storage devices. Your system covers are designed to accommodate a cable cover that helps guard against someone using the connectors in this way.

The cable-cover option prevents the cables from being removed from the rear of the server, and prohibits other computers or devices from being attached to the unused connectors.

Note: The cable cover does not protect against unauthorized access through devices attached outside of the system.

The cable-cover option is not a standard feature of your system; it is available from your IBM authorized reseller or marketing representative.

Unauthorized-Access Monitor

The unauthorized-access monitor feature works with the administrator (also known as privileged-access) password, when you are using DOS. If you set an administrator password and the unauthorized-access monitor feature is enabled, your server will alert you if an attempt has been made to tamper with the locked covers. If it detects any tampering, the system will stop all operations or display an error message. If the system stops, any data in memory waiting to be stored might be lost. (The response to tampering varies with the operating system you are using.)

When you set an administrator password, the unauthorized-access monitor is automatically enabled. If you do not want the system to stop operations if the covers are tampered with, set the unauthorized-access monitor to *Disable.* You can change the setting of this feature through the Change Configuration screen.

Removable Media

If you are using your server in an environment where a diskette is left in the drive while the server is unattended, an electronic-eject diskette drive can provide the security you need to protect the diskette from theft or damage. This separately purchasable diskette drive allows you to lock a diskette in the drive or prohibit someone from inserting a diskette in the drive, while the drive is not in use. However, the operating system you are using must provide a means of activating this feature (such as entering a password).

Not all operating systems support this feature. Contact your IBM authorized reseller or marketing representative for details.

Securing Data

It is very difficult to secure shared information from theft, but an effective method is to limit user access. Only users who need the software or data should be able to access it. This level of access control is provided as a standard feature of most network programs. Another common feature of network programs that helps secure data from theft is password protection. Passwords are easy to use and very effective. They help prevent unauthorized users from accessing sensitive files.

Some advanced network management programs can actually audit usage, based on names, adapter addresses, date, time, and unsuccessful attempts to access a file. This type of information can help you identify users who are attempting to access restricted data.

To secure extremely sensitive data, you might want to consider using a commercially available data-encryption tool. These tools encode the data files so that they are unintelligible, thus useless if stolen. There are two ways that you can encrypt data: by using a program or using an encryption device. The software programs are usually less expensive than the hardware devices, but they also are slower.

Power-On Password

You do not need to use a power-on password to use your server, but a password helps protect the information you store in it. You can use any combination of up to seven characters (A–Z, a–z, and 0–9) for your power-on password. Keep a record of your password in a secure place. If you forget it, you will have to remove the server cover and change the position of a specific jumper to erase it.

How the Power-On Password Works

The power-on password locks the keyboard and mouse (if attached to the mouse port) to help prevent unauthorized use of your server. If you are using a mouse that is connected to the serial port, the mouse remains active.

After you set a power-on password, `Enter password` appears each time you turn on the server. Before you can use the server, you must enter the correct password. (The password does not appear on the screen as you type it.) When you enter the correct password, `Password accepted` appears on the screen, the keyboard and mouse are unlocked, and the system resumes normal operation. If you type the wrong password, `Incorrect password` appears on the screen and `Enter password` is again displayed. After three incorrect attempts, you must turn off the server and start again.

Unattended Start Mode

After you set a power-on password, you can enable the *unattended start mode.* This mode locks the keyboard and mouse, but allows the system to start the operating system and carry out the instructions in the CONFIG.SYS and AUTOEXEC.BAT files. Although `Enter password` does not appear, the keyboard and mouse remain locked until you enter the correct password. This mode is ideal for systems that operate unattended. If a power failure occurs, the system automatically restarts when power returns and resumes normal operation, without operator intervention.

ATTENTION MOUSE USERS:

* Trademark of the IBM Corporation.

The following statement applies only to those who use a PS/2*-style mouse; a serial mouse is not affected.

In the unattended start mode, the keyboard and mouse ports are disabled (locked). Because of this, the system cannot detect that a mouse is attached, and an error occurs.

You must do one of the following:

- In the CONFIG.SYS file, set the operating system so that it does not stop on a device-driver error. For example, when using the OS/2 operating system, use the PAUSEONERROR=NO statement.
- Remove the mouse driver statement from the CONFIG.SYS file if you do not want to use a mouse.

When using the OS/2 operating system, if you do not perform one of the previous steps, the system issues an error message, halts, and prompts you to press Enter to continue. Before pressing Enter, type the power-on password.

Refer to your operating-system documentation for information about modifying your CONFIG.SYS and AUTOEXEC.BAT files.

Administrator Password

Before setting an administrator (also known as privileged-access) password, it is a good idea to read this section in its entirety. Using this security feature requires some planning and ongoing administration.

