



PC Server 520

User's Reference

Note

Before using this information and the product it supports, be sure to read the general information under Appendix, "Product Warranties and Notices" on page 105.

Second Edition (May 1996)

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 reseller or IBM marketing representative.

© Copyright International Business Machines Corporation 1996. All rights reserved.

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

About This Book	v
How This Book is Organized	v
Related Publications	vii
Chapter 1. Introducing Your IBM PC Server 520	1
Features at a Glance	2
What Your IBM PC Server 520 Offers	3
Controls and Status Indicators	6
Expansion Bays	8
Expansion Slots	10
Input/Output Ports	12
International Capabilities	16
Operating System Support	16
Reliability, Availability, and Serviceability	17
Specifications	18
Chapter 2. Understanding Server Performance	19
Integrated Systems Management	19
Performance Tuning	19
Server Performance	20
Estimating Server Performance and Capacity	23
Improving Server Performance	23
Chapter 3. Microprocessors and Memory	25
Microprocessor	25
Memory Subsystem	28
Chapter 4. Bus Architecture	33
PCI Bus Overview	33
EISA Bus Overview	34
Micro Channel Bus Overview	34
PCI and Micro Channel Shared Features	35
Features Unique to the Micro Channel Bus	39
Chapter 5. Data Storage Devices	41
IBM Direct Access Storage Devices	41
Tape Drives	43
IBM External SCSI Storage Enclosures	44

Chapter 6. The SCSI Subsystem	45
Advantages of the SCSI Subsystem	45
Types of Interfaces	47
SCSI Physical Configuration	48
Cabling	51
Termination	52
Chapter 7. Disk-Array Subsystems	53
Internal Disk-Array Subsystem	53
RAID Technology	53
Disk Array Adapter Features	62
Chapter 8. Configuration	65
Configuration Programs	66
Configuration Process Overview	71
Chapter 9. Input/Output Ports	77
Serial Ports	77
Parallel Port	79
Video Port	79
Keyboard and Auxiliary-Device Ports	80
SCSI Ports	81
Chapter 10. Video Subsystem	85
Resolution and Dot Pitch	86
Vertical-Refresh Rate	87
Selecting a Monitor	88
Screen-Saver Programs	89
Brightness and Contrast Controls	89
Avoiding Fatigue	90
Chapter 11. Security Features	91
Securing Hardware	91
Securing Data	94
Securing Data from Loss	99
Redundant Power Supply	102
Backup Power Supply	103

Appendix. Product Warranties and Notices 105
Warranties 105
Notices 108

Glossary 113

Index 127

About This Book

This *User's Reference* provides 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 520.

This reference does not provide specific instructions for performing procedures. That information is in your *PC Server 520 User's Handbook*.

How This Book is Organized

Chapter 1, "Introducing Your IBM PC Server 520," provides an overview of the features and expansion capabilities of your server. This chapter also includes the PC Server 520 specifications.

Chapter 2, "Understanding Server Performance," contains information about server performance, the advantages of using a LAN, and tuning your server for optimal performance in your environment.

Chapter 3, "Microprocessors and Memory," contains general information about the microprocessors and types of memory in the PC Server 520 models.

Chapter 4, "Bus Architecture," contains information about the PC Server 520 dual bus architecture, which provides PCI/Micro Channel and PCI/EISA buses for high performance and a variety of expansion options. This chapter provides a general description of each bus, including information about direct memory access (DMA), arbitration, burst data transfers, and the streaming-data procedure used in your server.

Chapter 5, "Data Storage Devices," explains PC Server 520 data storage capabilities and includes information about the storage devices that you can install, such as diskette drives, hard disk drives, CD-ROM drives, and tape drives.

Chapter 6, "The SCSI Subsystem," describes the SCSI (small computer system interface) subsystem.

Chapter 7, “Disk-Array Subsystems,” contains information about the disk-array subsystem shipped in some PC Server 520 models. It describes the different RAID (redundant array of independent disks) levels as well as the advantages of each.

Chapter 8, “Configuration,” provides information about the configuration process. It also contains a description of the configuration programs provided with your PC Server 520.

Chapter 9, “Input/Output Ports,” provides basic information about the serial, parallel, keyboard, mouse, video, and SCSI-2 input/output (I/O) connectors on the PC Server 520 models. Pin assignments for each I/O connector are provided to help you determine whether a device is compatible with your server.

Chapter 10, “Video Subsystem,” contains information about the super video graphics array (SVGA) controller (or adapter) and video device drivers. This chapter also includes detailed information about some of the displays that you can attach to your server.

Chapter 11, “Security Features,” describes how to secure the server hardware and network data stored on the server, and the selectable-drive startup feature.

An Appendix, “Product Warranties and Notices,” 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 follows the glossary.

Related Publications

The Setup sheet provided with this server contains information about unpacking and setting up your server.

The *User's Handbook* provided with this manual contains detailed information about 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 *IBM Personal System/2 Hardware Maintenance Manual Supplement* is available for purchase. It contains a parts catalog, error codes, and advanced diagnostic procedures for PC Servers. This manual is intended for trained service technicians. (Diagnostic diskettes are not included.)

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 reseller or marketing representative.



Chapter 1. Introducing Your IBM PC Server 520

The IBM PC Server 520 is designed to meet your needs today and provide plenty of expansion capabilities for the future. Your system is a high-performance server, where microprocessor performance, memory, system expansion, flexibility, and large amounts of data storage are high priorities.

Your PC Server 520 is 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. Hot swapping is standard on all models of the PC Server 520.

Your server is built around a powerful Intel Pentium microprocessor. You can install an optional second Pentium microprocessor. Your server also supports advanced 32-bit operating systems, such as the IBM Operating System/2 (OS/2) operating system and Microsoft Windows NT. With an operating system that supports *symmetric multiprocessing (SMP)*, a second Pentium microprocessor can increase server performance.

As always, your IBM server meets stringent worldwide certifications for power, EMC (electromagnetic compatibility), and safety.

Features at a Glance

<p>Microprocessor</p> <ul style="list-style-type: none"> • Intel Pentium microprocessor with 16 KB of internal cache; clock rate varies by model • A processor upgrade, which adds another Pentium microprocessor, is available for symmetric multiprocessing <p>Cache Memory</p> <ul style="list-style-type: none"> • On each processor: 16 KB of level-1 cache • 512 KB of level-2 cache <p>Memory</p> <ul style="list-style-type: none"> • Error correcting code on SIMM (EOS), 70 ns • 32 MB minimum of EOS memory, expandable to 256 MB • Eight 72-pin, single inline memory-module (SIMM) sockets <p>Diskette Drives</p> <ul style="list-style-type: none"> • Standard: One 3.5-inch, 1.44 MB • Optional (internal): <ul style="list-style-type: none"> – 3.5-inch, 2.88 MB – 3.5-inch, 1.44 MB – 5.25-inch, 1.2 MB <p>Hard Disk Drives</p> <ul style="list-style-type: none"> • Number and capacity of standard drives varies by model • Supports up to 20 internal hard disk drives, 18 of which are hot swappable <p>CD-ROM Drive</p> <ul style="list-style-type: none"> • Standard: SCSI-2 	<p>Keyboard and Auxiliary Device</p> <ul style="list-style-type: none"> • 101-key keyboard • Mouse <p>Expansion Bays and Slots (PCI/Micro Channel models)</p> <ul style="list-style-type: none"> • Six 32-bit Micro Channel slots, two 32-bit PCI slots • 22 drive bays, 18 of which are hot swappable <p>Expansion Bays and Slots (PCI/EISA models)</p> <ul style="list-style-type: none"> • Five EISA slots, two PCI slots, and one PCI/EISA shared slot • 22 drive bays, 18 of which are hot swappable <p>Video</p> <ul style="list-style-type: none"> • Super video graphics array (SVGA) controller • Compatibility: <ul style="list-style-type: none"> – Video graphics adapter (VGA) – Color graphics adapter (CGA) – Multicolor graphics array (MCGA) <p>Security Features</p> <ul style="list-style-type: none"> • Bolt-down capability • Door and front cover lock • Administrator and power-on passwords (PCI/Micro Channel models) • Supervisor and user passwords (PCI/EISA models) • Selectable drive-startup • Security cover (optional) • LogicLock with active security (PCI/Micro Channel models) 	<p>Integrated Functions</p> <ul style="list-style-type: none"> • LED usability support • Video connector • Two serial connectors • Parallel connector • Auxiliary device connector • Keyboard connector • Battery-backed clock and calendar • SCSI-2 connectors on backplane • SCSI-2 (PCI/Micro Channel models) <ul style="list-style-type: none"> – One 16-bit external connector – Two internal connectors (one 8-bit and one 16-bit) <p>SCSI-2 Controller</p> <ul style="list-style-type: none"> • IBM PC ServeRAID Adapter (disk-array models only) <ul style="list-style-type: none"> – Three 16-bit internal connectors – One 16-bit external connector – Support for RAID levels 0, 1, and 5 • IBM SCSI-2 Fast/Wide PCI Adapter (PCI/EISA non-disk-array models) <p>Power Supply</p> <ul style="list-style-type: none"> • 434 watt with automatic range voltage selection (115–230 V ac) • Built-in overload and surge protection • 220 watt upgrade • 780 watt redundant power supply upgrade <p>Upgradable POST and BIOS</p> <ul style="list-style-type: none"> • Upgradable EEPROM on the system board • POST/BIOS upgrades (when available)
--	---	---

What Your IBM PC Server 520 Offers

The PC Server 520 offers impressive performance using the Pentium microprocessor. You can install an optional second Pentium microprocessor for increased performance when used with an operating system supporting SMP. Your server also provides large data-storage capabilities, enhanced system expandability, and SVGA video graphics. This system is designed to be used as a powerful, cost-effective server.

The PC Server 520 takes advantage of advancements in microprocessor speed, memory capacity, application programs, communication devices, networks, and operating environments.

IBM ServerGuide

The IBM ServerGuide package contains several compact discs (CDs). These CDs contain menu-driven programs that make it easy for you to install your operating system, network program, and systems-management software. These CDs also contain a performance tuning program. Refer to the documentation packaged with the ServerGuide CDs for more detailed information.

IBM NetFinity

The IBM ServerGuide package gives you the option of installing IBM NetFinity. NetFinity is an integrated systems-management program that enables you to use several of the built-in systems-management features in your server. To find out more about NetFinity and how it can help you, view the demonstration program provided in the ServerGuide package and read its accompanying information.

Getting Help on the World Wide Web

You can access the latest information about product compatibility and configuration on the World Wide Web. The information is updated as new PC Server products are announced. For information, refer to the *Getting Help, Service, and Information* section in your *User's Handbook*.

IBM PC Server Startup Support

The IBM PC Server Startup Support program provides comprehensive telephone assistance 24 hours a day, seven days a week during your first 90 days after installation.*

IBM provides you direct access to trained specialists who can assist you with the setup, installation, and configuration of your server. Assistance is available for IBM and non-IBM network operating systems, network interface adapters, and other optional peripherals. To receive a current listing of the network products supported by the IBM PC Server Startup Support program, call the IBM Automated Fax System in the U.S. at 1-800-426-3395 and ask for document number 16125. For more information about this exceptional program, or for assistance with the installation of your server:

- In the U.S., call 1-800-772-2227
- In Canada, call 1-800-565-3344
- In all other countries, contact your IBM reseller or marketing representative.

Ethernet PeerMaster Adapter

Some PC Server 520 models come with a preinstalled IBM Ethernet Quad-BT PeerMaster Adapter. This 10Base-T Ethernet adapter (hereafter referred to as PeerMaster adapter) provides the ability to connect the server to four separate 10Mbps Ethernet LAN segments.

* Response time may vary depending on the number and nature of calls received.

This adapter:

- Performs high performance data transfers between the host PC and adapter
- Provides high performance data transfers among ports on an adapter
- Provides high performance data transfers among ports across multiple adapters with no host processor intervention
- Provides a cost efficient, high performance means to forward cross-segment traffic without impacting server performance
- Performs 64-bit data transfers across the Micro Channel at burst rates of 640 Mbps, which equals the total bandwidth of 64 Ethernet ports

Instructions for configuring and cabling the PeerMaster adapter are in your *User's Handbook*.

PCI/Micro Channel Models


Combining the Micro Channel and peripheral component interconnect (PCI) bus architectures, the PC Server 520 provides compatibility with a wide range of existing and future hardware devices and software applications. These bus architectures define the size, physical arrangement, and logical relationships of the connectors. The PC Server 520 dual bus architecture provides a high speed, high throughput PCI local bus and the advanced Micro Channel expansion bus.


PCI/EISA Models

Using both peripheral component interconnect (PCI) and extended industry standard architecture (EISA), the PC Server 520 provides compatibility with a wide range of hardware devices and software applications. The EISA expansion slots support both EISA and industry standard architecture (ISA) devices.

Controls and Status Indicators

The most commonly used controls and indicators are on the front of your server.





Power On/Off Switch: Press this to turn your server on or off. Do not turn off your server if any drive activity indicator is on. This might damage the information stored on a hard disk or diskette.

Power OK Indicator: This lights when you turn on your server.


Hard Disk Drive Activity Indicator: This lights when the drive is accessing a disk.

CD-ROM Tray Release Button: Push this to release a CD-ROM from the drive.

