



PC Server 704

User's Reference

Note

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

Third Edition (March, 1997)

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.

This publication was developed for products and services offered in the United States of America. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change without notice. Consult your local IBM representative for information on the products, services, and features available in your area.

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

© Copyright International Business Machines Corporation 1995, 1997. 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.

About This Book	v
How This Book is Organized	v
Related Publications	vii
Chapter 1. Introducing Your IBM PC Server 704	1
Features at a Glance	2
What Your IBM PC Server 704 Offers	3
IBM ServerGuide	4
IBM TME 10 NetFinity	4
IBM PC Server Startup Support	4
World Wide Web	5
Server Controls	6
Status Indicators	8
Expansion Bays	10
Expansion Slots and Input/Output Ports	12
Padlock Loops and Power Supplies	14
International Capabilities	16
Operating-System Support	17
Reliability, Availability, and Serviceability Features	18
Specifications	19
Chapter 2. Server Performance	21
Understanding Server Performance	21
Network and Application Programs	21
Network Adapter	22
System Memory	22
System Microprocessor	23
Hard Disk Subsystem	23
Estimating Server Performance and Capacity	24
Improving Server Performance	25
Chapter 3. Microprocessors and Memory	27
Microprocessor	27
Microprocessor Clock	27
Microprocessor Performance	27
Microprocessor Registers and Interfaces	28
Real Mode and Protected Mode	29
Numeric Computing	30
Memory Subsystem	31
Memory Controller	31

Random Access Memory and Read-Only Memory	32
Physical Memory and Virtual Memory	33
Cache Memory	34
Chapter 4. Bus Architectures	37
EISA Bus Overview	37
PCI Bus Overview	38
Bus Parity Checking	38
EISA and PCI Shared Features	39
Expansion Bus	40
Address and Data Bus	40
Bus Master Implementation	40
Bus Arbitration	41
Chapter 5. Data Storage Devices	43
Hard Disk Subsystem	43
Internal Bays	44
Direct Access Storage Devices	46
SCSI Hard Disk Drives	46
CD-ROM Drive	46
Tape Drives	47
The IBM 24/48 GB Internal Tape Autoloader	47
Chapter 6. SCSI Subsystem	49
Advantages of the SCSI Subsystem	49
Types of Interfaces	50
SCSI ID	51
Logical Units	51
Assigning SCSI IDs	52
Cabling	53
Termination	53
Chapter 7. Disk Arrays	55
Improved System Performance	55
RAID Technology	56
Logical Drives	56
Additional Storage Capacity	57
Hard Disk Drive Capacities	58
Disk-Array Classifications	58
RAID Level 0	58

RAID Level 1	59
RAID Level 5	60
RAID Adapter Features	61
Overlapped Input/Output Operation	61
Interleave Depth/Stripe Unit	61
Queue Depth	62
Cache	62
Hot-Spare Drive Replacement	62
Data Protection	62
Chapter 8. Configuration	65
Power-On Self-Test	65
Configuration Programs	66
The SCSISelect Utility Program	68
Configuration Process Overview	69
Configuring SCSI Devices	69
Configuring EISA, ISA, and PCI Adapters	71
Chapter 9. Input/Output Ports and Connectors	75
Serial Ports	75
Parallel Port	77
Video Port	78
Keyboard and Auxiliary-Device Ports	79
UltraSCSI Ports	79
Chapter 10. Video Subsystem	81
How Monitors Operate	82
Resolution and Dot Pitch	82
Vertical-Refresh Rate	83
Selecting a Monitor	84
Using Your Monitor	85
Brightness and Contrast Controls	85
Screen-Saver Programs	85
Avoiding Fatigue	86
Chapter 11. Security Features	87
Hardware Security Features	87
Door Lock and Key	87
Padlock Loops	87
Software Security Features	88

User Password	88
Administrative Password	89
Security Subsystem Features	89
Drive-Startup Sequence	91
Securing Data	92
Protecting Data from Loss	92
Appendix A. Product Warranties and Notices	97
Warranties	97
Notices	100
Trademarks	101
Important Notices	102
Electronic Emission Notices	102
Power Cords	104
Glossary	107
Index	115

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 704.

This book does not provide specific instructions for performing procedures. That information is in your:

Setup sheet

User's Handbook

SCSI Software User's Guide

RAID adapter documentation

The documentation that comes with your ServerGuide package.

How This Book is Organized

Chapter 1, "Introducing Your IBM PC Server 704," provides an overview of the features and expansion capabilities of the IBM PC Server 704. This chapter also includes the system specifications.

Chapter 2, "Server Performance," contains information about server performance, and includes tips on improving server performance.

Chapter 3, "Microprocessors and Memory," explains microprocessors, and discusses the types of microprocessors and memory in your server.

Chapter 4, "Bus Architectures," provides information about the bus architectures in your server. These include the EISA bus and the PCI bus.

Chapter 5, "Data Storage Devices," explains the data storage capabilities of your server. It includes information about direct access storage devices (DASD), such as diskette drives, hard disk drives, CD-ROM drives, and tape drives.

Chapter 6, "SCSI Subsystem," discusses the small computer system interface (SCSI) subsystem and its advantages.

Chapter 7, “Disk Arrays,” contains information about the RAID (redundant array of independent disks) adapter available for your PC Server 704. It also provides an overview of disk array technology and RAID levels, as well as the advantages of each level.

Chapter 8, “Configuration,” provides information about the configuration process and describes the system configuration programs provided with your server.

Chapter 9, “Input/Output Ports and Connectors,” discusses the integrated serial, parallel, keyboard, mouse, video, and UltraSCSI input/output (I/O) ports on the IBM PC Server 704. This chapter also includes detailed pin assignments for each I/O connector, to help you determine if a device is compatible with your server.

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

Chapter 11, “Security Features,” explains securing network hardware, securing data on the network, and the drive startup feature.