The administrator password allows you to control who has access to the system programs. If an administrator password is set, you must enter it to use the system programs in the System Partition on the hard disk or on the Reference Diskette. The administrator password also can be used to override the power-on password. After an administrator password is set, only those who know the password can perform tasks such as:

- Altering computer settings or features controlled by the system programs

- Running diagnostic tests
- Resetting the system after a forced entry (If the server is forced open, a POST error occurs. To clear the error, you must enter the administrator password.)

Your server is shipped with the administrator password feature *Disabled.* You must move a jumper on the system board before an administrator password can be set for the first time. The jumper has two positions:

- The *locked state* prevents an administrator password from being set, changed, or removed. This is the position set at the factory.
- The *change state* allows an administrator password to be set, changed, or removed.

The location of the administrator password jumper is shown in the *User's Handbook.*

Warning: If an administrator password is set, then forgotten, it cannot be overridden or removed. The system board must be replaced in order to access the system programs.

How the Administrator Password Works

Use the administrator password when highly classified information must be protected. Although it can be used with your power-on password, it provides a much higher level of security when used with an operating system that controls access through the use of passwords. This type of operating system is called a *secured* or *trusted computing base.*

You can use any combination of up to seven characters (A–Z, a–z, and 0–9) for your administrator password, just as you can with your power-on password. For additional security, the two passwords should not be the same.

One important difference between the power-on password and the administrator password is that a forgotten administrator password cannot be overridden or disabled. The single most important reason for setting an administrator password is that when one is set, only those who know the password can access the system programs and modify the hardware or change any of the settings.

If you type the wrong password, `Incorrect password` appears and `Enter the privileged-access password` is again displayed. After three incorrect attempts, `The system is locked` message is displayed and you must turn off the server and start again.

Setting an Administrator Password

In order to set an administrator password, you must first set your selectable drive-startup sequence. (If you try to set this password and the drive sequence is not set, an error code and text message with instructions appear.) You can select the default drive-startup sequence, which includes a diskette drive. To have a totally secure system, you can remove the diskette drive from the drive-startup sequence.

Forgotten Administrator Password

If an administrator password is set, then forgotten, it cannot be overridden or removed. If you forget your administrator password, you must replace the system board to gain access to the system programs.

Keyboard Password

A keyboard password allows you to lock the keyboard while the system is turned on. This is useful if you leave your system on and unattended. With the keyboard locked, no one can use or restart your system.

The way you set the keyboard password depends on the operating system you are using. The OS/2 operating system provides keyboard-password protection as a standard feature.

If you forget your keyboard password, turn off the server for at least 5 seconds; then turn it on. The keyboard password is erased from memory when you turn off the server.

Selectable Drive Startup

Selectable drive startup (sometimes referred to as selectable *boot*) allows you to control the startup sequence of the drives in your server. Each time you turn on the server, it checks the drives as it looks for the operating system. The order in which the system checks the drives is the *drive-startup sequence.*

In most cases, you do not need to change the *default* drive-startup sequence. However, if you set an administrator password, or are working with multiple hard disk drives, multiple operating systems, or different sized diskette drives, you might want to change the default drive-startup sequence.

The default drive-startup sequence checks the primary diskette drive for a self-starting (*bootable*) diskette. If one is present, the operating system or program is loaded from the diskette. If not, the system checks the primary hard disk for an operating system. If one is present, it is loaded from that hard disk.

If you start the system from a diskette, the drive containing the diskette becomes drive A, regardless of the defined sequence, and the first hard disk selected in the startup sequence becomes drive C. You can choose a startup sequence of up to four drives.

You can *customize* the startup sequence by changing the order in which the system checks the drives. You decide which four drives are the first to be checked, *and* the order in which the system checks them.

Notes:

1. If a Reference Diskette is present in a diskette drive which is not in the selectable boot sequence, POST will boot that diskette *only* if you press F1 or (Alt+Ctrl+Ins) at the server startup.
2. When you change the startup sequence, the drive letters also might change. The operating system assigns the drive letters when the system starts. Letters A and B *always* are assigned to diskette drives. Subsequent drive letters can be assigned to any type of drive based on the operating system or the device drivers used.

Warning: If you changed your startup sequence, you must be extremely careful when doing *write* operations (for example, copying, saving, or formatting). You can accidentally overwrite data or programs if you select the wrong drive.

Securing Data from Loss

There are several methods you can use to prevent the loss of data. You should consider these methods carefully, because some of them affect the performance, reliability, and hardware requirements of the server.

There are primarily three ways that you can protect your data from loss. You can wait until the data accumulates on the server, and then make backup copies of all the hard disks. You can configure a disk array to duplicate the data (create a redundant copy) as it is entered, and then store the duplicate copy on a separate hard disk. Or, you can configure a disk array to store the parity information about the data on the array as the data is being entered.

Backup Copies

A good method for preventing the loss of data is to make copies of the data. That way, in the event of a hardware failure, you can always recover using the copies.