Diskette Eject Button: Push this to release the diskette from the drive.

Expansion Bays

Your server comes with one 3.5-inch, 1.44 MB diskette drive. The number of hard disk drives that come with your server varies by model; some models do not come with a hard disk drive.



Open Bays: Your server has a special design that gives you the flexibility to use up to 18 hot-swap, 3.5-inch hard disk drives. There is additional open space in the top left bay for two more diskette drives or an additional SCSI drive. For installation instructions and information on the types of drives that you can install in each bay, see your *User's Handbook*.

CD-ROM Drive: All models come with a CD-ROM drive.

Diskette Drive: The 3.5-inch, 1.44 MB diskette drive uses 1 MB and 2 MB diskettes.

Hard Disk Drives: The number of preinstalled drives and their capacities vary by model.


Note: See Chapter 5, "Data Storage Devices," for more information about data storage devices you can install in your server. For the latest information on available options, contact your IBM reseller or marketing representative.

Expansion Slots

The PCI/Micro Channel model of the PC Server 520 has six Micro Channel expansion slots and two PCI slots. The PCI/EISA model has six expansion slots and three PCI slots. On disk-array models, a PC ServeRAID Adapter is installed in one of the slots. Some Micro Channel models come with a PeerMaster adapter preinstalled in a Micro Channel slot.


The remaining slots are available for future expansion. 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.

PCI/Micro Channel Model



PCI/EISA Model

On the PCI/EISA model of the PC Server 520, the EISA and PCI slots closest to each other are called *shared* slots, meaning you can install one EISA or PCI device in either slot, but not both at the same time. (See the following illustration.)



Note: You also can install ISA devices in the slots designated as EISA slots.


The video adapter is installed in an EISA slot. Your server comes with a SCSI-2 Fast/Wide PCI Adapter or an IBM PC ServerRAID Adapter, depending on model. This adapter is installed in a PCI slot.

Input/Output Ports

PCI/Micro Channel and PCI/EISA models provide the same types of input/output ports (connectors). However, the arrangement of the connectors, which are located on the rear panel of your server, depends on which model you have.

PCI/Micro Channel Model

The following illustration shows the input/output ports (connectors).



Note: Some Micro Channel models come with a preinstalled PeerMaster adapter in a Micro Channel expansion slot.

Power Connector: The power cable connects here.



SCSI-2 Connector: Attach external SCSI devices here.



Monitor Connector: The display signal cable attaches to this connector.



Parallel Connector: This is where the signal cable for a parallel printer or other parallel device connects to your server.



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 connector sometimes is called an auxiliary-device or pointing-device connector.



Keyboard Connector: The cable from your keyboard connects here.

PCI/EISA Model

The following illustration shows the input/output ports (connectors).



Note: Some PCI/EISA models come with a preinstalled PeerMaster adapter in a PCI expansion slot.

Power Connector: The power cable connects here.



Keyboard Connector: The keyboard cable connects here.



Mouse Connector: The mouse cable connects here. This connector also is called an auxiliary-device or pointing-device connector.



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.



Parallel Connector: One 25-pin parallel connector is provided. This is where the signal cable for a parallel printer or other parallel device connects to your server.



SCSI Connector: External SCSI devices attach here.



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

International Capabilities

The IBM PC Server 520 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 many languages.



Operating System Support

You can use the PC Server 520 in a variety of environments to do many different tasks. To meet the requirements of these environments, your server supports several different operating systems. Some operating systems, such as OS/2 2.1 or later, provide support for greater than 16 MB of memory. Check your operating-system documentation or contact your IBM reseller or marketing representative for more information.

OS/2 for SMP supports the dual-processing capability of your server. A second microprocessor is available as an upgrade.

Reliability, Availability, and Serviceability

Three of the most important factors in server design are *reliability*, *availability*, and *serviceability* (RAS). These factors are measurements that help you determine how well your server will: protect the integrity of the data stored on your server; be available when you want to use it; and enable you to diagnose and repair with minimal inconvenience any failure that occurs.

The following is a list of some of the RAS features of your PC Server 520. Many of these features are described in later chapters of this book, or in your *User's Handbook*.

- Menu-driven configuration programs
- Power-on self-test (POST)
- 24-hour customer support center
- Disk mirroring (some models)
- Hot-swap drive bays
- Menu-driven diagnostic programs
- Error codes and messages
- Upgradable POST and BIOS code
- Error correcting code on SIMM (EOS)
- Upgradable redundant power supply
- Parity checking on the SCSI bus, Micro Channel bus, keyboard interface, and serial connectors
- Vital product data (VPD) including serial number information and replacement part numbers, stored in nonvolatile memory

Specifications

<p>Size</p> <ul style="list-style-type: none"> • With pedestal: <ul style="list-style-type: none"> - Depth: 755 mm (29.7 in.) 835 mm (32.8 in.) with cable cover - Height: 622 mm (24.5 in.) - Width: 353 mm (13.9 in.) <p>Weight</p> <ul style="list-style-type: none"> • Minimum configuration as shipped: 31.29 kg (69 lb) <p>Environment</p> <ul style="list-style-type: none"> • Air temperature: <ul style="list-style-type: none"> - 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.) • Humidity: <ul style="list-style-type: none"> - System on: 8% to 80% - System off: 8% to 80% • Maximum altitude: 2133 m (7000 ft) 	<p>Electrical Input</p> <ul style="list-style-type: none"> • Sine-wave input (50 to 60 Hz) is required • Input voltage: <ul style="list-style-type: none"> - Low range: <ul style="list-style-type: none"> - Minimum: 90 V ac - Maximum: 137 V ac - High range: <ul style="list-style-type: none"> - Minimum: 180 V ac - Maximum: 265 V ac - Input kilovolt-amperes (kVA) approximately: <ul style="list-style-type: none"> - Minimum configuration as shipped: 0.16 kVA - Maximum configuration: 1.0 kVA <p>Total Power Available for Drives</p> <ul style="list-style-type: none"> • Nominal Operating Current allowed: <ul style="list-style-type: none"> - +5 V dc line: 10 A base; 23 A with upgrade - +12 V dc line: 10 A base; 23 A with upgrade 	<p>Heat Output</p> <ul style="list-style-type: none"> • Approximate heat output in British Thermal Units (BTU) per hour: <ul style="list-style-type: none"> - Minimum configuration: 150 BTU (44 watts) - Maximum configuration: 3600 BTU (1055 watts) <p>Acoustical Noise Emission Values</p> <ul style="list-style-type: none"> • Average sound pressure levels at operator position: <ul style="list-style-type: none"> - 1 hard disk drive: <ul style="list-style-type: none"> - 43 db operating - 42 db idle • Average sound pressure at bystander position (1 meter): <ul style="list-style-type: none"> - 1 hard disk drive: <ul style="list-style-type: none"> - 38 db operating - 38 db idle • Declared (upper limit) sound power levels: <ul style="list-style-type: none"> - 1 hard disk drive: <ul style="list-style-type: none"> - 5.6 bels operating - 5.6 bels idle
--	--	--

Chapter 2. Understanding Server Performance

Your productivity and satisfaction, as well as everyone else's on your network, depend upon server performance. Several 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 includes suggestions for improving server performance.

One of the main capabilities of a local area network (LAN) 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 improve work flow, work quality, 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.

Integrated Systems Management

The SystemView package that comes 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 demonstration program and read the information in the SystemView package.

Performance Tuning

A performance tuning feature available on the ServerGuide CD (provided with your server) enables you to 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:

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

Network Operating System

Always check the documentation that comes with the network operating system to determine which parameters affect performance and which, if any, parameters you can change to increase performance. The original parameter settings provided by most network operating systems typically support a network with 12 to 18 concurrent users. If 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 good practice to record the original settings before making any changes. If performance decreases, you can return the settings to their original values. It is also 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. Such 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 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 “Network Operating System,” it is 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.

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. Bus-master network adapters are more suited for use in network servers. They provide higher throughput rates because they handle most of the network processing and leave the system microprocessor free for other tasks.

The appropriate network adapter for your server depends on the type of network your server is attached to (broadband, baseband, token-ring, or Ethernet). IBM and other companies offer baseband, broadband, token-ring, and Ethernet network adapters for PCI and Micro Channel applications.

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. However, 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.

For a current list of the network adapters that are available for your model, contact your 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. Additional memory might be needed to support additional application programs, buffers, or control blocks.

System Microprocessor

One way to improve system microprocessor performance is to lighten its load by using bus-master adapters. Bus-master adapters handle tasks, such as the network processing or data transfers, with minimal processor involvement.

You can install an additional Pentium microprocessor in the socket next to the system microprocessor. Used with an operating system that supports symmetric multiprocessing (SMP), a second microprocessor can increase server performance.

Hard Disk Subsystem

Your server has a SCSI-2 bus-master controller on the system board or on an adapter. Disk-array models also have a PC ServeRAID Adapter. These controllers significantly increase 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 SCSI-2 controller. Also, you can install additional SCSI controllers for further expansion. You also can use SCSI storage enclosures to house additional devices.

Disk arrays use multiple hard disk drives that provide faster input and output rates than a single large-capacity drive, because files are stored on individually addressable disks and can be found more quickly during read requests. The amount of performance improvement depends upon the application programs that you run on the server and the type of disk array. (See Chapter 7, “Disk-Array Subsystems” on page 53 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 approach follows:

1. Define the functional requirements of your server for your working environment.
2. Define what you consider to be adequate performance, such as response times, transaction rates, and number of users.
3. Understand what factors affect server performance.
4. Determine an approximate configuration of applications for your work environment.
5. 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.
6. Ensure that the performance of the server will still be acceptable if a component 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 suit your environment; concentrate on obtaining improvements for those applications. It is possible that some changes might improve the performance of one application, but degrade the performance for others. 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 the bottleneck is the hard disk subsystem, 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. Microprocessors and Memory

The processing center of your server contains devices that run programs and store data used by programs. These devices include the microprocessor and the components of the memory subsystem.

Microprocessor

The *microprocessor* is an integrated circuit that performs most of the control and computing functions of your server. An Intel Pentium microprocessor is installed in your IBM PC Server 520.

The system board is equipped with a socket for an optional second Pentium microprocessor. You can increase server performance by installing a second microprocessor and an operating system that supports *symmetric multiprocessing (SMP)*, such as OS/2 for SMP.

Note: Operation of a dual processor system is transparent to adapters and application software. An operating system supporting SMP is not required if the second microprocessor is installed. If an optional second processor is installed and the operating system does not support SMP, the server operates as a single-microprocessor (uniprocessor) machine.

Microprocessor Clock

The rate at which the microprocessor performs its basic operations is called the *clock speed*. Examples of basic operations include fetching an instruction from memory, decoding it, executing it, writing the result to memory, and so on. The process of executing an instruction always consists of several basic operations. 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. Clock speed is measured in megahertz (MHz).

A *clock cycle* is the time unit that measures how long the microprocessor takes to perform a basic operation at a given clock speed. Clock cycles are measured in nanoseconds (ns) and are inversely related to the clock speed. For example, at a clock speed of 25 MHz, one clock cycle takes 40 ns; at a clock speed of 100 MHz, one clock cycle takes 10 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 internal microprocessor operations.

Microprocessor Performance

The Pentium microprocessor gains its high performance through its *superscalar* design. The microprocessor contains three *execution units*; each execution unit can issue an instruction in one clock cycle. Each execution unit contains a *pipeline*. In a pipeline, multiple instructions are in various stages of execution. For example, while one instruction is fetched, two are decoded, another is executed, and the results of another are written to memory. This parallelism potentially speeds up instruction execution by a factor of 15 when compared to a processor that is running at the same speed, but that does not use pipelining and multiple execution units.

In practice, the speedup is much less. In many clock cycles, the microprocessor does not issue an instruction. Memory accesses, even to very fast cache memory, take multiple cycles to finish. Some instructions spend many clock cycles in the execution stage of a pipeline. Still, performance is substantially improved.

Microprocessor Registers and Interfaces

The microprocessor needs data or instructions, or both, for each basic operation. Instructions and data are copied from memory into data-storage locations, known as *registers*, in the microprocessor. The microprocessor can read from and write to its registers very quickly. Registers also store the results of processing operations until the results are needed for other operations or are written to memory.

The microprocessor *data interface* is the connection between the microprocessor and the data bus. The data interface of the microprocessor is 64 bits wide.

In PCI/Micro Channel models, both expansion buses in the server are 32 bits wide. These models support 32-bit PCI and Micro Channel adapters and 16-bit Micro Channel adapters. The 32-bit adapters transfer 32 bits at a time; the 16-bit adapters transfer 16 bits at a time.

In PCI/EISA models, the PCI expansion bus is 32 bits wide and the EISA bus is 16 bits wide. These models support 32-bit PCI adapters and 16-bit and 8-bit EISA adapters. The 32-bit PCI adapters transfer 32 bits at a time; the 16-bit EISA adapters transfer 16 bits at a time, and the 8-bit EISA adapters transfer 8 bits at a time.

The set of instructions that the microprocessor can perform determines whether a particular program will run in the server. Programs 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 connects the microprocessor and the address bus, determines the width of the address bus and the amount of physical memory the microprocessor can address. The Pentium microprocessor in your server has a 32-bit address interface and can address up to 4 GB of physical memory. (GB equals approximately 1 000 000 000 bytes.)