Appendix A, “Product Warranties and Notices,” contains product warranties and special notices, such as a list of trademarks used in this book.

See the glossary in the back of this book for definitions of technical terms. An index is also provided.

Related Publications

The *IBM PC Server 704 Setup* sheet that comes with your server contains instructions for cabling your PC Server 704.

The *IBM PC Server 704 User's Handbook* that comes with your server 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 complete description of what that book contains.

The *SCSI Software User's Guide* comes with all PC Server 704 models. This manual contains detailed information about configuring devices attached to the UltraSCSI controllers on the system board.

RAID adapter documentation comes with disk array models. This documentation contains information about the PCI RAID adapter and instructions for configuring disk arrays.

The *IBM PC Server 704 Safety Information* booklet comes with all PC Server 704 manuals. This booklet provides translated copies of the safety and information notices found in your PC Server 704 manuals.

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

For a complete listing of publications available in the U.S., Canada, and Puerto Rico, call 1-800-879-2755.

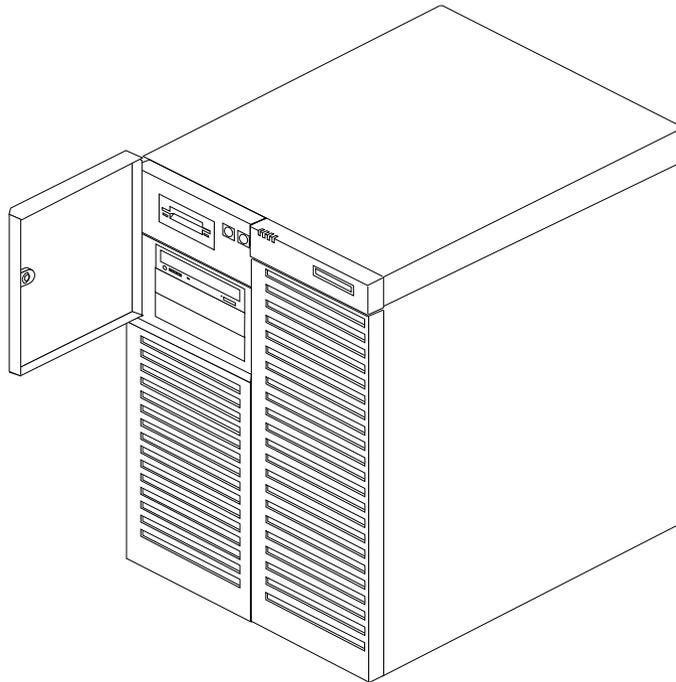
In all other countries, contact the IBM support organization that services your area, your IBM marketing representative, or your IBM authorized reseller.

Chapter 1. Introducing Your IBM PC Server 704

Your PC Server 704 is a high-performance, symmetric multiprocessing (SMP) server. It is ideally suited for networking environments that require superior microprocessor performance, efficient memory management, flexibility, and large amounts of reliable data storage.

Performance, ease of use, reliability, and expansion capabilities were key considerations during the design of your server. These design features make it possible for you to customize the system hardware to meet your needs today, while providing flexible expansion capabilities for the future.

Your PC Server 704 comes with a three-year limited warranty and IBM PC Server Startup Support. All PC Server 704 models also come with ServerGuide. The PC Server Startup Support program provides comprehensive telephone assistance for the installation of your hardware and your software during your first 90 days.



Features at a Glance

<p>Microprocessors</p> <ul style="list-style-type: none"> • Intel® Pentium® Pro with: <ul style="list-style-type: none"> – Level-1 cache, 16 KB – Level-2 cache, 512 KB • Expandable to four Pentium Pro microprocessors <p>Memory</p> <ul style="list-style-type: none"> • 256 MB of system memory, expandable to 2 GB • Industry standard, 60 ns, parity • Sixteen memory sockets • Two-way and four-way interleaved • Error correcting code support <p>Diskette Drive</p> <ul style="list-style-type: none"> • Standard: One 3.5-inch, 1.44 MB <p>CD-ROM Drive</p> <ul style="list-style-type: none"> • Standard: IDE CD-ROM Drive <p>Keyboard</p> <ul style="list-style-type: none"> • Standard: 101-key or 102-key <p>Mouse</p> <ul style="list-style-type: none"> • Standard: IBM Mouse <p>Upgradable POST and BIOS</p> <ul style="list-style-type: none"> • Flash ROMs on the system board <p>Information Panel</p> <ul style="list-style-type: none"> • Two 16-character lines • LCD display of status information 	<p>Expansion Slots and Bays</p> <ul style="list-style-type: none"> • Ten expansion slots: <ul style="list-style-type: none"> – Six 32-bit PCI slots – Four 32-bit EISA/ISA slots • Seventeen expansion bays: <ul style="list-style-type: none"> – Twelve hot-swap hard disk drive bays – Four removable-media drive bays – One dedicated, 3.5-inch diskette drive bay <p>Video</p> <ul style="list-style-type: none"> • SVGA controller • Video memory, 512 KB, expandable to 1 MB • Compatibility: <ul style="list-style-type: none"> – Color graphics adapter (CGA) – Enhanced graphics adapter (EGA) – Video graphics array (VGA) – Hercules graphics <p>Hard Disk Drives</p> <ul style="list-style-type: none"> • Number of drives and drive capacities are model dependent • Can support up to 12 hot-swap hard disk drives <p>Security Features</p> <ul style="list-style-type: none"> • Door lock • Padlock loops 	<p>Integrated Functions</p> <ul style="list-style-type: none"> • LED usability support • SVGA video connector • Two serial connectors • Parallel connector • Mouse connector • Keyboard connector • Battery-backed clock and time/date calendar • Two UltraSCSI connectors <p>Hot-Swap Power Supplies</p> <ul style="list-style-type: none"> • Two or three 420 watt supplies standard with: <ul style="list-style-type: none"> – Automatic range voltage selection (115–230 V ac) – Built-in overload and surge protection • For redundant power on two-supply models, an optional 420 watt supply with: <ul style="list-style-type: none"> – Automatic range voltage selection – Built-in overload and surge protection <p>SCSI Controller</p> <ul style="list-style-type: none"> • Two built-in bus master UltraSCSI controllers (all models) • PCI RAID adapter (some models)
---	---	--