Backing up the entire contents of a hard disk to diskettes can be very time-consuming and, in the case of a network server that has multiple hard disks, might require hundreds of diskettes. A faster and more efficient way to back up the data is to use a tape-backup drive. Using a tape-backup drive, you can copy several billion bytes of data from the hard disks to a single tape.

There are two problems with making backup copies. The first problem is someone has to remember to do it. The second problem is that you usually need to disconnect the server from the network to make the backup copies.

Redundant Copies

A convenient way to copy the contents of a hard disk is to duplicate all data as it is being entered. This duplicate copy is then stored on another hard disk. If one hard disk fails, you still have access to the data on the other. This method is more expensive than backing up to diskettes or tapes, because it doubles the amount of disk-storage space needed.

Parity Information

Storing parity information about the data as the data is being entered is an efficient, cost effective, and reliable method commonly used to guard against the loss of data. In the event of a hard disk failure, the parity information and the data on the remaining functional disks are used to reconstruct the data that was stored on the failed drive.

Viruses

Computer viruses are a problem that exists within the personal computer community. A computer virus is a program (or instructions hidden within a program) that *infects* other programs by modifying them without your knowledge. Like any other program, the virus can do anything it is programmed to do. Some viruses are practical jokes, causing unusual or erratic screen behavior. Others are destructive, erasing or damaging files or overloading memory and communication networks.

Viruses are difficult to detect. Many stay inactive until triggered by a specified event such as a date, command, or some other operation. Others are activated when an infected program is started a specified number of times. When the symptoms of the virus appear, it might be difficult to determine if the problem is a hardware failure, a problem in the software, or a virus in action.

Several programs are available that can detect the presence of many known viruses. These programs typically examine files and look for patterns associated with these viruses, or look for changes in the size of files. These programs are best used as a preventive measure to detect a virus before it becomes widespread or causes damage. Many computer users check for viruses on a regular basis.

When a virus is found, it must be removed. This might be as simple as replacing a file, or it might require the assistance of a trained technician.

Viruses are generally spread unknowingly from computer to computer when programs are exchanged or shared. If you don't know where a program came from, be careful. Most reputable program distributors and bulletin-board owners scan their files to guard against viruses and maintain records to identify program owners.

Here are a few tips to help guard against computer viruses:

- Write protect original program diskettes before using them. This ensures the diskettes will not be infected if a virus is already present in your system.
- Avoid using programs and diskettes from unknown sources.
- If you use programs provided from public-domain software distributors or bulletin-board services, find out what precautions they take to guard against viruses.
- Back up all critical data and programs regularly.
- Watch out for changes in file sizes (when they shouldn't change).
- Use password protection to limit access to your system or network.

Remember, not every problem is caused by a virus. If your system starts acting erratically, refer to your *troubleshooting charts* in the *User's Handbook* to test the system.

Erased Files

Erasing a file from a diskette or hard disk does not destroy the file. With the right type of software, all or part of an erased file can be reclaimed. This can cause a security risk if you sell a system or give someone a diskette that once contained classified or confidential information.

One way to help ensure that no readable information is left on a hard disk is to do a low-level format. An operating-system format operation does not remove all information from a hard disk.

The operating-system format operation works a little differently with diskettes. It writes a repeated pattern over the entire surface. Any information that was on the diskette becomes unreadable.

Depending on the type of information stored, your company might require additional safeguards.

Backup Power Supply

To prevent the loss of data caused by a power fluctuation, you might want to consider installing a backup power supply. There are two basic types of commercially available backup power supplies: *uninterruptible* and *standby*.

Uninterruptible Power Supplies

This type of power supply works as an interface between the main power source and the server, and it automatically takes over when the main power source is lost. An uninterruptible power supply connects directly to the main power source (electrical outlet), and then the server connects to the power supply. The power supply transfers the power from the main power source to its internal battery; then its internal battery provides the power to the server. The advantages of uninterruptible power supplies are that they require no switching time (because they are always on), and they protect the server from power surges. A disadvantage is that they usually cost more than standby power supplies.

Standby Power Supplies

This type of power supply waits until it detects a drop in voltage, and then it switches on and provides power to the server until its battery power is depleted. An advantage of a standby power supply is low cost. It is usually less expensive than an uninterruptible power supply. A disadvantage of a standby power supply is the amount of time that it needs to switch itself on; the more time it takes, the greater the chance of losing data.

Appendix A. Product Warranties

International Business Machines Corporation   Armonk, New York, 10504

Statement of Limited Warranty

The warranties provided by IBM in this Statement of Limited Warranty apply only to Machines you originally purchase for your use, and not for resale, from IBM or an IBM authorized reseller. The term "Machine" means an IBM machine, its features, conversions, upgrades, elements, or accessories, or any combination of them. Machines are subject to these terms only if purchased in the United States or Puerto Rico, or Canada, and located in the country of purchase. If you have any questions, contact IBM or your reseller.