Real Mode and Protected Mode Operation

The microprocessor has a *real mode* with a 1 MB address space. (MB equals approximately 1 000 000 bytes.) When the microprocessor is in the real mode, your server can run programs written for earlier models of the IBM personal computer. The real mode is the default condition for the microprocessor; the microprocessor is always in the real mode after the server is turned on and before any operating system is started. The real mode does not support virtual memory (see “Physical Memory and Virtual Memory” on page 29).

The microprocessor also has a *protected mode*, which supports virtual memory and allows the microprocessor to address more than 1 MB of memory. In protected mode 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 you install determines 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 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 uses the virtual memory and multitasking capabilities provided by the protected mode of the microprocessor.

Note: The OS/2 operating system 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.

Numeric Computing

The microprocessor in your server has an execution unit that executes only floating-point arithmetic and mathematical instructions. This execution unit, called the *floating-point unit (FPU)*, executes the types of instructions previously executed by the separate math coprocessors used in earlier personal computers. These instructions, generally called *numeric instructions*, perform basic floating-point arithmetic and a variety of mathematical functions.

The control circuits of the microprocessor pass data and instructions to the FPU, which then performs the mathematical operations. The two other execution units, called *integer units*, can execute instructions concurrently with the FPU. The FPU is pipelined, so several floating point instructions can be in various stages of execution simultaneously.

Programs that require extensive mathematical calculations run much faster when the FPU performs the calculations, because the FPU is designed specifically for that purpose. A program that is intended to use the FPU must contain specific instructions that are recognized by the FPU. Otherwise, the program will not activate the FPU and the integer units will perform 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* (also called *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 RAM, so no data is retained.

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. It is used to store the power-on password, the supervisor or administrator password, the time and date, and system configuration information.

Read-only memory (ROM) stores programs that the server uses for startup procedures and other internal operations. These programs are permanently encoded in the ROM module. ROM is *nonvolatile* memory, and 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 (or *flash memory*) 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. The disk-array models also have an EEPROM on the ServeRAID adapter that stores disk array configuration information.

Physical Memory and Virtual Memory

Physical memory consists of all the writable memory locations in the server, most of which are in 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. The part of physical memory available for instructions and data that the server uses to run programs is *system memory*.

You can add memory to your server by installing 16 MB or 32 MB single inline memory modules (SIMMs). Your server supports up to 256 MB of 70 ns EOS (error correcting code on SIMM) memory. See your *User's Handbook* for more information about installing additional memory.

Parity memory module kits are supported in the PC Server 520, but must not be used in connection with EOS memory module kits.

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. The portion of the hard disk used for virtual memory is often called *swap space*. The amount of swap space on the disk determines the amount of available virtual memory, which is typically much larger than the actual amount of installed physical memory.

Virtual memory supports very large programs and *multitasking*. In multitasking, multiple programs run simultaneously. The memory requirements of these programs can exceed the amount of available physical memory. The operating system allocates the memory required by each program as virtual memory and manages the available physical memory, swapping portions of each program's address space to the hard disk as required.

The maximum amount of virtual memory that an operating system can allocate is determined by the microprocessor and limited by the amount of available swap space. The microprocessor in your server can address up to 64 TB of virtual memory in the protected mode. (TB equals approximately 1 000 000 000 000 bytes.)

Memory Controller

The memory controller is a device 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.

Cache Memory

Cache memory is high-speed memory that can respond to requests to read from or write to memory without imposing *wait states* on the microprocessor. A wait state is a clock cycle during which the microprocessor suspends processing and waits for a response to a request to read to or write from system memory (a *memory I/O operation*). Wait states cause the server to operate less efficiently than it would if the microprocessor did not suspend processing. Using cache memory is another method of reducing the need for wait states. It improves server performance by temporarily storing frequently used data and instructions in a *cache*, a relatively small buffer between the microprocessor and system memory.

The PC Server 520 system board in your server contains level-1 and level-2 caches. Cache performance depends on the size of the cache and the application programs you are running on your server.

Level-1 Cache

The level-1 cache is implemented as an area of very fast RAM on the microprocessor. This type of cache memory is also referred to as an internal cache. A Pentium microprocessor provides 16 KB of level-1 cache memory: 8 KB for data and 8 KB for instructions. (1 KB equals approximately 1000 bytes.)

The level-1 cache stores the data and instructions most frequently used by the microprocessor in the current execution context. (The execution context is the executable code and associated data used by the program running on the microprocessor.) Because the microprocessor sends many fewer read and write requests to system memory, a level-1 cache enables a microprocessor to execute programs much faster than if it had to access system memory for each request. When the execution context changes, the level-1 cache controller copies new data and instructions into the cache, replacing data and instructions that are no longer needed.

Server performance improves each time the microprocessor finds what it needs in the cache (a *cache hit*). The microprocessor is designed so that the cache hit rate is very high, and the penalty for a *cache miss* is minimized. A cache miss occurs when the information needed by the microprocessor is not available in the cache. In the case of a cache miss, the level-1 cache controller must request the needed data from the level-2 cache or system memory. One or more wait states are imposed on the microprocessor while the data is located. The cache controller on the microprocessor manages the use of the cache so that the number of cache hits far exceeds the number of cache misses.

Level-2 Cache

The system board in your server has a 512 KB level-2 cache. The level-2 cache consists of high-speed static random access memory (SRAM). The level-2 cache provides an intermediate level of storage between the level-1 cache and system memory. Although level-2 cache accesses are slower than level-1 cache accesses, level-2 cache accesses are still much faster than system memory accesses.

In the case of a level-1 cache miss, the microprocessor can access the level-2 cache relatively quickly. 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.

Write-Back Mode

For maximum performance, both PC Server 520 caches are set to the write-back mode. In the write-back mode, the microprocessor updates the cache in just over one clock cycle. After the microprocessor updates the cache, the cache controller updates system memory. The microprocessor is free to perform other operations while the cache controller updates system memory. Because the cache controller updates system memory, which is a relatively lengthy process, the microprocessor is free to execute instructions.

Chapter 4. Bus Architecture

A *bus* is an electrical pathway for transferring signals and power among microprocessors, memory, adapters, and I/O devices. The signals include addresses, data, and control information.

Your PC Server 520 provides a dual bus architecture. A peripheral component interface (PCI) bus provides a high-throughput local bus, and an EISA or Micro Channel bus provides an expansion bus that supports PC compatibility and a wide range of adapters.

PCI Bus Overview

The PCI bus is a local bus, intended for interconnecting a system microprocessor with tightly-coupled peripherals and memory. A local bus becomes a system requirement as microprocessor speeds increase. In previous systems, bus performance more closely matched microprocessor performance. As microprocessor performance improves, the microprocessor spends too much time waiting for I/O to move across the relatively slow system bus. A local bus is much more tightly coupled to the system microprocessor, and its speed much more closely matches the microprocessor's.

PCI architecture uses the burst mode for all transfers. PCI supports a 132 megabyte-per-second (MBps) peak transfer rate for 32-bit PCI devices and a 264 MBps peak transfer rate for 64-bit PCI devices. The current design supports PCI bus speeds up to 33 MHz.

Examples of devices that might be attached to a local bus include:

- Video controllers
- Disk controllers
- High-speed network interface adapters

EISA Bus Overview

One of the benefits of EISA is that it allows you to install and use both EISA and ISA adapters and the devices that they support.

An EISA Configuration program, supplied with your server, automatically detects and configures EISA adapters. On most ISA adapters, you must set switches and jumpers. The program assists you in determining the necessary switch and jumper settings.

Data transfers on the EISA bus are synchronized to an 8.33 MHz bus clock signal. On the EISA bus, *if the target device supports EISA burst-mode transfers*, the data transfer requires only one cycle of the bus clock. This results in a data transfer rate of 8.33 Mbps.

Because the data path on the EISA bus is 32-bits wide, four bytes of data (maximum) can be transferred during each data transaction. This allows a maximum data transfer rate of 33.32 Mbps.

ISA Bus Overview

Data transfers on the ISA bus are synchronized to an 8 MHz (more typically 8.33 MHz) bus clock signal. If the target device is a zero-wait-state device, the data transfer requires a minimum of two cycles of the bus clock. This results in a data transfer rate of 4.165 MBps.

Because the data path on the ISA bus is only 16-bits wide, two bytes of data (maximum) can be transferred during each data transaction. This allows a maximum data transfer rate of 8.33 MBps.

Micro Channel Bus Overview

In your server, the Micro Channel bus serves as an expansion bus supporting Micro Channel adapters.

PCI and Micro Channel Shared Features

The PCI and Micro Channel implementations in your PC Server 520 share some fundamental features. Adapters for both buses are automatically configured by your server. Configuration programs and files are used to assign interrupt levels, I/O addresses, memory addresses, arbitration levels, and other resources.

Both buses are implemented as 32-bit buses in your server. Both have 32 signal paths for data; 32 bits of address or data can be sent across each bus simultaneously.

Both buses support burst data transfers; an address and multiple data words (each containing 32 bits) can be sent in each transfer. The address specifies the location of the data to be sent; a bus command specifies the number of data words to send.

Each bus supports data transfers from only one device at a time. A *bus arbitration* mechanism determines which device can access a bus at a given time.

Both buses have several functions; in your server each is used as an expansion bus, an address bus, and a data bus.

Expansion Bus

An expansion bus supports the addition of adapters to add new functions to your server. Adapters plug into connectors, called *expansion slots*, on the system board.

The PCI expansion bus provides two connectors for adapters; the Micro Channel expansion bus provides six connectors. These connectors provide paths for transferring data to and from the adapters.

Address and Data Bus

When a device sends data, it identifies where the data is stored (source) and where it is to be sent (destination). Each location is identified by a unique number, known as an *address*. Each memory location and I/O device that is attached to the bus has an address. The sending device uses the bus as an address bus to send the target address to the receiving device. Then the sending device uses the bus as a data bus to send the data.

The width of a bus determines the size of the address space available to the bus. The PCI and Micro Channel buses each have a 32-bit address bus, enabling devices on the bus to address up to 4 GB of memory.

The width of the data bus determines how much data can be transferred simultaneously. PCI, as implemented in your server, provides a 32-bit data path. The Micro Channel architecture, as implemented in your server, 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 the full width of the 32-bit data bus used.

In your server, the PCI bus connects to the Micro Channel bus through a *bridge*, a hardware component that provides a path between the system microprocessor, the PCI bus, and the Micro Channel bus.

Bus Master Implementation

Both the PCI bus and the Micro Channel bus support *bus masters*. A bus master is an on-board device or controller, or an adapter, that can take control of (own) the bus. The system microprocessor is the first owner of the bus at system start-up. Other bus masters receive ownership of the bus by halting the system microprocessor.

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.

Typical bus masters include:

- Memory and cache controllers
- Disk controllers
- DMA (direct memory access) controllers
- Video controllers and adapters
- Network interface adapters

Once a bus master has control of the bus, it can issue bus commands to other devices on the bus. The target of a bus master, called a *subordinate*, executes bus commands issued by a bus master. Subordinates can initiate service requests, such as interrupts, but subordinates never own the bus.

A bus master can also be a subordinate, but, at a given time, it must be either a bus master or a subordinate. It cannot be both simultaneously.

Bus Arbitration

The PCI and Micro Channel bus architectures support arbitration, a process that determines which of two or more competing devices has control of the bus at any given time. When multiple devices need to own a bus at the same time, these devices participate in *arbitration*.

The PCI and Micro Channel architectures each implement a *central arbitration control point*. This is where contending devices send their arbitration signals. The central arbitration control point does not actually decide which device should control the bus; contending devices make that determination among themselves, using 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.

Every Micro Channel 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 a bus 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.

Both buses have *fairness features* to ensure that all devices get a turn to own the bus, even those devices having low priority levels. 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 Micro Channel 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.

Bus-Parity Checking

The PCI and Micro Channel buses support *bus-parity checking*, 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 is a common feature in 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.

Features Unique to the Micro Channel Bus

The Micro Channel bus provides features not implemented in the PCI bus. These include:

- Direct memory access (DMA)
- Streaming-data procedure
- Channel-check reporting and error logging

Direct Memory Access

Direct memory access (DMA) is used to transfer large amounts of data between system memory and I/O devices without the intervention assistance of 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 need to interrupt its processing activities to manage data transfers.

The *DMA controller* is a bus master that manages data transfers between DMA subordinates and memory subordinates. 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.

Streaming-Data Procedure

The *streaming-data procedure* allows high-speed data transfers between bus masters and subordinates. This procedure supports high-speed transfers of large blocks of data for devices such as disk controllers 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; 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 megabytes 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.

Chapter 5. Data Storage Devices

IBM data storage devices and IBM external storage enclosures are available for the PC Server 520. You can use these storage devices, along with a variety of industry standard storage devices, to expand the data storage capacity of your server.

Note: Some storage devices described in this chapter are not supported in all IBM PC Server 520 models. See your *User's Handbook*, or contact your IBM reseller or marketing representative, for more information.

IBM Direct Access Storage Devices

In a *direct access storage device (DASD)*, the time required to access data does not depend on the location of the data on the storage medium. These devices include diskette drives, hard disk drives, CD-ROM drives, and rewritable optical drives.