What Your IBM PC Server 704 Offers

The unique design of your PC Server 704 takes advantage of advancements in symmetric multiprocessing (SMP), data storage, and memory management. Your server combines:

- Impressive performance using an innovative approach to symmetric multiprocessing

The PC Server 704 supports up to four Pentium Pro microprocessors. Your server contains two processor boards, and each processor board supports up to two Pentium Pro microprocessors.

- Large data-storage and hot-swap capabilities

All models of the PC Server 704 support up to 12 hot-swap hard disk drives. This *hot-swap* feature enables you to remove and replace hard disk drives without turning off the server.

- Integrated disk-array subsystem

Although many operating systems provide software fault tolerance through mirroring, IBM provides hardware fault tolerance through the redundant array of independent disks (RAID) adapter. The PCI RAID adapter comes preinstalled as a standard feature on disk-array models, and it supports RAID levels 0, 1, and 5.

- Large system memory

The memory bus in your server supports up to 2 GB (GB equals approximately 1,000,000,000 bytes) of system memory. The memory controller on the system board provides error correcting code (ECC) support for industry standard, 60 ns parity memory. Your server supports two-way and four-way interleaved memory, which is shared among all microprocessors.

Your PC Server 704 is designed to be cost effective, powerful, and flexible. It uses both peripheral component interconnect (PCI) and extended industry standard architecture (EISA) bus architectures, to provide compatibility with a wide range of existing hardware devices and software applications.

IBM ServerGuide

The IBM ServerGuide package comes with all PC Server 704 models. It 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 online publications. Refer to the documentation packaged with the ServerGuide CDs for more detailed information.

IBM TME 10 NetFinity

The IBM ServerGuide package that comes with your server gives you the option of installing IBM TME 10 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 on the ServerGuide CD and read its accompanying information.

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.

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

For more information about this exceptional program, or for assistance with the installation of your server:

- In the U.S. and Puerto Rico, call 1-800-772-2227.
- In Canada, call 1-800-565-3344.

In all other countries, contact the IBM support organization that services your area, your IBM marketing representative, or your IBM authorized reseller.

World Wide Web

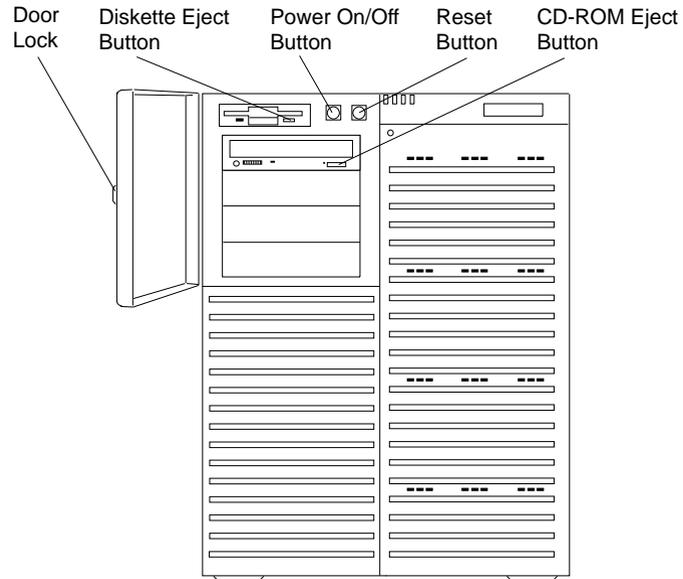
If you have access to the World Wide Web, you can obtain up-to-date information about your PC Server model and other IBM PC Server products at the following address:

<http://www.pc.ibm.com/servers/>

See your *User's Handbook* for more information.

Server Controls

The most commonly used controls on the front of the server appear in the following illustration.



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

Diskette Eject Button: Press this button to remove a diskette from the drive.

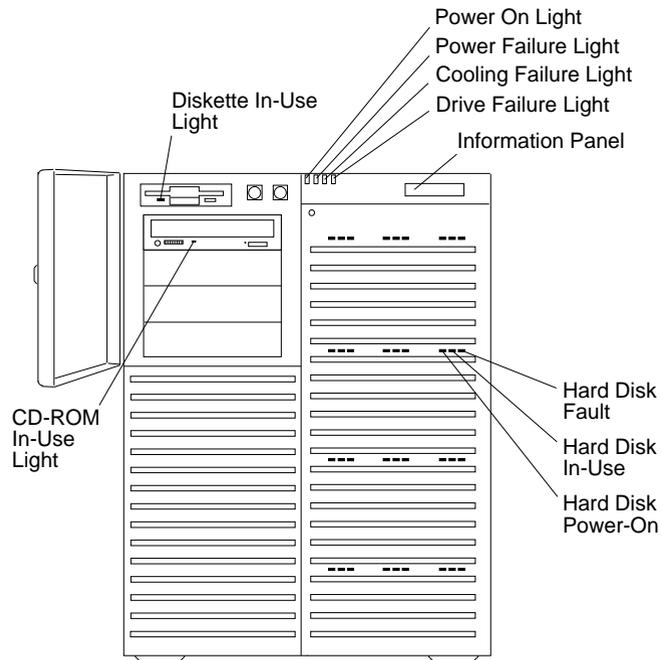
Power On/Off Button: Press this convex button to manually turn the server on or off.