Machine: IBM PC Server 500

Warranty Period*: Three Years

**Elements and accessories are warranted for three months. Contact your place of purchase for warranty service information.*

Production Status

Each Machine is manufactured from new parts, or new and serviceable used parts (which perform like new parts). In some cases, the Machine may not be new and may have been previously installed. Regardless of the Machine's production status, IBM's warranty terms apply.

The IBM Warranty

IBM warrants that each Machine 1) is free from defects in materials and workmanship and 2) conforms to IBM's Official Published Specifications. IBM calculates the expiration of the warranty period from the Machine's Date of Installation. The date on your receipt is the Date of Installation, unless IBM or your reseller informs you otherwise.

During the warranty period, IBM or your reseller will provide warranty service under the type of service designated for the Machine and will manage and install engineering changes that apply to the Machine. IBM or your reseller will specify the type of service.

For a feature, conversion, or upgrade, IBM or your reseller may require that the Machine on which it is installed be 1) the designated, serial-numbered Machine and 2) at an engineering-change level compatible with the feature, conversion, or upgrade. Some of these transactions (called "Net-Priced" transactions) may include additional parts and associated replacement parts that are provided on an exchange basis. All removed parts become the property of IBM and must be returned to IBM.

Replacement parts assume the remaining warranty of the parts they replace.

If a Machine does not function as warranted during the warranty period, IBM or your reseller will repair or replace it (with a Machine that is at least functionally equivalent) without charge. If IBM or your reseller is unable to do so, you may return it to your place of purchase and your money will be refunded.

If you transfer a Machine to another user, warranty service is available to that user for the remainder of the warranty period. You should give your proof of purchase and this Statement to that user.

Warranty Service

To obtain warranty service for the Machine, you should contact your reseller or call IBM. In the United States, call the IBM Personal Systems HelpCenter at **1-800-772-2227**. In Canada, call the IBM Personal Systems HelpCenter at **1-800-465-4955**. You may be required to present proof of purchase.

Depending on the Machine, the service may be 1) a "Repair" service at your location (called "On-site") or at one of IBM's or a reseller's service locations (called "Carry-in") or 2) an "Exchange" service, either On-site or Carry-in.

When a type of service involves the exchange of a Machine or part, the item IBM or your reseller replaces becomes its property and the replacement becomes yours. The replacement may not be new, but will be in good working order and at least functionally equivalent to the item replaced.

It is your responsibility to:

1. obtain authorization from the owner (for example, your lessor) to have IBM or your reseller service a Machine that you do not own;
2. where applicable, before service is provided —
   a. follow the problem determination, problem analysis, and service request procedures that IBM or your reseller provide,
   b. secure all programs, data, and funds contained in a Machine,
   c. inform IBM or your reseller of changes in a Machine's location, and
   d. for a Machine with exchange service, remove all features, parts, options, alterations, and attachments not under warranty service. Also, the Machine must be free of any legal obligations or restrictions that prevent its exchange; and
3. be responsible for loss of, or damage to, a Machine in transit when you are responsible for the transportation charges.

Extent of Warranty

IBM does not warrant uninterrupted or error-free operation of a Machine.

Misuse, accident, modification, unsuitable physical or operating environment, improper maintenance by you, or failure caused by a product for which IBM is not responsible may void the warranties.

THESE WARRANTIES REPLACE ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. HOWEVER, SOME LAWS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES. IF THESE LAWS APPLY, THEN ALL EXPRESS AND IMPLIED WARRANTIES ARE LIMITED IN DURATION TO THE WARRANTY PERIOD. NO WARRANTIES APPLY AFTER THAT PERIOD.

In Canada, warranties include both warranties and conditions.

Some jurisdictions do not allow limitations on how long an implied warranty lasts, so the above limitation may not apply to you.

Limitation of Liability

Circumstances may arise where, because of a default on IBM's part (including fundamental breach) or other liability (including negligence and misrepresentation), you are entitled to recover damages from IBM. In each such instance, regardless of the basis on which you are entitled to claim damages, IBM is liable only for:

1. bodily injury (including death), and damage to real property and tangible personal property; and
2. the amount of any other actual loss or damage, up to the greater of $100,000 or the charge for the Machine that is the subject of the claim.

Under no circumstances is IBM liable for any of the following:

1. third-party claims against you for losses or damages (other than those under the first item listed above);
2. loss of, or damage to, your records or data; or
3. economic consequential damages (including lost profits or savings) or incidental damages, even if IBM is informed of their possibility.

Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so the above limitation or exclusion may not apply to you.

This warranty gives you specific legal rights and you may also have other rights which vary from jurisdiction to jurisdiction.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights or other legally protectable rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, programs, or services, except those expressly designated by IBM, are the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood NY 10594, U.S.A.

Trademarks

The following terms, denoted by an asterisk (*) in this publication, are trademarks of the IBM Corporation in the United States or other countries:

IBM	Micro Channel
NetFinity	Operating System/2
OS/2	PS/2
ServerGuide	

Trademarks

The following terms, denoted by a double asterisk (**) in this publication, are trademarks of other companies as follows. Other trademarks are trademarks of their respective companies.