Diskette Drives

The PC Server 520 supports three diskette drives: the 3.5-inch 1.44 MB, the 3.5-inch 2.88 MB, and the 5.25-inch 1.2 MB. The 3.5-inch 1.44 MB diskette drive comes as standard equipment. This drive uses 3.5-inch diskettes with a formatted capacity of 720 KB or 1.44 MB. (KB equals approximately 1000 bytes, and MB equals approximately 1 000 000 bytes.) The 3.5-inch 2.88 MB diskette drive uses 3.5-inch diskettes with a formatted capacity of 720 KB, 1.44 MB, or 2.88 MB. To prevent the loss of data, always format 1 MB diskettes to store 720 KB of data, 2 MB diskettes to store 1.44 MB of data, and 4 MB diskettes to store 2.88 MB of data. All of these diskette drives are available for purchase as options.

SCSI Hard Disk Drives

IBM offers a variety of hard disk drives of various capacities and performance capabilities. The number of hard disk drives that your server can support is determined by your operating system. Refer to your operating-system documentation or contact your IBM reseller or marketing representative to determine the number of hard disk drives that your operating system supports.

If your PC Server 520 uses Micro Channel architecture, it has an integrated SCSI controller on the system board. If your PC Server 520 uses EISA architecture, it has a SCSI-2 Fast/Wide PCI Adapter. Both the integrated SCSI controller and the SCSI adapter provide the same support: up to 15 SCSI devices, such as hard disk drives, optical drives, and tape drives. Disk-array models of the server have an IBM PC ServeRAID Adapter that supports up to 45 devices. For more information about the SCSI subsystem and the devices it supports, see Chapter 6, “The SCSI Subsystem” on page 45. All models come with a SCSI CD-ROM drive. Some models come with one or more SCSI hard disk drives.

All models of the server contain hot-swap bays. The hot-swap bays allow you to install up to 18 hot-swappable hard disk drives in your server. Hot-swapping allows you to replace disk drives without turning off the server.

You can install SCSI drives inside the server or in an external storage enclosure. (For more information about SCSI storage enclosures, see “IBM External SCSI Storage Enclosures” on page 44.) The PC Server 520 supports drives that comply with ANSI SCSI Standards X3.131-1986 or X3.131-1994.

SCSI CD-ROM Drive

Compact disc-read only memory (CD-ROM) drives use a compact disc (CD) that the system can read from but not write to. Data can be viewed but not updated or changed. While diskette drives, hard disk drives, and tape drives use a magnetic recording technique, CD-ROM drives use an optical technique that 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 formatted capacities of approximately 680 MB of data.

IBM CD-ROM drives comply with ANSI SCSI Standards X3.131-1986 or X3.131-1994 (SCSI-2). They support CD-ROM Extended Architecture (XA), a standard for interleaving different types of data, such as video, audio, and text. This provides you with the option of expanding your system to work with advanced multimedia applications.

Tape Drives

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

Backing up the entire contents of a hard disk to diskettes can be very time-consuming and, in the case of a network server with multiple hard disks, can require hundreds of diskettes. A faster and more efficient way to back up the data is to use a tape-backup drive.

Tape drives provide a cost-effective way to back up your files. You can copy several billion bytes of data from the hard disks to a single tape. Data is stored in the same sequence as it was entered.

The tape drives available for use with your server include:

- 4/10 GB 4-mm 3.5-Inch DAT Tape Option Kit
- 4/10 GB 4-mm 5.25-Inch DAT Tape Option Kit
- 3445 External 5/10 GB Tape Drive
- 24/48 GB Internal Tape Autoloader

The 3.5-Inch and 5.25-Inch 4/10 GB DAT Tape Drives

These two internal digital audio tape (DAT) drives use 4-mm tape cartridges and feature data compression, which allows you to store more data. The main difference between these drives is their external size.

The 3445 External 5/10 GB Tape Drive

This tape-backup drive uses 8-mm tape cartridges and also offers the data compression feature. You can attach it to the external SCSI-2 connector.

The 24/48 GB Internal Tape Autoloader

You can install this autoloading tape-backup drive in an internal, 5.25-inch bay. It has six cartridges that can each store up to 8 GB of data. In the native mode, this drive supports a data-transfer rate of 510 KBps, but this data-transfer rate doubles when the

data-compression feature is used. This drive also supports burst-data transfers of 3 MBps in the asynchronous mode and 7.5 MBps in the synchronous mode. A convenient feature of this tape-backup drive is its ability to perform up to 8 GB of backup every six days completely unattended.

Contact your IBM reseller or marketing representative for more information about the tape drives available for the PC Server 520.

IBM External SCSI Storage Enclosures

A SCSI storage enclosure is used to increase the data-storage capacity of any server that has a SCSI controller. The IBM Server Enterprise Expansion Enclosure is one of several SCSI storage enclosures that you can attach to the external SCSI-2 connector on your server.



This storage enclosure is shipped without 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 manual that comes with the enclosure includes information about setting it up, installing and removing drives, and attaching the enclosure to the server.

Chapter 6. The SCSI Subsystem

Your PC Server 520 comes with a *small computer system interface (SCSI)* controller. Your server uses the SCSI interface to communicate with devices such as hard disk drives, read-and-write optical drives, CD-ROM drives, tape drives, scanners, and printers.

In the PCI/EISA model, the SCSI controller is an adapter: either the SCSI-2 Fast/Wide PCI Adapter (referred to as the SCSI adapter) or the IBM PC ServeRAID Adapter (referred to as the ServeRAID adapter). (RAID is an acronym for *redundant array of independent disks*).

The PCI/Micro Channel models contain an integrated SCSI controller located on the system board. Servers that use the disk-array subsystem also have a ServeRAID adapter installed in a PCI expansion slot. When a ServeRAID adapter is installed, it and the integrated SCSI controller operate independently.

Advantages of the SCSI Subsystem

Before the SCSI standards were developed, many physical devices required controllers (usually adapters) that could support only one device or multiple devices of the same type. Because each device required support from the system microprocessor, every device added to the system increased the workload on the microprocessor. Also, because almost every device required its own adapter in a separate expansion slot, flexibility and growth were limited by the number of expansion slots in the system.

The SCSI standards define the interface through which the system communicates with different types of devices. Each SCSI device has a built-in *device controller* that interprets SCSI commands. The SCSI standards define the commands and messages used in the interface.

Several SCSI devices can operate concurrently and share the interface, achieving significantly higher data-transfer rates. Faster, more efficient, I/O management for network file servers can be realized using the SCSI subsystem. The SCSI subsystem produces significant improvements in system performance and flexibility over other I/O interfaces.

You can connect internal and external physical devices to the same SCSI subsystem. The integrated SCSI controller and the SCSI adapters each have internal and external connectors.

The integrated SCSI controller and the SCSI adapter support up to 15 SCSI devices. Up to seven 8-bit devices can be attached to the internal, narrow connector. If there are no devices attached to the internal 8-bit connector, up to fifteen 16-bit and 8-bit SCSI devices can be attached to the internal, wide connector or to the external, wide connector.

To make effective use of your SCSI connectors, you might consider using the *daisy-chain* connectors on the backplanes. Daisy chaining allows you to connect adjacent backplanes serially. In that way, you can control up to 12 hard disk drives connected to the backplanes with a single SCSI channel. Refer to your *User's Handbook* for more information.

Note: An 8-bit to 16-bit converter is required to attach an 8-bit device to a 16-bit (wide) cable.

The ServeRAID adapter supports up to 45 physical devices. The ServeRAID adapter has three wide, internal connectors and one wide external connector. If you use all three internal connectors, you cannot use the external connector.

The SCSI controller and adapters support:

- SCSI devices that conform to ANSI SCSI Standard X3.131-1986 or X3.131-1994 (SCSI-2)
- Devices 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 provided by the device manufacturer

Types of Interfaces

The two most common types of interfaces for attaching devices to a server are *device-level* interfaces and *bus-level* interfaces. The SCSI is a bus-level interface.

Device-Level Interfaces

Each device connected to a device-level interface has a controller that is mounted on a separate adapter, away from the physical device. Functions such as formatting, head selection, and error detection are directed by the controller and not by the device itself. The device is limited to performing the mechanical operations of reading and writing data, and it can do only what the controller allows it to do. There are limitations on the types of devices you can attach to the same controller.

Bus-Level Interfaces

Each device connected to a bus-level interface has a *device controller* that provides the details of how to accomplish the tasks performed by the device. A set of commands, the Common Command Set, specifies the tasks that devices are to perform. Each device controller determines how the tasks should be performed. The SCSI controller imposes few restrictions on the number and types of devices that you can attach to it.

The SCSI standard allows great flexibility in the design of devices. Hardware manufacturers can develop SCSI devices to perform functions without requiring that the SCSI controller support those functions. Functions unique to a particular SCSI device are supported by the device driver provided by the device

manufacturer. As new SCSI devices become available, there is no need to install new adapters. You can attach many different devices, such as hard disk drives, CD-ROM drives, scanners, and printers, to the same SCSI controller.

SCSI Physical Configuration

The following describes a possible configuration of the SCSI subsystem in a PC Server 520.

- In PCI/Micro Channel models, the integrated SCSI controller provides one external, wide connector (sometimes called the SCSI-2 port); and two internal connectors, one wide and one narrow. Only two connectors can be used at a time, regardless of whether the external SCSI connector is used.
- In PCI/EISA models, the SCSI-2 controller (on a PCI adapter) is connected to the PCI bus.
- SCSI devices are attached to a SCSI connector by a cable. This can be a SCSI cable inside the system that provides connectors for internal devices, or a SCSI cable (or series of cables) that connects to the external SCSI-2 connector.
- You can attach up to 15 devices to the integrated SCSI controller or SCSI adapter using the 16-bit internal and external connectors. You can attach up to 7 devices using the 8-bit internal connector. The ServeRAID adapter supports up to 45 physical devices.
- If any SCSI device connected to the controller has a data transfer rate faster than 5 MB/second, the maximum length of the SCSI bus (cable) must not exceed 3 meters (approximately 10 feet). Otherwise, the maximum bus length is 6 meters (approximately 20 feet). These lengths include internal and external cables. Following these standards will ensure proper server operation.
- The first external device is connected to an external SCSI cable, which connects to the external SCSI-2 connector. (On the PCI/EISA model, the external SCSI-2 connector is a 68-pin connector on the adapter.) Additional external devices are connected to each other with SCSI option-to-option cables. The last external device must have a terminator; if there is only one

external device, it must have a terminator. See “Termination” on page 52 for more information.

- For the integrated SCSI-2 controller, one internal connector is for an 8-bit cable, and the other internal connector is for a 16-bit cable. You can use any two connectors simultaneously.

For the SCSI-2 adapter, one internal connector is for an 8-bit cable, and the other internal connector is for a 16-bit cable. Again, you can use any two connectors simultaneously (that is: both internals; or one internal and the external—but not all three).

For the ServerAID adapter, the internal and external connectors are for 16-bit cables. And again, you cannot use the internal channel 1 and the external channel 1 at the same time.

For the integrated SCSI controller, the SCSI-2 adapter, and the ServerAID adapter, the 8-bit connectors support 7 devices, the 16-bit connectors support 15 devices.

Note: Each device attached to the interface must conform to ANSI SCSI Standard X3.131-1986 or X3.131-1994 (SCSI-2). These standards specify the mechanical, electrical, and functional requirements and the command sets that must be used by devices in the SCSI subsystem.

SCSI Device IDs

The SCSI ID is also known as the *SCSI address* or *physical unit number (PUN)*.

Each SCSI device must have a unique SCSI ID. The unique ID prevents two devices from sending data to the same SCSI controller at the same time. In all models, the SCSI IDs are set automatically in 18 hot swap bays. SCSI devices use a variety of mechanisms to set the SCSI ID, such as jumpers, thumbwheels, or switches. Devices attached to a wide connector can use IDs 0 through 15, except 7, which is reserved; devices attached to the narrow connector can use SCSI IDs 0 through 6 only.

The integrated SCSI-2 controller (or SCSI-2 adapter) is always preset to SCSI ID 7. The CD-ROM drive is preset to ID 6. The hot-swap

drives in banks C, D, and E are automatically assigned SCSI IDs 0 to 5 from left to right in each bank. Removable-media read and write devices, such as tape drives, should be assigned higher IDs than hard disk drives.

All three 16-bit connectors on the ServeRAID adapter support SCSI IDs 0 through 14. The ServeRAID adapter reserves SCSI IDs 7 and 15 for the controller. You can change the SCSI IDs that ServeRAID adapter reserves for the controller using the IBM PC ServeRAID Adapter Option Diskette. Devices attached to one channel can use SCSI IDs 0 through 14. Devices attached to another channel also can use SCSI IDs 0 through 14.

Note: If you connect backplanes using the daisy-chain connectors, you must set the SCSI ID address jumper on each of the backplanes to define how the system will assign SCSI IDs. Refer to the *User's Handbook* for instructions.

When more than one physical device is connected to the same SCSI-2 controller, the devices compete for support from the SCSI-2 controller. Because the SCSI-2 controller communicates with only one device at a time, the priority of a device is based on the ID of the device.

In summary, SCSI IDs perform two functions:

- They allow the SCSI-2 controller to distinguish one device from another.
- They allow the SCSI-2 controller to determine the priority of each device.