Reset Button: Press this concave button to reset the system and run the power-on self-test (POST).

CD-ROM Eject Button: Press this button to eject the CD-ROM tray from its drive so that you can insert or remove a CD.

Status Indicators

The most commonly used status indicators on the front of the server appear in the following illustration.



CD-ROM In-Use Light: A green light indicates that power is present. An amber light indicates CD-ROM drive activity.

Diskette In-Use Light: This green light indicates diskette-drive activity.

Power On Light: This green light indicates that the power-on button is in the on position.

Power Failure Light: This amber light indicates there is no alternating current (ac) present.

Cooling Failure Light: This amber light indicates a fan failure.

Drive Failure Light: This amber light indicates a drive failure.

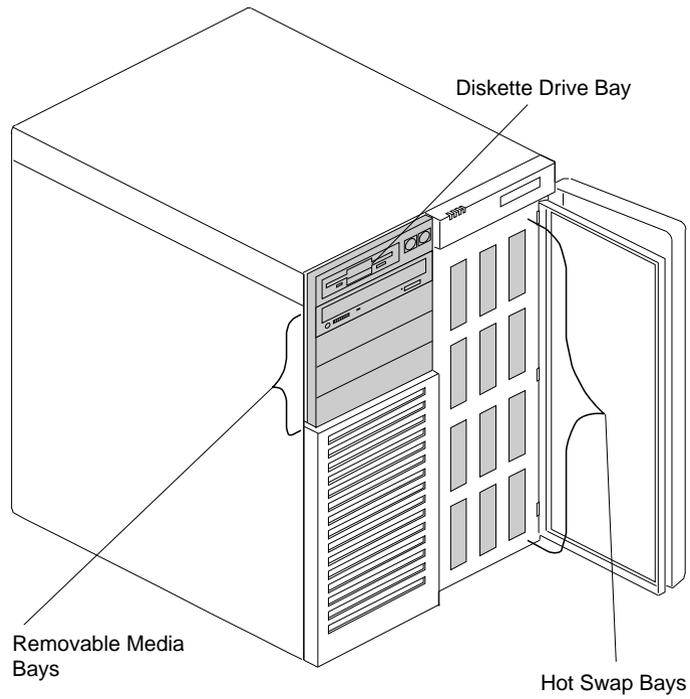
Information Panel: System monitoring information will appear on this liquid crystal display (LCD) if you install a system-management program and enable the information panel using the System Configuration Utility (SCU) program.

Hard Disk Status Lights: Each of the 12 hot-swap drive bays has a set of three status lights. These status lights indicate the following:

- **Hard Disk Fault Light:** When lit continuously, this amber light indicates a faulty hard disk drive.
- **Hard Disk In-Use Light:** When flashing, this green light indicates that the server is writing to or reading from the hard disk.
- **Hard Disk Power-On Light:** When lit continuously, this green light indicates that a drive is installed and power is present.

Expansion Bays

The following illustration shows the location of the expansion bays in your server. For installation instructions and information about the types of drives that you can install in these bays, refer to the “Installing Options” chapter in your *User's Handbook*.



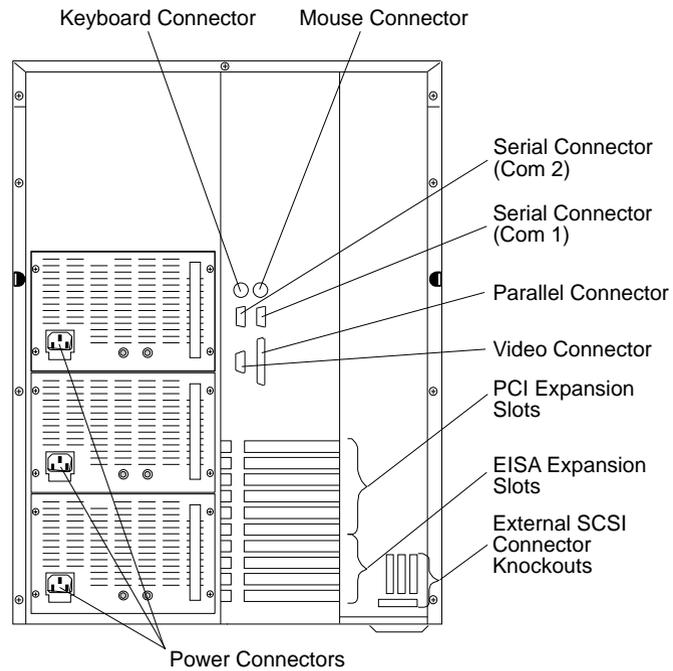
Removable Media Bays: Your PC Server 704 has four drive bays for removable-media devices, such as tape-backup drives, compact disc-read only memory (CD-ROM) drives, optical drives, or additional diskette drives. One of these drive bays contains an IDE CD-ROM drive, which comes standard in all PC Server 704 models. For information about inserting CDs and using the CD-ROM drive, refer to your *User's Handbook*.

Diskette Drive Bay: This dedicated bay contains a 3.5-inch, 1.44 MB diskette drive, which comes standard in all PC Server 704 models. This drive uses 1 MB and 2 MB diskettes (MB equals approximately 1,000,000 bytes). For optimum use, format 1 MB diskettes to 720 KB and format 2 MB diskettes to 1.44 MB.

Hot-Swap Bays: The hot-swap bays support up to 12 SCSI hard disk drives. The number of preinstalled drives and their capacities vary by model. The hot-swap feature enables you to remove and replace hard disk drives without turning off the server.

Expansion Slots and Input/Output Ports

The following illustration shows the expansion slots and the input/output connectors (ports) on the rear of the server.



Keyboard Connector: The cable from your keyboard connects here.

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

Serial Connectors: Your server has two, 9-pin serial connectors (COM1 and COM2). 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: This is where the signal cable for a parallel printer or other parallel device connects to your server.