Intel	Intel Corporation
Kodak	Kodak Corporation
Novell	Novell, Inc.
Pentium	Intel Corporation

Power Cords

For your safety, IBM provides a power cord with a grounded attachment plug to use with this IBM product. To avoid electrical shock, always use the power cord and plug with a properly grounded outlet.

IBM power cords used in the United States and Canada are listed by Underwriter's Laboratories (UL) and certified by the Canadian Standards Association (CSA).

For units intended to be operated at 115 volts: Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a parallel blade, grounding-type attachment plug rated 15 amperes, 125 volts.

For units intended to be operated at 230 volts (U.S. use): Use a UL-listed and CSA-certified cord set consisting of a minimum 18 AWG, Type SVT or SJT, three-conductor cord, a maximum of 15 feet in length and a tandem blade, grounding-type attachment plug rated 15 amperes, 250 volts.

For units intended to be operated at 230 volts (outside the U.S.): Use a cord set with a grounding-type attachment plug. The cord set should be marked <HAR> and have the appropriate safety approvals for the country in which the equipment will be installed.

IBM power cords for a specific country are usually available only in that country:

IBM Power Cord Part Number	Used in These Countries
13F9940	Argentina, Australia, China (PRC), New Zealand, Papua New Guinea, Paraguay, Uruguay, Western Samoa
13F9979	Afghanistan, Algeria, Andorra, Angola, Austria, Belgium, Benin, Bulgaria, Burkina Faso, Burundi, Cameroon, Central African Rep., Chad, Czech Republic, Egypt, Finland, France, French Guiana, Germany, Greece, Guinea, Hungary, Iceland, Indonesia, Iran, Ivory Coast, Jordan, Lebanon, Luxembourg, Macau, Malagasy, Mali, Martinique, Mauritania, Mauritius, Monaco, Morocco, Mozambique, Netherlands, New Caledonia, Niger, Norway, Poland, Portugal, Romania, Senegal, Slovakia, Spain, Sudan, Sweden, Syria, Togo, Tunisia, Turkey, former USSR, Vietnam, former Yugoslavia, Zaire, Zimbabwe
13F9997	Denmark
14F0015	Bangladesh, Burma, Pakistan, South Africa, Sri Lanka
14F0033	Antigua, Bahrain, Brunei, Channel Islands, Cyprus, Dubai, Fiji, Ghana, Hong Kong, India, Iraq, Ireland, Kenya, Kuwait, Malawi, Malaysia, Malta, Nepal, Nigeria, Polynesia, Qatar, Sierra Leone, Singapore, Tanzania, Uganda, United Kingdom, Yemen, Zambia
14F0051	Liechtenstein, Switzerland
14F0069	Chile, Ethiopia, Italy, Libya, Somalia
14F0087	Israel
1838574	Thailand
62X1045	Bahamas, Barbados, Bermuda, Bolivia, Brazil, Canada, Cayman Islands, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Guyana, Haiti, Honduras, Jamaica, Japan, Korea (South), Liberia, Mexico, Netherlands Antilles, Nicaragua, Panama, Peru, Philippines, Saudi Arabia, Suriname, Taiwan, Trinidad (West Indies), United States of America, Venezuela

Electronic Emission Notices

Industry Canada Compliance Statement

This digital apparatus does not exceed the Class B limits for radio noise emissions from digital apparatus as set out in the interference-causing equipment standard entitled: "Digital Apparatus", ICES-003 of Industry Canada.

Avis de conformité aux normes d'Industrie Canada

Cet appareil numérique respecte les limites de bruits radioélectriques applicables aux appareils numériques de Classe B prescrites dans la norme sur le matériel brouilleur : "Appareils numériques", NMB-003 édictée par Industrie Canada.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult an IBM authorized dealer or service representative for help.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. Proper cables and connectors are available from IBM authorized dealers. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Glossary

This glossary includes terms and definitions from:

The *American National Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42 Street, New York, New York 10036. Definitions are identified by the symbol (A) after the definition.

The ANSI/EIA Standard—440-A: *Fiber Optic Terminology*. Copies may be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue, N.W., Washington DC 20006. Definitions are identified by the symbol (E) after the definition.

The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

A

access time. The time interval between the instant at which a call for data is initiated and the instant at which the delivery of data is completed.

adapter. A printed circuit board that modifies the system unit to allow it to operate in a particular way.

address. (1) A value that identifies a register or a particular part of storage. The value is represented by one or more characters. (2) The location in the storage of a computer where data is stored. (3) To refer to a specific storage location by specifying the value that identifies the location.

address bus. The path used for the transmission of address information in a computer.

analog. (1) Pertaining to data consisting of continuously variable physical quantities. (T) (2) Contrast with digital, discrete.