SCSI Logical Unit IDs

Each physical device has a physical unit number and a *logical unit number (LUN)*. If a physical device supports no logical devices, its LUN is 0. Most SCSI devices support only LUN 0. Certain types of physical SCSI devices (*bridge controllers*) can support more than one logical unit (device).

For example, a printer controller might support up to eight printers. The printer controller is a physical device, and each printer is a

logical device. If the printer controller is assigned to ID 6, the ID of the first printer appears in the SCSI device configuration as PUN 6, LUN 0. The IDs of the other printers are PUN 6, LUN 1; PUN 6, LUN 2; and so on.

Cabling

An interface that conforms to the SCSI standard consists of a set of eight data lines, one data-parity line, and various control lines. For cables that support SCSI wide devices, the interface consists of a set of 16 data lines, two data-parity lines, and various control lines. Consider the following factors before selecting the cables for your system:

- The integrated SCSI-2 controller, and the SCSI-2 adapter that comes with non-disk-array models of the PC Server 520, have one 8-bit 50-pin narrow, internal connector; one 16-bit 68-pin wide, internal cable connector; and one 16-bit wide, external connector.
- All models come with a 16-bit wide cable. A 16-bit to 8-bit SCSI Internal Converter connects the 8-bit CD-ROM drive to the wide cable.

Contact your IBM reseller or marketing representative for more information about SCSI cables.

The SCSI standard defines the signals that are carried on the interface. In all SCSI configurations, the interface must be terminated at both ends of the cable. When both an internal cable and an external cable are present, the entire configuration is seen as one long bus. In this case, there are only two termination points, at the far ends of the internal and external cables.

Termination

SCSI devices are cabled to each other and to the SCSI adapter or controller. The device at each end of the interface must be terminated.

If you have only internal SCSI devices connected to the SCSI-2 controller, the controller becomes the termination at one end of the cable.

For non-disk-array servers, if you use both internal connectors, or one internal connector *plus* the external connector, you must disable the termination by using the SCSI utility program included with the server. See the “Installing Options” chapter of the *User's Handbook* for more information.

For PCI/Micro Channel models, if you have attached SCSI devices to the SCSI controller located on the system board, and installed the devices in the non-hot-swappable bays, you *must* set termination.

For disk-array models, the ServeRAID adapter sets termination automatically.

The following rules apply to SCSI-device termination:

Note: For non-disk-array models, refer to the instructions supplied with the device for specific information about termination.

- If you have both internal SCSI devices and external SCSI devices, the devices at the ends of the internal and external cables must be terminated. The adapter must not be terminated. If you have only one external device, that device must be terminated.

Note: If you connect external devices, you can attach internal SCSI devices to only one of the internal SCSI connectors.

- If you have more than one internal SCSI device, the device farthest from the SCSI-2 controller must be terminated. Other internal drives connected to the same SCSI-2 controller must not be terminated.

Chapter 7. Disk-Array Subsystems

IBM provides an integrated disk-array subsystem as a standard feature on some PC Server 520 models. Each disk-array model of the PC Server 520 contains a PC ServeRAID Adapter.

All PC Server 520 models support up to 18 *hot-swappable* hard disk drives. A hot-swappable drive can be removed and replaced while the server remains powered on.

Internal Disk-Array Subsystem

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 improve security, performance, or reliability. They provide faster input and output rates than single large-capacity drives, because the files are stored on multiple individually addressable disks and can be found more quickly during I/O requests.

The amount of improvement in security, performance, and 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 520 supports RAID levels 0, 1, and 5.

RAID Technology

RAID technology groups two or more hard disk drives in a server into an *array* that is defined as a single *logical drive*. (Sections or *chunks* of hard disk drives can also comprise logical drives.) This logical drive appears to the operating system as a single *physical drive*. RAID technology enhances capacity, access times, and data transfer rates beyond the physical limitations of existing hard disk drives.

In addition, RAID technology can provide enhanced fault tolerance. If a physical drive fails in a RAID level 1 or level 5 disk array, the

system continues to run, with no operator intervention required. The defective (*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.

Note: RAID levels 1 and 5 offer data redundancy and the performance features mentioned above. RAID level 0 provides maximum hard disk drive capacity, but no data redundancy. See “Disk-Array Classifications” for more information about RAID levels.

When you install internal SCSI drives in your server, you can connect them to either the channel 1, channel 2, or channel 3 connector on the ServeRAID adapter (see the *User's Handbook* for more information). The *external* connector on the ServeRAID adapter also is designated “channel 1.” You cannot connect devices to both the internal channel 1 and external channel 1. This means that you can connect SCSI devices to all three internal connectors and not the external connector; or, you can connect SCSI devices to two of the internal connectors and the external connector but not to all four connectors simultaneously.

To make effective use of your SCSI connectors, you might consider using the *daisy-chain* connectors on the backplanes. Daisy chaining allows you to connect adjacent backplanes serially. In that way, you can control up to 12 hard disk drives connected to the backplanes with a single SCSI channel.

Disk-Array Classifications

The term “RAID” is commonly qualified by a number that classifies the different types of data-storage patterns used by disk arrays. There are six basic RAID classifications, RAID-0 through RAID-5; your server supports the classifications RAID-0, RAID-1, and RAID-5.

RAID Level 0

RAID level 0 stripes the data across all the drives in the array. This offers substantial speed enhancement, but provides for no data redundancy. The RAID adapter automatically assigns RAID level 0 to all logical drives in an array containing only one hard disk drive.

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*. If you have logical drives assigned RAID level 1 or 5 in the same array, they will not lose data.

When you replace the failed drive, the RAID adapter automatically rebuilds all the logical drives assigned RAID levels 5 and 1 onto the replacement hard disk drive and defines the level-0 logical drive. However, the data that was in the failed level-0 logical drive is lost.

Though the risk of data loss is present, you might want to assign RAID level 0 to one of the logical drives to take advantage of the speed offered with this level. You could use this logical drive to enter data that you back up each day and for which safety is not of primary importance; that is, data that you can re-create easily. You also might want to use a level-0 logical drive when the work you are doing requires maximum capacity. Level 0 provides the largest capacity of the three RAID levels offered, because no room is taken up for redundant data or data parity storage.

In the event of a drive failure in an array that contains a logical drive assigned RAID level 0, the RAID adapter will not activate a hot-spare drive.

RAID Level 1


RAID level 1 provides 100% data redundancy and requires two hard disk drives. With RAID level 1, the first stripe is the data stripe; the second stripe is the *mirror* (copy) of the first stripe, but written to another drive. If one of the hard disk drives fails, the RAID adapter switches read and write requests to the remaining functional drive in the array.

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, the RAID adapter switches automatically to the Enhanced RAID level 1. The Enhanced RAID level 1 requires a minimum of three drives and supports a maximum of 16 drives.

In an Enhanced RAID level 1 array, the RAID adapter stripes data and copies of the data across all the drives in the array. As with the standard RAID level 1, the data is mirrored, and the capacity of the logical drive is 50% of the physical capacity of the grouping of hard disk drives in the array.

The following illustration shows data arranged in an array with three hard disk drives. The logical drive is assigned RAID level 1. Notice that the mirror of the first data stripe (XXX YYY ZZZ) is shifted one drive. The other data stripes follow the same pattern.



If you have only two drives available, the second drive is a mirror copy of the first drive.

RAID Level 5

RAID level 5 stripes data and parity across all drives in the array. When an 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 backup the data stored in the array, and then redefine the level to level 1; or, you can use a logical drive that you have assigned level 1.

Note: Before you reconfigure or redefine an existing array, be sure to back up your data. When you redefine an existing array, all data in that array is lost.

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.

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 1 GB or 2 GB), 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 1 GB, 1 GB, 1 GB, and 2 GB grouped in one disk array, the total capacity of the array will be 1 GB times 4, or 4 GB (instead of the 5 GB physically available).

Note: The total capacity of the example array is 4 GB; a RAID Level 5 array consisting of the same drives has a usable capacity of 3GB.

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

* GB=1,000 million bytes.

Hard Disk Drive Mapping

The drive locations shown on your IBM ServeRaid Controller Disk Array 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 following table shows the server configuration as it is shipped, with channel 1 and channel 2 not connected.

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

Table 1. One Channel Mapping

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

Displayed Channel Number	Displayed Bay Number	Bank and Bay Physical Location	SCSI ID
3	1	C1	0
3	2	C2	1
3	3	C3	2
3	4	C4	3
3	5	C5	4
3	6	C6	5
3	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 2. Two Internal Channels Mapping

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

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

Table 3. One Internal and One External Channel Mapping

The following table shows an example of using all three channels on the ServeRAID adapter. The first channel is attached to an external DASD storage enclosure, the second channel is connected to the optional backplane behind bank D, and the third channel is connected to bank C. Refer to the documentation that comes with the storage enclosure for physical locations.

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

Table 4. All Channel Mapping

If you install another PC ServeRAID Adapter in your server, you can use the same combinations shown here, depending upon your space availability.

Additional Storage Capacity

When you add one or more hard disk drives to your server, you must configure a new disk array before you can use the added storage capacity. You can either reconfigure the existing disk array to include the capacity offered with the added drive or 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. You can then use the array as a single logical drive or you can subdivide it 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 ServeRAID adapter supports up to eight independent arrays and a total of eight logical drives. You can combine a maximum of 16 drives into an array. An array can span both channels. (For information about *physical drives* supported, see Chapter 5, “Data Storage Devices” on page 41.)

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.

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 non-arrayed 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 also 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 PC ServeRAID Adapter supports:

- A 32-bit PCI data bus
- Local 32-bit PCI bus
- Dual bus architecture
- Change RAID levels and add logical drives, dynamically
- Advanced recovery from RAID configuration conflicts
- 4 MB 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
- Three internal 16-bit SCSI-2 channels and one external 16-bit SCSI-2 channel, which can support eight arrays. (The external channel 1 and internal channel 1 are shared and cannot both be used at the same time.)

(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 ServeRAID adapter provides multiple data paths to and from arrayed drives, your server can respond to requests from several users simultaneously. With its overlapped input and 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 8 KB to 64 KB 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 64. 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 ServeRAID adapter has 4 MB of memory. This cache memory operates by default in write-through mode, but can be changed to write-back mode. The cache memory has parity to detect memory errors and uses retry algorithms to recover from sporadic errors.

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 have the capacity to recover the data from the drives that it replaces. For example, if each drive in an array contains 200 MB of data, a hot-spare drive for the array must have a minimum capacity of 200 MB. Even if each drive in the array has a 300 MB capacity. 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 program, the system automatically defines the replacement drive as a hot spare. The system defines the defunct drive as a defunct hot-spare drive.

Data Protection

The ServeRAID adapter provides the server with a data redundancy technique that distributes data across all the drives in the array. If a drive fails, 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, 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.

Configuration Data

If you install additional ServeRAID adapters, you can optionally delete the new ServeRAID adapters' configuration data. Your system requires only one set of configuration data for all ServeRAID adapters in your server. This allows you to conserve space in flash memory.

Chapter 8. Configuration

You play a key role in how your server allocates resources to hardware devices and software programs. This resource-allocation process is referred to as *configuration*. The steps required to configure your server depend on the PC Server 520 model you have, the number and variety of devices you install, and the software programs you install.

Your server has the flexibility and power to support several types of adapters and devices. Depending upon your model, this flexibility lets you choose from among thousands of adapters and devices that comply with the following standards:

- Peripheral component interconnect (PCI)
- Micro Channel architecture
- Enhanced industry standard architecture (EISA)
- Industry standard architecture (ISA)
- Small computer system interface (SCSI)

In general, the greater the number and variety of hardware devices and software programs that you install, the more you will have to interact with your server and your devices to correctly configure your system.

Configuration Programs

All PC Server 520 models come with menu-driven configuration programs. These programs enable you to view or change:

- System-board settings and assignments
- EISA and ISA device assignments
- Micro Channel device assignments
- Disk array configuration assignments
- SCSI device assignments
- PeerMaster adapter assignments

The types of configuration programs provided with your server depend on your model. All PC Server 520 models come with built-in setup and configuration utility programs, which you can access during system startup. These programs are stored in nonvolatile RAM on the system board.

The built-in setup and configuration programs enable you to view or change system-board settings, such as interrupt request levels (IRQs), memory-address assignments, I/O port assignments, and the startup sequence for non-disk-array drives. These programs also allow you to set the date and time, set passwords, and define system security.

Step-by-step instructions for using all of the setup, configuration, and utility programs are provided in your *User's Handbook*.

In most cases, the server will operate using the default settings. Before you install a new device or program, or make changes to your system-board settings, be sure to start the built-in utility programs; then record the current setting. In the unlikely event that a conflict occurs, you will be able to change the settings back to their original values.

Configuration Programs for PCI/EISA Models

In addition to the built-in setup and configuration utility programs, PCI/EISA models come with two of the following configuration programs and diskettes:

- EISA Configuration Diskette

This diskette comes with all PCI/EISA models. It contains configuration and utility programs for EISA and ISA adapters. You can use the programs on this diskette to configure EISA and ISA adapters. (Instructions for using this diskette are in your *User's Handbook*.)

- IBM PC ServeRAID Adapter Option Diskette

This diskette comes with disk-array models. It contains device drivers, monitoring programs, and menu-driven configuration programs that enable you to define, create, modify, or delete your disk-arrays. The programs on this diskette also enable you to view or change the SCSI device assignments for the drives attached to the IBM ServeRAID adapter.