Video Connector: The monitor signal cable attaches to this 15-pin connector.

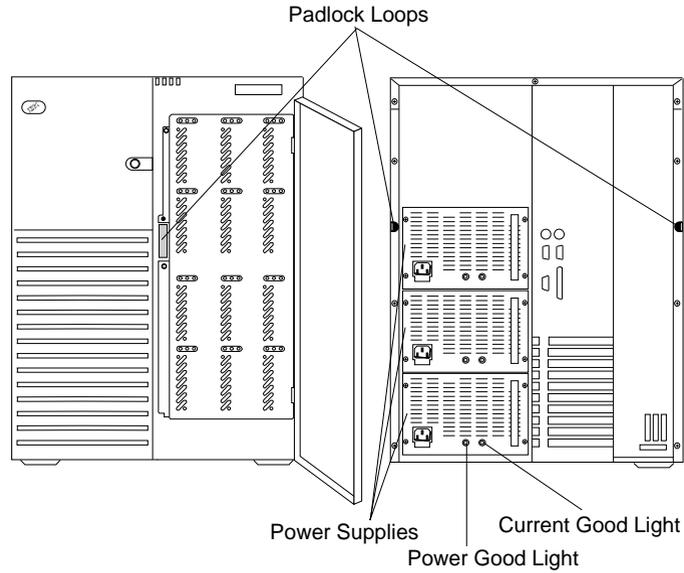
Expansion Slots: The PC Server 704 has four extended industry standard architecture (EISA) expansion slots and six peripheral component interconnect (PCI) expansion slots. You can install PCI or EISA adapters to provide communication, specialized graphics, and sound. Many adapters provide bus-master capabilities, which enable the adapters to perform operations without interrupting the system microprocessors.

External SCSI Connector Knockouts: An optional external SCSI cable can be attached to an internal SCSI connector by using one of these knockouts.

Power Connectors: The system power cords connect here. Your PC Server 704 supports up to three hot-swap power supplies.

Padlock Loops and Power Supplies

The following illustration shows the location of the padlock loops and hot-swap power supplies on your PC Server 704.



Padlock Loops: The two side cover padlock loops enable you to insert padlocks (not provided) to secure the top and side covers to the system unit. The electromagnetic interference (EMI) panel on the front of the server also has a padlock loop. You can insert a small padlock (not included) through this loop to secure the EMI panel to the system unit and help prevent unauthorized access to the hard disk drives in the hot-swap bays.

Hot-Swap Power Supplies: All PC Server 704 models support three 420 watt hot-swap power supplies. Power is drawn equally from all installed supplies. In a three-supply system, the third power supply provides redundant power. An isolating device on each power supply is designed to ensure that a failure of one supply does not affect the operation of the remaining functional supplies. And, if one of the power supplies fails in a three-supply system, the third (redundant) supply takes over immediately, so there should be no interruption to your server's availability or performance.

Note: Not all PC Server 704 models support hot-swap power supplies. You can identify a hot-swap supply by the HS on its label. If you are working with multiple models of the PC Server 704, be sure to check the power supply label before you attempt to hot-swap a power supply.

Power-Good Light and Current-Good Light: These lights provide status information about the power supply as follows:

Power-Good Light (PS)	Current-Good Light (I)	Description
On	On	Power supply on and OK
Off	Off <i>or</i> On	Power supply failure
On	Off	Current limit

International Capabilities

As always, your IBM server meets stringent world wide certifications for power, electromagnetic compatibility (EMC), and safety. It is suitable for use worldwide. The voltage-sensing feature allows the server to automatically set itself to match differing power sources, because electric power can vary from country to country.

Language is another factor to consider. English is the default setting for the basic input/output system (BIOS) messages, but you can select a different language to display using the built-in Setup program. Your server also can use special editions of operating systems that let the system and keyboard operate in a variety of languages.



Operating-System Support

You can use your server in a variety of environments to do many different tasks. To meet the requirements of these environments, your server supports several operating systems. Some of these operating systems support the high-performance, symmetric multiprocessing (SMP) capability of your server.

For the latest information about the operating systems supported by your PC Server 704:

- In the U.S. and Puerto Rico, call 1-800-772-2227.
- In Canada, call 1-800-565-3344.

In all other countries, contact the IBM support organization that services your area, your IBM marketing representative, or your IBM authorized reseller.

IBM and other manufacturers of operating systems, network programs, and application programs periodically make software updates available. These updates provide enhancements and corrections for the off-the-shelf versions of the software. To ensure that the software you install on your server functions properly, contact the manufacturers of the software that you intend to install and obtain the most current software updates.

If you intend to install an IBM operating system or network program, such as OS/2 for SMP or OS/2 LAN Server, you can obtain the most current software updates from the IBM Personal Software Solutions Center. These software updates are called ServicePaks or corrective-service diskettes (CSDs). In the U.S., you can call the IBM Personal Software Solutions Center at 1-800-992-4777.

Reliability, Availability, and Serviceability Features

Three of the most important factors in server design are *reliability*, *availability*, and *serviceability* (RAS). These factors help to ensure the integrity of the data stored on your server; that your server is available when you want to use it; and that should a failure occur, you can easily diagnose and repair the failure with minimal inconvenience.

The following is an abbreviated list of the built-in RAS features on the IBM PC Server 704. Many of these features are explained in later chapters of this reference manual, or in your *User's Handbook*.