ANSI. American National Standards Institute. An organization consisting of producers, consumers, and general interest groups, that establishes the procedures by which accredited organizations create and maintain voluntary industry standards in the United States.

application. The use to which an information processing system is put; for example, a payroll application, an airline reservation application, a network application.

application program. (1) A program that is specific to the solution of an application problem. Synonymous with application software. (T) (2) A program written for or by a user that applies to the user's work, such as a program that does inventory control or payroll. (3) A program used to connect and communicate with stations on a network, enabling users to perform application-oriented activities.

arbitration. A process that determines which device or subsystem gains control of a bus when two or more devices or subsystems simultaneously compete for control.

architecture. See computer architecture.

B

bank. An aggregation of similar devices, such as transformers or lamps, connected to each other and used cooperatively.

benchmark. A point of reference from which measurements can be made.

benchmark test. A test that uses a representative set of programs and data designed to evaluate the performance of computer hardware and software in a given configuration. (T)

bit. Either of the digits 0 or 1 when used in the binary numeration system. Synonymous with binary digit. (T)

bridge. A functional unit that interconnects two local area networks that use the same logical link control protocol but may use different medium access control protocols.

buffer. (1) A routine or storage used to compensate for a difference in rate of flow of data, or time of occurrence of events, when transferring data from one device to another. (A) (2) A portion of storage used to hold input or output data temporarily.

bus. One or more conductors used for transmitting signals, data, or power. See also address bus and data bus.

bus master. A device or subsystem that controls data transfers between itself and a slave.

byte. A string that consists of a number of bits, usually 8, that are treated as a unit and represent a character.

C

cable. The physical medium for transmitting signals; it includes copper conductors and optical fibers.

cache. A buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

client. A functional unit that receives shared services from a server. (T)

clock. A device that generates periodic, accurately spaced signals used for purposes such as timing, regulation of the operations of a processor, or generation of interrupts. (T)

code. A collection of instructions that is in a form that can be read and processed by a computer.

compatibility. The capability of a hardware or software component to conform to the interface requirements of a given computer without adversely affecting its functions.

computer architecture. (1) The logical structure and functional characteristics of a computer, including the interrelationships among its hardware and software components. (2) The organizational structure of a computer system, including hardware and software.

configuration. The manner in which the hardware and software of an information processing system are organized and interconnected. (T)

configure. To set up a computer for operation by describing to the system the devices, optional features, and programs installed in the computer.

control. The determination of the time and order in which the parts of a computer and the devices that contain those parts perform the input, processing, storage, and output functions.

controller. A device that coordinates and controls the operation of one or more input/output devices, such as workstations, and synchronizes the operation of such devices with the operation of the system as a whole.

CSMA/CA. Carrier sense multiple access with collision avoidance. A network protocol in which the transmitting workstation resends the data if the receiving workstation does not confirm receipt of the data within a given period of time.

CSMA/CD. Carrier sense multiple access with collision detection. A network protocol in which the transmitting workstation detects data collisions and waits a random length of time before retrying the transmission.

cycle. (1) An interval of space or time in which one set of events or phenomena is completed. (A) (2) A complete vibration, electric oscillation, or alternation of current.

D

DASD. Direct access storage device.

data. (1) A re-interpretable representation of information in a formalized manner suitable for communication, interpretation, or processing. Operations can be performed upon data by humans or by automatic means. (T) (2) Any representations such as characters or analog quantities to which meaning is or might be assigned. (A)

data bus. A bus used to communicate data internally and externally to and from a processing unit, storage, and peripheral devices. (A)

device. A mechanical, electrical, or electronic piece of equipment designed to serve a special purpose or perform a special function.

digital. (1) Pertaining to data in the form of digits. (A) (2) Contrast with analog.

direct access storage device (DASD). A nonvolatile-storage device, such as a diskette drive, hard disk drive, or CD-ROM drive, in which access time is effectively independent of the location of the data on the storage medium.

direct memory access (DMA). The transfer of data between memory and input/output devices without microprocessor intervention.

disk array. Two or more hard disks interconnected to increase security, performance, or reliability.

disk duplexing. A method of storing data whereby the data from one hard disk is duplicated on another, with each hard disk drive using its own hard-disk controller. Contrast with disk mirroring.

diskette. A small magnetic disk enclosed in a jacket. (T)

disk mirroring. A method of storing data whereby the data from one hard disk is duplicated on another, with both hard disk drives sharing a single hard-disk controller. Contrast with disk duplexing.

DMA. Direct memory access.

E

EEPROM. Electrically erasable programmable read-only memory.

electrically erasable programmable read-only memory (EEPROM). EPROM that can be reprogrammed while it is in the computer.