Note: The SCSI termination and SCSI IDs for hot-swap drives are automatically set by the backplane behind bank C.

See your *User's Handbook* for additional information about defining or changing your disk-array configuration.

- SCSISelect utility program

This built-in SCSI-2 program comes with non-disk-array models. It enables you to view or change the assignments of your SCSI devices. You can use this utility program to change default settings, resolve configuration conflicts, or perform a low-level format on a SCSI hard disk drive. See “The SCSISelect Utility Program” on page 69 for more information. (Instructions for using this program are in your *User's Handbook*.)

- IBM PeerMaster Server Adapter Option Diskette

This diskette is provided with models that come with an IBM PeerMaster Server Adapter preinstalled. It contains device drivers, setup programs, NetFinity enablers, and diagnostic programs for the PeerMaster adapter. Instructions for using the programs on this diskette are in your *User's Handbook*.

Configuration Programs for PCI/Micro Channel Models

In addition to the built-in setup and configuration utility programs, PCI/Micro Channel models come with three or more of the following configuration programs and diskettes:

- PC Server 520 Reference Diskette

This diskette comes with all PCI/Micro Channel models. It contains configuration, utility, and test programs for Micro Channel adapters. You can use the programs on this diskette to automatically configure, view, or change the settings of Micro Channel adapters. This diskette also contains the Copy an Option Diskette program, which enables you to copy option-specific configuration programs from Option Diskettes to your backup copy of the Reference Diskette. (When you install additional Micro Channel adapters, you must use the Copy an Option Diskette program to update your backup copy of the Reference Diskette.) (Instructions for using the programs on the Reference Diskette are in your *User's Handbook*.)

- IBM PC ServeRAID Adapter Option Diskette

This diskette comes with disk-array models. It contains device drivers monitoring programs, and menu-driven configuration programs that enable you to define, create, modify, or delete your disk-arrays. The programs on this diskette also enable you to view or change the SCSI device assignments for the drives attached to the ServeRAID adapter.

Note: The SCSI termination and SCSI IDs for hot-swap drives are automatically set by the backplane behind bank C.

See your *User's Handbook* for additional information about defining or changing your disk-array configuration.

- SCISSelect utility program

This built-in SCSI-2 utility program enables you to view or change the assignments of SCSI devices that are connected to the integrated SCSI-2 controller. You can use this utility program to change default settings, resolve configuration conflicts, or perform a low-level format on a SCSI hard disk drive that is connected to the integrated SCSI-2 controller. See “The SCISSelect Utility Program” for more information. (Instructions for using this program are in your *User's Handbook*.)

- IBM PeerMaster Server Adapter Option Diskette

This diskette is provided with models that come with an IBM PeerMaster Server Adapter preinstalled. It contains device drivers, setup programs, NetFinity enablers, and diagnostic programs for the PeerMaster adapter. Instructions for using the programs on this diskette are in your *User's Handbook*.

The SCISSelect Utility Program

The integrated SCSI-2 controller (on PCI/Micro Channel models) and the SCSI-2 Adapter (in non-disk-array PCI/EISA models) each have a built-in, menu-driven configuration program that enables you to view or change SCSI settings. You can use the SCISSelect program to change default settings, locate and correct configuration conflicts, and format a new SCSI hard disk.

The settings that you can view or change include:

- Host Adapter SCSI ID

You can change the SCSI ID of the SCSI-2 controller or adapter from its default value of 7.

- SCSI Parity Checking

You can enable or disable parity checking on the SCSI-2 adapter or controller. The default setting is *Enabled*.

- Host Adapter SCSI Termination

You can change the SCSI-2 adapter or controller termination. The default setting is *Low On / High On*.

- SCSI Device Configuration

This enables you to configure SCSI device parameters. Before you can make updates, you must know the ID of the device whose parameters you want to configure.

- Advanced Configuration Options

This enables you to view or change the settings for advanced configuration options. These include:

- Startable CD-ROM drive: When set to *On*, you can include the CD-ROM drive as a startup device. For this option to take effect, you must insert a startable CD in the CD-ROM drive.
- Init13 Extensions: When set to *On*, these BIOS extensions are enabled.

Configuration Process Overview

When you install a new device or change the position of an existing device, you might need to reconfigure your server. The steps required to configure your server depend on the option you install.

Configuring SCSI Devices

When you install hot-swap drives in banks C, D, or E, the SCSI ID for the drive is assigned by the backplane behind the bank. The backplane also manages the termination of the hot-swap drives. However, if you attach an external SCSI device or install a SCSI device in bay A or B, you will need to manually set the SCSI ID and termination for the device. Check the documentation that comes with the device for instructions to set the SCSI ID and termination for the device. (See also “Termination” on page 52 for more information.)

Configuring a PCI/EISA Model

The PCI/EISA models use jumpers, switches, and configuration programs to assign interrupt levels, input/output addresses, memory addresses, and other server resources.

Some options, such as memory and PCI adapters, are recognized and configured automatically by the system. However, other options—such as EISA adapters and SCSI devices—must be manually configured.

Configuring EISA Adapters

Before you install an EISA or ISA adapter in your server, be sure to check the documentation that comes with the adapter. Some adapters have jumpers or switches that you must set to assign system resources, such as interrupt levels and memory addresses. Before you assign a system resource to the adapter, you must verify that the resource is not already assigned to the system board or another device. To do this:

1. Check the adapter's documentation to see which system resources, such as interrupt level or memory-address space, it needs to function properly.
2. Start the built-in setup and configuration programs.
3. Record the settings of the existing devices.
4. Compare the settings of the existing devices to the settings required by the new adapter; then choose settings for the new adapter that will not conflict with an existing device.

Some adapter documentation might list only one choice for a system resource. For example, the adapter might operate only at interrupt level 3. If this occurs, you might need to change the assignment of an existing device to avoid a configuration conflict.

When you assign interrupt levels, remember that when you assign a very high interrupt level to a high performance device, performance can be adversely affected. For best performance results, assign a low interrupt level to a high performance device. For example, you might assign interrupt level 10 to an Ethernet PCI LAN adapter.

5. Set the jumpers or switches on the new adapter.
6. Install the new adapter in an available expansion slot.
7. Use the EISA Configuration Diskette to assign or verify the adapter's configuration.

Configuring a PCI/Micro Channel Model

The PCI/Micro Channel models use programmable configuration features, configuration programs, and configuration files to assign interrupt levels, input/output addresses, memory addresses, arbitration levels, and other server resources.

Some options, such as memory, are recognized and configured automatically by the system. However, other options, such as SCSI devices, must be manually configured.

Programmable Configuration

The programmable configuration hardware consists of:

- Nonvolatile RAM on the system board
- Hardware used to store the device ID
- Registers on the system board and adapters

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 on the server. The registers are volatile storage areas, and retain the configuration information only while the server is turned on. When you turn off the server, 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 some setup information, such as the power-on password.

Each adapter has a unique device ID. The POST routine and the configuration programs use the device ID to identify the type and location of each adapter in the server. The configuration programs also use the device ID during the configuration process to locate the appropriate configuration files for the system board and Micro Channel adapters.

The programmable configuration 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

Configuration Files: The Set/View Micro Channel Configuration program (hereafter called the Set Configuration program) on the Reference Diskette uses configuration files to assign interrupt levels, I/O addresses, memory addresses, arbitration levels, and other server resources.

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

The Set Configuration program uses two types of files to configure the server: adapter description files and adapter description programs. Some devices require only an adapter description file, but more complex devices sometimes require an adapter description program. Configuration files for optional devices usually are provided on an Option Diskette supplied with the device.

The following table shows the types of files 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 ID for the device.

File Name	Description
@xxxx.ADF	Adapter description file
Cxxxx.ADF	Adapter description 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

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 the device ID and several choices for:

- Interrupt levels
- I/O, RAM, and ROM addresses
- Arbitration levels
- Other server resources that the device can use

Usually, 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 program cannot find the adapter description file, the server prompts you to insert the Option Diskette for the device. If you insert the diskette, and the Set Configuration program still cannot find the adapter description file, the program disables the device.

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.

Adapter Description Programs: Some devices 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.

A device that needs an adapter description program has an entry in its adapter description file that tells the Set Configuration program to run the adapter description program. During configuration, the Set Configuration program reads the adapter description file for each device. When the Set Configuration program finds an entry for an adapter description program, it locates and runs it.

If the Set Configuration program cannot find the adapter description program, the server prompts you to insert the Option Diskette for the device. If you insert the diskette, and the Set Configuration program still cannot find the appropriate adapter description program, the program disables the device.

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, be sure to read the instructions that come with the device. You might need to copy configuration files from the Option Diskette to the backup copy of the Reference Diskette. Use the instructions provided with the device to install the configuration files. If no instructions are provided, use the Copy an Option Diskette program to install the configuration files.

Before copying the configuration files, the Copy an Option Diskette program checks the Option Diskette for valid configuration files. If the program does not find valid configuration files, a message appears. If the program finds valid configuration files, it checks 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 option-configuration files are copied to the Reference Diskette, they become an integrated part of the system programs.

Chapter 9. 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 ports include:

- Two serial ports
- One parallel port
- One video port
- One keyboard port
- One auxiliary-device (mouse) port

“Input/Output Ports” on page 12 shows the locations of these ports.

Serial Ports

Serial ports are used to communicate with printers, plotters, external modems, auxiliary terminals, and other computers. Your server provides two serial ports. You can add more serial ports by installing a serial adapter in one of the expansion slots.

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

The serial ports transmit and receive data and commands at a rate of 300 bits per second (bps) up to 345 600 bps. (To use a serial port at 345 600 bps, you need a special shielded cable. For information about this cable, contact your IBM reseller or marketing representative.) The bits-per-second measurement is commonly referred to as the *baud rate*.

Serial Port Connectors

Each serial port uses a 9-pin, male D-shell connector. The pin-number assignments conform to the RS-232C industry standard. The following table shows the pin-number assignments for the serial port connectors.

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 Pin-Number Assignments

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 adapters to support up to 16 serial ports in your server, named 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. In some configurations, you may need to change communication port address assignments to resolve conflicts.

Note: If you experience a problem with a non-IBM device when attaching it to the serial port, you might need to set the serial port direct memory access (DMA) arbitration level to "Disabled." See the "Configuring Your Server" section in the *User's Handbook* for more information about address assignments.

Parallel Port

Parallel ports are used to communicate with printers and other devices, such as some CD-ROM and tape drives. Your server provides one 25-pin parallel port.

The parallel port uses a 25-pin, female D-shell connector on the back of the server. The parallel port conforms to the Institute of Electrical and Electronics Engineers (IEEE) Standard 1284. The following table shows the pin-number assignments for the parallel port connector.

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 Pin-Number Assignments

Note: If you experience a problem with a non-IBM device when attaching it to the parallel port, you might need to set the parallel port arbitration level to “Disabled.” See the “Configuring Your Server” section in the *User's Handbook* for more information.

Video Port

The video port allows you to attach a video display to your server. Your server provides a 15-pin video port connector. In the PCI/Micro Channel model, the video port connector is located on the system board. In the PCI/EISA model, the video port connector is on an SVGA adapter.

The following table shows the pin-number assignments for the video port connector.

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 Port Connector Pin-Number Assignments

Keyboard and Auxiliary-Device Ports

Your server provides one keyboard port and one auxiliary-device port. An auxiliary device can be a mouse or other pointing device.

The following table shows the pin-number assignments for the keyboard and auxiliary-device port connectors.

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 Connector Pin-Number Assignments

SCSI Ports

The integrated SCSI-2 controller, and the SCSI-2 adapter that comes with non-disk-array models of the PC Server 520, have one 8-bit 50-pin narrow, internal connector; one 16-bit 68-pin wide, internal cable connector; and one 16-bit wide, external connector.

On non-disk-array models, the wide connectors can support up to fifteen 16-bit and 8-bit SCSI devices. The narrow connector can support up to seven 8-bit SCSI devices.

Disk-array models have an external 2×34 -pin port and three internal 2×34 connectors on the ServeRAID adapter.

The system can support any three connectors simultaneously; you can attach up to 15 drives to each connector.

The following table shows the pin-number assignments for 2 × 34-pin SCSI port connectors.

Pin	Signal	Pin	Signal
1	Ground	35	Data 12
2	Ground	36	Data 13
3	Ground	37	Data 14
4	Ground	38	Data 15
5	Ground	39	Data P1
6	Ground	40	Data 0
7	Ground	41	Data 1
8	Ground	42	Data 2
9	Ground	43	Data 3
10	Ground	44	Data 4
11	Ground	45	Data 5
12	Ground	46	Data 6
13	Ground	47	Data 7
14	Ground	48	Data P0
15	Ground	49	Ground
16	Ground	50	Ground
17	+Term power	51	+Term power
18	+Term power	52	+Term power
19	Reserved	53	Reserved
20	Ground	54	Ground
21	Ground	55	-Attention
22	Ground	56	Ground
23	Ground	57	-Busy
24	Ground	58	-Acknowledge
25	Ground	59	-Reset
26	Ground	60	-Message
27	Ground	61	-Select
28	Ground	62	-Control/Data
29	Ground	63	-Request
30	Ground	64	-Input/Output
31	Ground	65	Data 8
32	Ground	66	Data 9
33	Ground	67	Data 10
34	Ground	68	Data 11

Table 10. 2 × 34-Pin SCSI Port Connector Pin-Number Assignments

The following table shows the pin-number assignments for the 2 × 25-pin internal SCSI-2 connectors.