- Menu-driven system configuration utility (SCU) programs
- Menu-driven SCSI configuration utility programs
- Menu-driven ROM based Setup programs
- Menu-driven diagnostic programs
- Power-on self-test (POST)
- Customer support center 24 hours per day*
- Cooling fans with error-sensing capability
- Hot-swap drive bays (12)
- Hot-swap power supplies (3)
- Error checking and correcting (ECC) memory
- Two- and four-way interleaved memory
- Error codes and messages
- Upgradable POST and BIOS code
- Information panel
- Automatic restart after a power failure
- Parity checking on the SCSI bus, PCI bus, keyboard interface, and serial ports
- Three cooling fans (one redundant)
- Redundant hot-swap power supply

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

Specifications

<p>Size</p> <ul style="list-style-type: none"> • Depth: 635 mm (25.0 in.) • Height: 699 mm (27.5 in.) • Width: 457 mm (18.0 in.) • Front clearance: 305 mm (12.0 in.) • Rear clearance: 127 mm (5.0 in.) • Side clearance: 76 mm (3.0 in.) <p>Weight</p> <ul style="list-style-type: none"> • Unpacked, maximum: 65.8 kg (145 lb) <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: 132 V ac - High range <ul style="list-style-type: none"> - Minimum: 180 V ac - Maximum: 264 V ac 	<p>Environment</p> <ul style="list-style-type: none"> • Air temperature <ul style="list-style-type: none"> - System on: 5° to 35° C (41° to 95° F) derated 0.5° C for every 1000 ft. (305 m) above sea level - System off: -40° to 70° C (-40° to 158° F) • Humidity <ul style="list-style-type: none"> - System on: 85%, noncondensing at +40° C (108° F) <33° wetbulb (at 40° C ambient no peripherals) - System off: 95%, noncondensing at 55° C (131° F) • Altitude <ul style="list-style-type: none"> - System on: 0 to 3050 m (10,000 ft) - System off: 0 to 15,240 m (50,000 ft) 	<p>Acoustical Noise</p> <ul style="list-style-type: none"> • Bystander position: Less than 55 dBA at +20° to +25° C (65° to 75° F) <p>Electrostatic Discharge</p> <ul style="list-style-type: none"> • Tested to 15 KV <p>Immunity</p> <ul style="list-style-type: none"> • Verified to comply with EN 50082-2 <p>Electromagnetic Emissions</p> <ul style="list-style-type: none"> • Certified to FCC 47 Class B • Tested, CISPR 22/85 Class B, EN55022 • Registered with VCCI <p>Safety</p> <ul style="list-style-type: none"> • UL 1950 • CSA C22.2 No. 950-M93 by cUL • EN 60950 by TÜV • IEC 950 by TÜV • EN 60950 and Nordic deviations by NEMKO
--	---	---

Your productivity and satisfaction, as well as the productivity and satisfaction of everyone on your network, depends on server performance. Many factors can influence this performance, including the configuration of the hardware, application workloads, operating-system configuration parameters, and the server operating environment.

The information in this chapter can help you avoid potential performance bottlenecks, and includes suggestions to help you maximize the performance of your server.

Understanding Server Performance

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

- Network and application programs
- Network adapter
- System memory
- System microprocessor
- Hard disk subsystem (UltraSCSI controllers and drives)

Network and Application Programs

Most complex programs allow you to change specific parameters to improve performance in a network environment. Always check the documentation that comes with your software to determine which parameters affect performance, and which parameters you can change to increase performance.

The original parameter settings provided by network programs usually 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.

Make sure that you record the original settings before you make any changes. In the unlikely event that performance decreases, you can change the settings back to their original values. Also document changes by making remarks, such as the date and reason for the changes, above or below the values that you change. This

information can save you valuable time if you need to make additional adjustments.

Network Adapter

Several network adapters are available for your PC Server 704. PCI bus-master network adapters are best suited for use in network servers. The primary value of using a bus-master network adapter is added capacity or, in the case of a heavily loaded server, improved performance.

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

Adding a faster network adapter provides additional throughput. The total number of transactions that the server processes per second increases, and the application response time decreases. However, keep in mind that this reduction in response time is limited. As you add more users, the effects of the faster network adapter are offset by the increase in data-space access and write-data processing (which reduces the disk-cache-hit rate and overall server performance). See “Cache Memory” on page 34 for more information.

System Memory

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

The PC Server 704 supports up to 2 GB of system memory. Your server comes with a minimum of 256 MB of system memory. You can install additional 60 ns parity memory modules in the 16 sockets on the memory board.

System Microprocessor

One way to optimize the performance of your server is to install an additional microprocessor. The PC Server 704 comes with two processor boards, and each processor board supports up to two Pentium Pro microprocessors. Each Pentium Pro microprocessor has global access to all I/O devices and has an integrated:

- Level-1 cache, 16 KB
- Level-2 cache, 512 KB
- Cache controller
- Advanced programmable interrupt controller (APIC)

Integrating the level-2 cache, cache controller, and APIC into the Pentium Pro microprocessor improves performance and alleviates some of the performance bottlenecks experienced with earlier microprocessor designs. See “Microprocessor Performance” on page 27 for more information.

Another way to improve system microprocessor performance is to lighten its load by using bus masters. Bus masters can handle tasks, such as network processing or data transfers, and provide relief for an overloaded system microprocessor. See “Bus Master Implementation” on page 40 for more information.

Hard Disk Subsystem

Two UltraSCSI PCI bus-master controllers come standard in all PC Server 704 models. These bus-master controllers significantly increase performance by handling the processing of SCSI commands and data transfers, leaving the system microprocessors free to handle other tasks.

You can install up to 12 hot-swap hard disk drives to increase disk storage and performance. Your *User's Handbook* contains instructions for installing and configuring additional drives.

Some PC Server 704 models come with a redundant array of independent disks (RAID) adapter preinstalled as a standard feature. This PCI RAID adapter enables you to configure RAID levels 0, 1, and 5 disk arrays.

Disk arrays use multiple hard disk drives to provide faster input and output rates than a single large-capacity drive. Because the files are stored on individually addressable disks, they can be found more quickly during read requests. As always, the amount of performance improvement depends on the application programs that you run on the server and the type of array that you configure. See Chapter 7, “Disk Arrays” on page 55 and your RAID adapter documentation for more information.