EPROM. Erasable programmable read-only memory.

erasable programmable read-only memory (EPROM). A PROM that can be erased by a special process and reused. (T)

F

file. A named set of records stored or processed as a unit. (T)

file server. A high-capacity disk storage device or a computer that each computer on a network can use to access and retrieve files that can be shared among the attached computers.

frame. (1) A data structure that consists of fields, predetermined by a protocol, for the transmission of user data and control data. The composition of a frame, especially the number and types of fields, may vary according to the type of protocol. (T)

frequency. The rate of signal oscillation, expressed in hertz.

frequency-division multiplexing (FDM). Division of a transmission facility into two or more channels by splitting the frequency band transmitted by the channel into narrower bands, each of which constitutes a distinct channel. See also time-division multiplexing.

G

gateway. A functional unit that interconnects two computer networks with different network architectures. A gateway connects networks or systems of different architectures. A bridge interconnects networks or systems with the same or similar architectures.

GB. Gigabyte.

gigabyte. (1) For processor storage and real and virtual memory, 2^{30} or 1 073 741 824 bytes. (2) For disk storage capacity, 1 000 000KB. (3) For transmission rates, 1 000 000 000 bytes.

H

hard disk. A rigid magnetic disk such as the internal disks used in the system units of personal-computer systems and in external hard disk drives.

hard disk drive. A disk drive that reads and writes data on rigid disks and can be installed in or connected to a computer.

hardware. (1) All or part of the physical components of an information processing system, such as computers or peripheral devices. (T) (2) The equipment, as opposed to the programming, of a computer. (3) Contrast with software.

hertz (Hz). A unit of frequency equal to one cycle per second.

I

icon. A graphic symbol, displayed on a screen, that a user can point to with a device such as a mouse in order to select a particular function or software application. Synonymous with pictogram. (T)

initialization. Preparation of a system, device, or program for operation.

input/output. Pertaining to a device, process, or channel involved in data input, data output, or both.

input/output device. A device in a data processing system by means of which data can be entered into the system, received from the system, or both. (I) (A)

instruction. A statement that specifies an operation to be performed by a microprocessor, and that identifies data involved in the operation.

I/O. Input/output.

ISO. International Organization for Standardization. An organization of national standards bodies from various countries established to promote development of standards to facilitate international exchange of goods and services, and develop cooperation in intellectual, scientific, technological, and economic activity.

K

KB. Kilobyte.

kilobyte. 2^{10} or 1024 bytes.

L

LAN. See local area network.

load. To bring all or part of a computer program into memory from auxiliary storage so that the computer can run the program.

local area network (LAN). (1) A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary may be subject to some form of regulation. (T) (2) A network in which a set of devices are connected to one another for communication and that can be connected to a larger network.

logical. (1) Pertaining to content or meaning as opposed to location or actual implementation. (A)
(2) Pertaining to a view or description of data that does not depend on the characteristics of the computer system or the physical storage. (A) (3) Contrast with physical. (A)

M

math coprocessor. In personal-computer systems, a microprocessor that supplements the operations of the system microprocessor, enabling the computer to perform complex mathematical operations in parallel with other operations.

MB. Megabyte.

Mbps. One million bits per second.

medium. A physical material in or on which data may be represented.

megabyte. (1) For processor storage and real and virtual memory, 2^{20} or 1 048 576 bytes. (2) For disk storage capacity and transmission rates, 1 000 000 bytes.

megahertz (MHz). A unit of measure of frequency equal to 1 000 000 cycles per second.

memory. Addressable storage space in the computer that is used for temporary storage of instructions and data while a program is running, or for permanent storage of microcode. Contrast with auxiliary storage.

menu. A list of options displayed to the user by a data processing system, from which the user can select an action to be initiated. (T)

microprocessor. A processor whose elements have been miniaturized into one or a few integrated circuits. (T)

millimeter (mm). One thousandth of a meter.

millisecond (ms). One thousandth of a second.

milliwatt (mw). One thousandth of a watt.

modem (modulator/demodulator). (1) A functional unit that modulates and demodulates signals. One of the functions of a modem is to enable digital data to be transmitted over analog transmission facilities. (T) (A) (2) A device that converts digital data from a computer to an analog signal that can be transmitted on a telecommunication line, and converts the analog signal received to data for the computer.

multiplexing. In data transmission, a function that permits two or more data sources to share a common transmission medium so that each data source has its own channel. (I) (A)

multitasking. A mode of operation that provides for concurrent performance, or interleaved execution, of two or more tasks. (I) (A)

N

nanosecond (ns). One thousand millionth (10^{-9}) of a second.

network. (1) An arrangement of nodes and connecting branches. (T) (2) A configuration of data processing devices and software connected for information interchange.

network server. See server.

nonvolatile. (1) Pertaining to a storage device whose contents are not lost when power is cut off. (T) (2) Contrast with volatile.

normalization. The process of restructuring a relation for the purpose of reducing it to its simplest form, so that each of its attributes is based on a simple domain that consists of single, noncomposite values. (A)

ns. nanosecond.