Pin	Signal	Pin	Signal
1	Ground	26	Terminator power
2	Data 0	27	Reserved
3	Ground	28	Reserved
4	Data 1	29	Ground
5	Ground	30	Ground
6	Data 2	31	Ground
7	Ground	32	-Attention
8	Data 3	33	Ground
9	Ground	34	Ground
10	Data 4	35	Ground
11	Ground	36	-Busy
12	Data 5	37	Ground
13	Ground	38	-Acknowledge
14	Data 6	39	Ground
15	Ground	40	-Reset
16	Data 7	41	Ground
17	Ground	42	-Message
18	Data P	43	Ground
19	Ground	44	-Select
20	Ground	45	Ground
21	Ground	46	-Control/data
22	Ground	47	Ground
23	Reserved	48	-Request
24	Reserved	49	Ground
25	Not connected	50	-Input/output

Table 11. 2 × 25-Pin (Narrow) SCSI Port Connector Pin-Number Assignments

Chapter 10. Video Subsystem

Your PC Server 520 provides super video graphics array (SVGA) capability. The PCI/Micro Channel model contains a video controller on the system board. The PCI/EISA model contains a video adapter. Both the controller and the adapter use an external 15-pin analog video connector. This connector enables you to attach a video monitor to your server. For information about this connector, see “Video Port” on page 79.

The PC Server 520 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 monitors that meet ISO standards. For more information about noninterlaced monitors, see “Vertical-Refresh Rate” on page 87.

The video subsystem includes a video display buffer that contains 1 MB of *video dynamic random access memory (VDRAM)*. (MB equals approximately 1 000 000 bytes.) The amount of VDRAM in your server and the capabilities of the video driver determine the resolution and the number of colors that are supported. Your server supports the following screen resolutions and numbers of colors:

- 640 x 480 pels with up to 16.8 million colors (noninterlaced)
- 800 x 600 pels with up to 65 536 colors (noninterlaced)
- 1024 x 768 pels with up to 256 colors (noninterlaced)
- 1280 x 1024 pels with up to 16 colors (interlaced)

Note: Monitors supporting 1280 x 1024 pels and up to 16 colors require operating system support. See your operating system documentation for more information.

Additional features of the video subsystem include:

- Integrated VGA emulation
- Support for the 132-column text mode
- Reduced-flicker operation for noninterlaced monitors that meet ISO standards

You can install a video adapter in one of the expansion slots.

Monitors come in many sizes and types. The video controller supports a wide range of monochrome and color monitors having various capabilities. The most common distinguishing characteristics are *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 monitor screen is coated with light-emitting materials called *phosphors*. Color monitors use three phosphors: red, green, and blue. The phosphor coating is made up of dots arranged in a red-green-blue (RGB) 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 monitor with the highest resolution and finest dot pitch produces the sharpest image. A monitor with a fine dot pitch, such as 0.28, produces sharper, better defined characters than a similar monitor 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 monitor 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) monitor may work as well as a color monitor.

Vertical-Refresh Rate

The phosphors on the inside of the monitor screen hold their brightness and color for a very short time. The image must 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. Monitors that have a fast, noninterlaced refresh rate, such as the IBM 952x monitors, provide a very stable image with little or no flicker detectable. Monitors with slower refresh rates tend to have more noticeable flicker. This is particularly noticeable with white and other light-colored backgrounds. You can often minimize flicker by changing the colors you work with. Many programs allow you to select colors.

Sensitivity to flicker varies from person to person; young people are 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 Monitor

A monitor having a fast noninterlaced refresh rate fully uses the video capabilities of your server. The IBM 9524 or 9525 are examples of monitors that use your servers video capabilities fully.

However, as mentioned earlier, some applications demand less video support than others. IBM provides many monitors that you can use with your system. The following table provides a partial list of monitors and their characteristics.

Type	Dot Pitch	Resolution	Maximum Vertical Refresh Rate
G40	0.28	1024x768	43 Hz interlaced
		800x600	60 Hz noninterlaced
		640x480	75 Hz noninterlaced
G50	0.28	1024x768	72 Hz noninterlaced
		800x600	85 Hz interlaced
		640x480	85 Hz noninterlaced
P50	0.26	1024x768	85 Hz noninterlaced
		800x600	85 Hz noninterlaced
		640x480	85 Hz noninterlaced

Table 12. Monitor Types and Characteristics

For more details, refer to the information supplied with your monitor. See your IBM reseller or marketing representative for the current list of IBM monitors.

Screen-Saver Programs

Leaving the monitor turned on for prolonged periods without changing what appears on the screen can adversely affect the monitor. 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 or mouse has not been used for a defined period. To resume working with the monitor, you press a predefined key to restore the previous image to the screen. Several companies produce screen-saver programs. Contact your IBM reseller or marketing representative for details.

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 monitor, set the controls near their center position. Setting these controls to their maximum position can decrease the useful life of a monitor.

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

Avoiding Fatigue

The American Academy of Ophthalmology considers monitors safe for normal use and has declared they present no hazard to vision and do not damage the eye.

However, performing any visually demanding task for a long time, such as knitting, reading, or viewing a monitor screen, can cause eye fatigue. To avoid eye fatigue, you should reduce annoying glare and ensure that your vision is properly corrected for the task. If eye discomfort persists, consult your eye-care professional or physician.

Similarly, sore muscles and 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.

Modifying work practices can reduce muscle, joint, and eye fatigue. Taking breaks, rotating tasks, and adjusting the monitor, the lighting, and the chair height for your comfort usually result in less fatigue.

Additional information is available in your *User's Handbook*.

Chapter 11. 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 can implement the security features of your server to protect it and the data stored on it.

Before implementing any security features, evaluate your security requirements. Where will the system be located? Should it be secured to permanent furniture or fixtures? Should use of the system be limited? After you have identified your security needs, you can activate or implement the appropriate security features.

Keylocks and a system of passwords help deter problems caused by 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 security covers. (To find out whether your operating system supports a particular security feature, see the documentation that comes 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. For more protection, you can fasten your server to a stationary object, such as a table, a desk, or a wall. Another approach is to locate critical resources, including your server, in a limited-access area behind locked doors.

The hardware security features of your server include:

- The door lock and keys
- The optional U-bolt facility, which allows you to secure the system to a stationary object using a cable or chain
- The security-cover option

Door Lock and Keys


A door lock mechanically deters access to the front of the server. It controls access to the small door that covers the removable media drives as well as the front cover that secures the hard disk drives.

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.

Note: Anyone who has the key serial number and manufacturer's address can order duplicate keys, so store the tag in a safe place.

U-Bolt Facility


If your server is in an accessible area, you might need a method of securing it to a desk, table, or other stationary object. Two holes in the rear panel of the server can accommodate a U-bolt. After a 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 your *User's Handbook*.

Security-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. The rear cover of your server accommodates a security cover that prevents access to the connectors.



The security-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 security cover does not protect against unauthorized access through devices attached outside of the system.

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

LogicLock (PCI/Micro Channel Models)

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 systems are operating-system dependent, which means that (if you want to use them) you must use an operating system that supports them. Other security features are optional features, such as cable covers, that can be purchased. (To find out whether your operating system supports a particular security feature, see the documentation that came with your operating system.)

Securing Data

It is very difficult to protect 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 operating systems. Password protection is another common feature of network operating systems that helps protect data from theft. Passwords, which are easy to use and very effective, 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. Such information can help identify users who attempt 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 files so that they are unintelligible unless a special key is provided. You can use a program or an encryption device to encrypt data. Software programs are usually less expensive than hardware devices, but programs are slower.

Power-On Password (PCI/Micro Channel Models)

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. A power-on password can contain any combination of up to seven characters (A–Z, a–z, and 0–9). Keep a record of your password in a secure place.

If you forget the power-on password, you must remove the server cover and remove the battery for at least 15 minutes, unless you have set and know the administrator password. The administrator password can be entered at the power-on password prompt to unlock the keyboard and mouse. You can then use the Configuration/Setup utility program to set a new power-on password.

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

Administrator Password (PCI/Micro Channel Models)

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

The administrator password allows you to control access to the Configuration/Setup utility programs. If an administrator password is set, you must use it to access the system setup programs 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. The location of the administrator password jumper is shown on a label inside your server. 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 *User's Handbook* describes the procedure for changing the administrator password jumper and setting the administrator password.

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 the password prompt appears again. After three incorrect attempts, The system is locked message is displayed and you must turn off the server and start again.

Supervisor Password (PCI/EISA Models)

The supervisor password allows you to control who has access to the Setup program. If a supervisor password is set, you must enter it to use the Setup program. Only someone who knows the current supervisor password can change the supervisor password.

Before you can set a supervisor password, you must first set your selectable drive-startup sequence, as described in your *User's Handbook*.

Warning: If a supervisor password is set and then forgotten, it cannot be overridden or removed. If you forget your supervisor password, you must place a service call to regain access to your server.

The password becomes effective immediately after you enter it. After you set the supervisor password, you are prompted to enter it each time you attempt to start the Setup program. See your *User's Handbook* for details on the use of the supervisor password.

User Password (PCI/EISA Models)

You do not need to set a user password, but it helps protect the information you store in your server. The user password controls access to your operating system and the information stored on your hard disk drives.

The password becomes effective immediately after you enter it. You are prompted to enter the user password each time you start your operating system, after POST occurs. See your *User's Handbook* for details on the use of the user password.

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 (Micro Channel model only) 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 help 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 recover the data 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 large volumes 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.

Virus-detection programs 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.

Your server comes with a virus-detection program that examines files at system startup. On PCI/Micro Channel models, you can use the Configuration/Setup utility program to enable or disable the virus-detection programs. On PCI/EISA models, you can use the Setup utility program to enable or disable write protecting of the hard disk boot sector (see your *User's Handbook* for details). Several other programs are available that can detect the presence of many known viruses.

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. Some virus-detection programs can remove viruses and repair the damage caused by relatively benign viruses.

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 *User's Handbook* for information about testing 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. This format operation 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.

Redundant Power Supply

The IBM PC Server 780W Redundant Power Option gives you two power supplies in one self-contained option. If one power supply should fail, the other one immediately takes over so there should be no interruption in your server's performance. This gives you time to make plans and preparations for replacing and repairing the power supply without losing any valuable work time.

This power option also eliminates the need for you to have to purchase and install a second power supply, which you normally need to do to expand your server's storage capacity.

Backup Power Supply

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

Uninterruptible Power Supplies

An uninterruptible power supply provides power 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 and its internal battery then provides 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

A standby power supply waits until it detects a drop in voltage; it then 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. Product Warranties and Notices

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 520

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

Warranties

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 that IBM product, program, or service may be used. Subject to IBM's valid intellectual property or other legally protectable rights, any functionally equivalent product, program, or service may be used instead of the IBM product, program, or service. The evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, are the responsibility of the user.

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:

IBM Director of Licensing
IBM Corporation
500 Columbus Avenue
Thornwood, NY 10594
U.S.A.

Trademarks

The following terms are trademarks of the IBM Corporation in the United States or other countries or both:

IBM	Operating System/2	Personal System/2
LogicLock	OS/2	PS/2
Micro Channel	PeerMaster	ServerGuide
NetFinity	XT	HelpCenter

The following terms are trademarks of other companies:

Intel	Intel Corporation
Pentium	Intel Corporation
SCSIselect	Adaptec, Inc.
Adaptec	Adaptec, Inc.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Other company, product, and service names, which may be denoted by a double asterisk (**), may be trademarks or service marks of others.

Electronic Emission Notices

PCI/EISA models, and PCI/Micro Channel models are *Class A*:

Industry Canada Class A Emission Compliance Statement

This Class A digital apparatus meets the requirements of the Canadian Interference-Causing Equipment Regulations.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

For Class A models...

The following FCC statement applies to *Class A* models.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his own expense.

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. 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.

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 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

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.

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

back up. To copy information, usually to diskette or tape, for safekeeping.

backup. Pertaining to a system, device, file, or facility that can be used in the event of a malfunction or loss of data.

BIOS (Basic Input/Output System). Code that controls basic hardware operations such as

interactions with diskette drives, hard disk drives, and the keyboard.

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 bridge. A functional unit that interconnects two busses together, such as an EISA bus and a PCI bus.

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

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)

clock cycle. In microprocessors, a unit of time expressed in nanoseconds (ns) that measures the period between clock signals.

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.

connector. An electrical part used to join two or more other electrical parts. (Contrast with *port*.)

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.

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.

device driver. A file that contains the code needed to use an attached device.

diagnostic. Pertaining to the detection and isolation of errors in programs and faults in equipment.

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)

diskette drive. The mechanism used to seek, read, and write data on diskettes. It can be installed in, or attached to, a computer.

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*.

display. A component capable of displaying information on a viewing surface; for example, a cathode ray tube or a gas panel.

DMA. Direct memory access.

dot pitch. In color video displays, the distance on a scan line (expressed in millimeters) from the center of one phosphor dot to the center of the next phosphor dot of the same color.

E

ECC. Error correcting code.

EEPROM. Electrically erasable programmable read-only memory.

EISA. Extended industry standard architecture.