Estimating Server Performance and Capacity

There is no simple method 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, use a systematic approach, such as the following:

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

Improving Server Performance

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

Try to determine which subsystem is the most likely bottleneck for your environment, then concentrate on obtaining improvements for components within that subsystem. Some changes might improve the performance of one application, but degrade the performance for other applications. If this occurs, you will need to assess the trade-offs.

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

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

For example, if the bottleneck is memory, you could install additional memory; if it's the microprocessor, you could install an additional microprocessor; if it's the hard disk subsystem, you could install additional hard disk drives 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. Your server contains two processor boards, and each processor board can support two Pentium Pro microprocessors. You can install up to three more microprocessors in your server. Additional microprocessors enhance performance and provide symmetric multiprocessing (SMP) capability.

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 operations within the microprocessor.

Microprocessor Performance

The Pentium Pro microprocessor gains its high performance through its innovative design. *Dynamic execution* is the term given to the micro-architecture design of the Pentium Pro microprocessor. This dynamic design enables the Pentium Pro microprocessor to speculatively execute instructions in any order, but always complete

the instructions in their original program order. Looking 20 to 30 instructions ahead of the instruction pointer, the Pentium Pro microprocessor predicts program flow, and then analyzes the program data flow to determine the best order in which to execute the program's instructions. This approach eliminates the constraint for linear instruction sequencing and enables the microprocessor to schedule up to five instructions per clock cycle; although, a sustained rate of three instructions per clock cycle is more typical.

One of the design features that helps the Pentium Pro achieve its higher clock rate is its *pipeline* implementation. 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 speeds up instruction execution significantly when compared to a microprocessor that is running at the same speed, but does not use pipelining. The Pentium Pro microprocessor has a 12-stage, *superpipelined* implementation. This superpipelined implementation trades less work per pipestage for more stages, which results in a pipestage time that is 33 percent less than the pipestage time of a standard Pentium microprocessor.

Another design feature that enables the Pentium Pro to achieve its higher clock rate is the integration of several system components. Each Pentium Pro microprocessor in your server contains a 512 KB level-2 cache, a cache controller, and an advanced programmable interrupt controller (APIC). In earlier personal computers, these system components were not contained within the microprocessor. Instead, they were separate components on the system board, processor board, or memory board. This added level of integration improves performance by alleviating some of the performance bottlenecks experienced with earlier system designs. This integration also results in a microprocessor bus that more closely resembles a symmetric multiprocessing (SMP) system bus rather than resembling a previous generation processor-to-cache bus.

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 width of the data bus determines which type of adapters (16-bit or 32-bit) the server supports. Your server has a 32-bit data bus, and it uses a microprocessor that has a 64-bit data interface that supports both 16-bit and 32-bit adapters.

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

The size of the microprocessor *address interface*, which is the connection between the microprocessor and the address bus, determines the width of the address bus and the amount of physical memory the microprocessor can address. The Pentium Pro microprocessor in your server can address up to 4 GB of physical memory. Physical memory is described in “Physical Memory and Virtual Memory” on page 33.

Real Mode and Protected Mode

The microprocessor has a *real mode* with a 1 MB address space. 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 33).

The microprocessor also has a *protected mode* that 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) starts, 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 performs 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 within the microprocessor pass data and instructions to the FPU, which then performs the mathematical operations. Other units, such as the *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 microprocessor. Otherwise, the microprocessor will not assign the operation to 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, and cache memory.

Memory Controller

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

A 32-bit *direct-memory-access* (DMA) controller is integrated onto the system board in your PC Server 704. This DMA controller is a *bus master* that manages all DMA data transfers between system memory and I/O devices. See “Bus Master Implementation” on page 40 for more information.

Direct Memory Access

Direct memory access is a method used to transfer large amounts of data between system memory and I/O devices without the 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 have to interrupt its processing activities to manage data transfers.

Transferring data between system memory and an I/O device requires two steps. First, 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 bus to send or receive data, it competes with all the other devices that are trying to gain control of the bus, in a process known as *arbitration*. The DMA controller does

not arbitrate for control of the bus; instead, the I/O device that is sending or receiving data (the *DMA subordinate*) participates in arbitration. The DMA controller takes control of the bus when the central-arbitration-control point grants the request of the DMA subordinate. For more information about bus masters, subordinates, and arbitration, see “Bus Master Implementation” on page 40.

The functions of the memory controller vary and can include memory interleaving, which is a method of reducing the time the microprocessor has to wait for system memory to respond during memory I/O operations.

Memory Interleaving

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

The microprocessors in your server are so fast that system memory cannot respond to read and write requests as quickly as the microprocessor can send them. To reduce the time that the microprocessor spends waiting for system memory to respond, your server supports two- and four-way interleaved 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* (SIMMs). It is *volatile* memory, which means that to retain its contents, it must be constantly refreshed by an electrical current. While the server is turned off, no current is supplied to the single inline memory modules, so no data is retained in RAM.

Your server supports up to 2 GB of 60 ns parity memory. You can add memory to your server by installing matched sets of 60 ns parity memory modules in the 16 connectors on the memory board.

In addition to volatile RAM, the system board has a small area of *nonvolatile RAM* (NVRAM). Using a small current from the system battery, NVRAM retains its contents while the server is turned off. Your passwords, the time and date, and the system configuration information are stored in NVRAM.

Read-only memory (ROM) is used for storage of programs that the server uses for startup procedures and other internal operations. These programs are permanently encoded in 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.

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*.

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 physical memory, swapping portions of each program's address space to the hard disk as required.

The maximum amount of physical memory that you can install in your server, and the maximum amount of virtual memory that the operating system can allocate, are determined by the microprocessor.