O

operating system. Software that controls the execution of programs and that may provide services such as resource allocation, scheduling, input/output control, and data management. Although operating systems are predominantly software, partial hardware implementations are possible. (T)

P

parity bit. A check bit appended to an array of binary digits to make the sum of all the binary digits, including the check bit, always odd or always even. (A)

performance. One of the two major factors, together with facility, on which the total productivity of a system depends. Performance is largely determined by a combination of throughput, response time, and availability.

physical. (1) Pertaining to actual implementation or location as opposed to conceptual content or meaning. (A) (2) Contrast with logical. (A)

port. An access point for data entry or exit.

processing. The performance of logical operations and calculations on data, including temporary retention of data in microprocessor storage while the data is being operated on.

processor. A functional unit that interprets and executes instructions. A processor consists of at least an instruction control unit and an arithmetic and logic unit. (T) See microprocessor and central processing unit.

program. (1) A sequence of instructions that a computer can interpret and execute. (2) To design, write, modify, and test computer programs. (I) (A)

programmable read-only memory (PROM). A storage device that, after being written to once, becomes read-only memory. (T) (A)

PROM. Programmable read-only memory.

prompt. A visual or audible message sent by a program to request the user's response. (T)

R

random access memory (RAM). (1) A storage device in which data can be written and read. (2) A storage device into which data is written and from which data is read in a nonsequential manner.

read. To acquire or interpret data from a storage device, from a data medium, or from another source.

read-only memory (ROM). Memory in which stored data cannot be modified by the user except under special conditions. See also EEPROM, EPROM, and PROM.

record. (1) A set of data treated as a unit. (2) A set of one or more related data items grouped for processing.

refresh. (1) To recharge a memory location in volatile memory with an electric current so that it retains a state or binary value. (2) In computer graphics, the process of repeatedly producing a display image on a display surface so that the image remains visible.

register. (1) An integrated circuit that contains 8, 16, or 32 storage locations, each of which can store 1 bit of binary data. See also binary. (2) An area that stores binary data while it is being processed by the computer.

S

SCSI. Small computer system interface.

seek time. The time required for the access arm of a direct access storage device to be positioned on the appropriate track.

sequential. Pertaining to a process in which all events occur one after the other, without any time lapse between them.

sequential access. The capability to enter data into a storage device or a data medium in the same sequence as the data is ordered, or to obtain data in the same order as it has been entered.

server. (1) A functional unit that provides shared services to workstations over a network. (2) In a network, a data station that provides facilities to other stations.

small computer system interface (SCSI). A standard input/output interface used by personal computers.

socket. A receptacle for a microchip.

software. (1) All or part of the programs, procedures, rules, and associated documentation of a computer. Software is an intellectual creation that is independent of the medium on which it is recorded. (2) Contrast with hardware.

SQL. Structured query language.

startup sequence. In personal-computer systems, the order that the computer uses to search the direct access storage devices for an operating system.

storage. A functional unit into which data can be placed, in which it can be retained, and from which it can be retrieved.

subsystem. In computers, a secondary or subordinate system, usually capable of operating independently of a controlling system, and usually having a single purpose, such as displaying video or reading from and writing to hard disks. A subsystem can be integrated into the system board or on an adapter.

system board. In a system unit, the main circuit board that supports a variety of basic system devices, such as a keyboard or a mouse, and provides other basic system functions.

system unit. In personal-computer systems, the part of the computer that contains the processor circuitry, read-only memory (ROM), random access memory (RAM), and the I/O channel.

T

TB. Terabyte.

terabyte. (1) For processor storage and real and virtual memory, 2^{40} or 1 099 511 627 776 bytes. (2) For disk storage capacity, 1 000 000 000KB. (3) For transmission rates, 1 000 000 000 000 bytes.

throughput. A measure of the amount of work performed over a period of time, for example, number of jobs per day.

time-division multiplexing. Division of a transmission facility into two or more channels by allotting the common channel to several different information channels, one at a time. See also frequency-division multiplexing.

token. In a local area network, the symbol of authority passed successively from one data station to another to indicate the station temporarily in control of the transmission medium. Each data station has an opportunity to acquire and use the token to control the medium. A token is a particular message or bit pattern that signifies permission to transmit.

transaction. An exchange between a workstation and another device that accomplishes a particular action or result.

U

utility program. (1) A computer program in general support of computer processes; for example, a diagnostic program, a trace program, a sort program. (2) A program designed to perform an everyday task such as copying data from one storage device to another.

V

virtual. Pertaining to a functional unit that appears to be real, but whose functions are accomplished by other means.

volatile. (1) Pertaining to a storage device whose contents are lost when power is cut off. (2) Contrast with nonvolatile.

W

workstation. (1) A functional unit at which a user works. A workstation often has some processing capability. (2) A terminal or microcomputer, usually one that is connected to a mainframe or to a network, at which a user can perform applications.

write. To make a permanent or transient recording of data in a storage device or on a data medium.

Index

C

D

E

F

G

H

I

N

O

P

Q

R

S

T

U

V

W

X

IBM

Part Number: 83G9764

Printed in U.S.A.

September 1994

83G9764