EOS. Error correcting code on SIMM.

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)

Error correcting code on SIMM (EOS). A method that provides increased system reliability, not only by detecting single-bit memory errors (as in parity memory) but also by correcting them in real time.

execution unit. A logical unit of a microprocessor that performs the arithmetic and logical operations specified by the microprocessor instruction set.

F

FDM. Frequency-division multiplexing.

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.

floating-point unit (FPU). In microprocessors, an execution unit that executes floating-point arithmetic instructions.

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. Approximately 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.

hot swap. To replace a hard disk drive while the system is turned on.

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)

IEEE. Institute of Electrical and Electronics Engineers.

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.

integer unit. In microprocessors, an execution unit that executes logical and integer arithmetic instructions.

I/O. Input/output.

ISA. Industry standard architecture.

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.

J

jumper. A device, such as a short wire, used to open or close a circuit.

K

KB. Kilobyte.

kilobyte. Approximately 1000 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. Approximately 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*.

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

parallel port. An access point through which a computer transmits or receives data that consists of several bits sent simultaneously on separate wires. Contrast with *serial port*.

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)

PCI. Peripheral component interconnect.

pel. (1) picture element. (2) In computer graphics, the smallest element of a display surface that can be independently assigned color and intensity.

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)

pin. One of the conducting contacts of an electrical connector.

pipeline. In microprocessors, a technique used to shorten the critical path of instruction execution by performing multiple execution stages simultaneously.

port. An access point for data entry or exit. (Contrast with *connector*.)

POST. Power-on self-test.

power-on self-test (POST). A series of diagnostic tests that are run automatically by a device when the power is turned on.

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)

PUN. Physical unit number.

R

RAID. Redundant array of independent disks.

RAM. Random access memory.

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.

RAS. Reliability, availability, and serviceability.

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 *EPROM*, *EPRM*, 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.

resolution. In video displays, a measure of the sharpness of an image, expressed as the number of lines and columns on the display screen or the number of pels per unit of area.

ROM. Read-only memory.

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.

serial port. An access point through which a computer transmits or receives data, one bit at a time. Contrast with *parallel port*.

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.

SIMM(s). Single inline memory module(s).

slot. (1) A position in a device used for removable storage media. (2) One of several receptacles in the rear panel of the system unit into which a user can install an adapter.

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.

superscalar. In a microprocessor, a technique in which multiple independent instructions are issued per clock cycle.

SVGA. Super video graphics array.

swap space. In virtual memory systems, external storage allocated to program and memory images.

symmetric multiprocessing. In personal-computer systems, a multiprocessing design that enables two or more microprocessors to run concurrently and work independently, with each microprocessor capable of performing any application task.

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. Approximately 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

VDRAM. Video dynamic random access memory.

vertical-refresh rate. The rate, measured in hertz, at which a display image is redrawn from top to bottom.

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

wait state. A period during which a microprocessor suspends processing while waiting for a response to a request for an unavailable resource.

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.

A

- about this book vii
- acoustical noise emission 18
- adapter description file (ADF) 38
- adapters
 - See bus-master adapter
 - See network adapter
 - See SCSI adapter
 - See ServeRAID adapter
 - See video adapter
- adding storage capacity 60
- address bus 35
- address interface 27
- ADF (adapter description file) 38
- administrator password 95
- advanced configuration options 70
- altitude of machine 18
- American National Standards Institute (ANSI)
 - SCSI standards 45
- application programs
 - See also multimedia applications
 - effect on performance 20, 23
 - mathematical calculations in 28
 - supported by server 27
- arbitration
 - between buses 35
- architecture
 - dual bus 33
 - EISA 34
 - local bus (PCI) 33
 - SCSI 45

B

- battery
 - information saved by 29
- bays
 - See expansion bays

- bolt-down capability 92
- bootable diskette 98
- bridge controllers 50
- brightness and contrast controls 89
- bulletin-board services 101
- burst mode 33, 35
- bus
 - See also address bus
 - See also data bus
 - See also expansion bus
 - See also local bus
 - See also Micro Channel bus
 - See also PCI bus
 - definition of 33
 - width 36
- bus architecture
 - EISA 34
 - expansion (Micro Channel)
 - local (PCI) 33
 - overview of 33
- bus master 36
- bus-level interface 47
- bus-master adapter
 - purpose of 22

C

- cable-down feature 92
- cables
 - attaching SCSI devices 51
 - daisy-chain hook-up 46, 54
 - option-to-option 48
 - SCSI, maximum length 48
 - terminating 52
- cache memory
 - description of 31
 - level-1 31
 - level-2 32
- CD (compact disc)
 - capacity of 42

- CD-ROM drive
 - description of 42
- central arbitration control point 37
- chapter descriptions vii
- chip
 - See microprocessor
- clock cycle 25
- clock speed 25
- Common Command Set 47
- communication ports (COM1 through COM16) 78
- compact disc (CD)
 - capacity of 42
- computer viruses 100
- configuration
 - errors 66
 - interrupt levels 72
 - of SCSI devices 48
 - peripheral component interconnect (PCI) 66
 - power-on self-test (POST) 66
 - utility programs 66
- configuring your server 66
- conflicts 66
- connectors
 - 2×34-pin SCSI 81
 - external device 81
 - for SCSI devices 46
 - internal SCSI 81
 - keyboard and auxiliary device 80
 - number of, for each bus 35
 - parallel port 79
 - parallel ports 79
 - security 92
 - serial ports 77, 78
 - video 79
- contrast and brightness controls 89
- controls, user-operated 6
 - brightness and contrast 89

cords, power 110

D

- daisy-chain hook-up 46, 54
- DASD (direct access storage device)
 - description of 41
- DAT (digital audio tape) 43
- data
 - backing up 43
- data bus 35
- data interface 26
- data transfer
 - burst mode 33, 35
 - effect of bus width 36
 - error checking 38
- device controller 45, 47
- device-level interface 47
- devices
 - assigning IDs to 49
 - connecting to SCSI subsystem 46
- diagnostic aids 17
- digital audio tape (DAT) 43
- direct access storage device (DASD)
 - description of 41
- direct memory access (DMA) 39
- disk arrays
 - effect on performance 22
- diskette drives
 - capacity of 41
- diskettes
 - formatting 41
- DMA (direct memory access) 39
- door lock 91
- dot pitch 86
- drives
 - See CD-ROM drive
 - See diskette drives
 - See hard disk drives

drives (*continued*)
 See optical drives
 See SCSI drives
duplicate data 55
duplicating data 100

E

EEPROM (electrically erasable programmable ROM) 29
effect of SCSI subsystem on
 cable length 48
 EISA bus 34
 ISA bus 34
 role of DMA 39
 streaming-data procedure 40
electrical requirements 18
electrically erasable programmable ROM (EEPROM) 29
electronic emission notices 109
encryption, data tool 94
erased files 102
even parity 38
expansion bays
 illustration of 8
expansion bus
 description of 34, 35
expansion slots
 number of connectors 35
 shared 11

F

fatigue, avoiding 90
fault-tolerance 64
FDISK 61
features of PC Server 520 2
flash memory 29
flexibility of system 47
flicker, reduced monitor 85

floating-point unit (FPU) 28
forgotten administrator
 password 96
formatting diskettes 102
FPU (floating-point unit) 28

G

glossary 113

H

hard disk drives
 operating system support for 41
heat output of machine 18
hot swap 53
hot-drive standby 62
hot-spare drive 64
hot-swap 53
hot-swap bays
 description of 42
humidity of machine
 environment 18

I

identifiers for SCSI devices 49
input/output management 46
input/output overlapped
 operation 63
input/output ports
 illustration of 12
 keyboard and auxiliary-device
 ports 80
 mouse port 80
 parallel port 79
 SCSI 81
 serial port 78
 video connector 79
instructions
 execution process for 25

integer units
 role in math computations 28
interfaces
 See address interface
 See bus-level interface
 See data interface
 See device-level interface
interlaced image 87
interleave depth/stripe unit 63
interrupt levels 72
introduction 54

K

keys, door lock 91

L

LAN management 19
languages, changing display for 16
level 1 56
level-0 data loss 55
lights, status 6
local bus
 description of 33
lock, door 91
locked state 95
locking the keyboard 94
logical unit number (LUN)
 for SCSI devices 50
low-level format 102
LUN (logical unit number)
 for SCSI devices 50

M

mathematical computations 28
memory
 See also cache memory
 See also electrically erasable
 programmable ROM (EEPROM)

memory (*continued*)

 See also flash memory
 See also nonvolatile memory
 See also physical memory
 See also random access memory
 (RAM)
 See also static random access
 memory (SRAM)
 See also system memory
 See also virtual memory
 See also volatile memory
adding to server 30
effect on performance 22
memory controller 30
memory subsystem 28
Micro Channel bus
 description of 34
microprocessor
 definition of 25
 design of 26
 installing dual 25
 performance of 22
 relation to local bus 33
mirroring data 55
modem 78
modes
 protected 27
 real 27
 write-back 32
multimedia applications
 See also application programs
 support for 42
multitasking 30

N

network adapter
 effect on performance 21
noninterlaced image 87

- nonvolatile memory
 - description of 29
- notices 108
- numeric instructions 28

O

- odd parity 38
- operating systems
 - DOS 28
 - effect on microprocessor mode 27
 - effect on performance 20
 - OS/2 28
 - support for hard disk drives 41
 - supported by server 16
 - Symmetric Multiprocessing 16
- optical drives 41
- option-to-option cables 48
- optional security features 91
- overlapped operation 63

P

- parallel port 79
- parallel transfer 61
- parity checking 38, 100
- passwords 94, 95, 97
- PC Server 520
 - introduction to 1
 - overview of features 2
 - specifications 18
- PC ServeRAID Adapter,
 - description 62
- PCI bus
 - description of 33
- PeerMaster adapter
 - description 4
 - option diskette 67, 69
- pels 86

- performance 61
 - effect of cache memory on 32
 - effect of direct memory access on 39
 - effect of local bus on 33
 - evaluating 23
 - improving 23
- phone, warranty service 105
- physical drive 53
- physical memory
 - description of 29
- physical unit number (PUN) 49
- pipeline 26
- ports
 - illustration of 12
- power cords 110
- power-on password 94
- preface vii
- privileged-access password 95
- processor
 - See microprocessor
- protected mode 27
- protection, data 64
- public-domain software 101
- publications, related ix
- PUN (physical unit number) 49

Q

- queue depth 63

R

- RAID 54
- RAID adapter
 - See ServeRAID adapter
- RAID technology 62
 - critical status 55
 - description 53
 - levels 0-5 55—57
 - speed enhancement 55

- RAM (random access memory)
 - purpose of 29
- random access memory (RAM)
 - purpose of 29
- read-only memory (ROM)
 - purpose of 29
- real mode 27
- redundant copies, data 100
- redundant power supply 102
- refresh rates, SVGA adapter 85
- registers
 - definition of 26
 - tuning for performance 19
- related publications ix
- resolution, monitor 86
- response time 61
- responses, simultaneous 61
- ROM (read-only memory)
 - purpose of 29

S

- screen-saver programs 89
- SCSI adapter
 - description of 45
 - where to install 11
- SCSI bus
 - description of 45
- SCSI controller, integrated 42, 45
- SCSI devices
 - assigning device IDs 49
 - cabling for 51
 - how to connect 48
 - quantity supported 42, 46
- SCSI drives
 - installing additional drives 41
- SCSI interface
 - cable requirements for 51
 - description of 45
- SCSI subsystem
 - description of 45
- secured computing base 96
- security-cover option 92
- selectable drive startup 98
- self-starting diskette 98
- sequential access storage device 43
- serial port 77
- ServeRAID adapter
 - daisy-chain hook-up 46, 54
 - maximum capacity 55
 - SCSI ID address jumper 46, 54
 - where to install 11
- SIMM (single inline memory module) 29
- simultaneous responses 61
- single inline memory module (SIMM) 29
- size of machine 18
- slots
 - See expansion slots
- small computer system interface (SCSI)
 - advantages of 45
 - Common Command Set 47
- specifications for PC Server 520 18
- SRAM (static random access memory) 32
- standby power supplies 103
- static random access memory (SRAM) 32
- status indicators 6
- storage devices
 - CD-ROM drives 42
 - diskette drives 41
 - hard disk drives 41
 - hot-swap bays 42
 - tape drives 43
- storage enclosures 44

- streaming-data procedure 40
- striping 62
- subsystem
 - SCSI 45
 - SVGA 85
 - video, description of 85
- swap space 30
- system management 19
- system memory 29

T

- tape drives
 - description of 43
- technology, raid 54, 62
- temperature of machine environment 18
- termination of cables 52
- terms, glossary of 113
- theft, deterring 91
- throughput 61
- tool, data-encryption 94
- trademarks 108
- transmit and receive data 77
- trusted computing base 96
- tuning 19

U

- u-bolt 92
- uninterruptible power supplies 103

V

- vertical-refresh rate 86, 87
- video adapter
 - where to install 11
- virtual memory
 - description of 30
 - maximum amount 30

- viruses, computer 100
- volatile memory
 - description of 29
- voltage
 - minimum and maximum 18
 - variation by country 16

W

- wait state 31
- warranty service 105
- weight of machine 18
- write-back mode 32



Part Number: 78H6340

Printed in U.S.A.

May 1996

78H6340