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

Cache Memory

Cache memory is high-speed memory, known as static random access memory (SRAM), that can respond to memory operations without imposing *wait states* on the microprocessor. A wait state is a period of time (one microprocessor clock cycle) during which the microprocessor suspends processing and waits for system memory to respond to a read or write operation. Wait states cause the server to operate less efficiently than it would if the microprocessor were able to continue processing data during read or write operations.

To reduce the time that the microprocessor spends waiting for system memory to respond, your server uses cache memory. Cache memory improves server performance by temporarily storing frequently used data and instructions in a *cache*, a small buffer between the microprocessor and system memory.

Each Pentium Pro microprocessor in your server contains a cache controller, level-1 cache, and level-2 cache. 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, which is part of the microprocessor, consists of very fast RAM. This type of cache memory is also referred to as internal cache. The Pentium Pro microprocessor contains 16 KB (KB equals approximately 1000 bytes) of internal level-1 cache.

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 cache controller copies new data and instructions into the level-1 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 a 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 integrated cache controller quickly locates the information needed in the level-2 cache or system memory, and then copies the information into the level-1 cache. The cache controller manages the use of the cache so that the number of cache hits far exceeds the number of cache misses.

Level-2 Cache

Each Pentium Pro microprocessor in your server also contains 512 KB of level-2 cache. The level-2 cache provides an intermediate level of storage between 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; then copies it into one of the caches.

Write-Back and Write-Through Modes

For maximum performance, your PC Server 704 caches are set to the write-back mode. In the write-back mode, the microprocessor updates the cache, and then the cache controller updates system memory. While the cache controller is updating system memory, which is a relatively lengthy process, the microprocessor is free to perform other operations.

You can use the built-in Setup program to change the cache settings to the write-through mode. The write-through mode, however, is not recommended because it provides substantially reduced performance. In the write-through mode, the microprocessor updates the caches, and then updates system memory. Because the microprocessor must write directly to system memory, which runs about five times slower than the cache, performance is substantially reduced.

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.

The PC Server 704 incorporates two bus-architecture designs: the extended industry standard architecture (EISA), which supports PC compatibility and a wide range of adapters; and the peripheral component interconnect (PCI), which provides a high-throughput local bus.

You can install a variety of EISA, ISA Plug-and-Play, and PCI adapters in the *expansion slots* on the system board. These expansion slots provide paths for transferring data to and from the adapters.

The System Configuration Utility is PCI and Plug-and-Play aware, and it complies with the EISA Specification (version 3.12) and the ISA Plug-and-Play Specification (version 1.1).

EISA Bus Overview

The PC Server 704 EISA bus provides:

- Memory addressing, 32-bit
- Type A transfers at 5.33 MBps
- Type B transfers at 8 MBps
- Burst-mode transfers at 33 MBps
- Data transfers of 8-, 16-, or 32-bits
- Interrupt sharing
- Automatic translation of bus cycles between EISA and ISA masters

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. Because EISA is fully backward compatible, you can install both old and new ISA adapters and software in your server.

The System Configuration Utility (SCU) programs, supplied with your server, automatically detect and configure EISA adapters. However, with some older ISA adapters, you might need to set

switches or jumpers. The SCU programs will help you determine the necessary switch and jumper settings.

PCI Bus Overview

The PCI bus is a 32-bit local bus that provides a high-speed data path between the processor boards and peripheral devices, such as the video and hard disk subsystems. The implementation of this industry-standard bus in your server supports PCI bus speeds up to 33 MHz, and allows PCI masters direct access to system memory.

You can install PCI adapters in the six, 32-bit PCI expansion slots on the PC Server 704 system board. Your PC Server 704 has two independent PCI buses. Each bus provides three, 32-bit PCI expansion slots.

The first PCI bus operates at 33 MHz and connects the microprocessor bus to the PCI-to-EISA bridge card, and three PCI expansion slots. The second PCI bus connects the microprocessor bus to the integrated UltraSCSI controllers and three PCI expansion slots.

The PC Server 704 PCI bus architecture provides support for:

- Memory addressing, 32- and 64-bit
- Burst transfers at 133 MBps
- Data transfers of 8-, 16-, and 32-bits
- Plug-and-Play peripherals
- Interrupt control
- Arbitration for six PCI masters
- General purpose I/O ports
- Keyboard emulation
- The front-panel interface and LCD
- Security features (such as system board temperatures and intrusion-detection switches)

Bus Parity Checking

The PCI bus supports *bus parity checking*, which is a method of verifying that data was not 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 the sending device to send the data 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.

EISA and PCI Shared Features

The EISA and PCI implementations in your PC Server 704 share some fundamental features. Your server automatically configures adapters for both buses. The System Configuration Utility (SCU) programs can automatically assign interrupt levels, I/O addresses, memory addresses, arbitration levels, and other system resources to the adapters in your server.

Both buses are implemented as 32-bit buses. 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 the bus at a given time.

Both buses have several functions. In your server, each bus 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 EISA expansion bus provides four connectors for adapters; the PCI 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 a 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 EISA bus and the PCI bus both have a 32-bit address bus, enabling devices on the bus to address up to 2 GB of memory.

The width of the data bus determines how much data can transfer simultaneously. Both the EISA bus and the PCI bus in your server provides a 32-bit data path. However, the full width of the 32-bit data bus is used only when both the sending device and the receiving device are 32 bits wide.

Bus Master Implementation

A *bus master* is a device or an adapter that can take control of (own) the bus. The system microprocessor is the first owner of the bus at system startup. Other bus masters receive ownership of the bus by halting the system microprocessor. All four of the EISA expansion slots support bus-master adapters. The PCI bus provides bus-master support for six masters.

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 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 EISA and PCI 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. Every device that can arbitrate for control of the bus is assigned a priority level, known as an *arbitration level*, that is used to determine which device should control the bus next.

EISA masters arbitrate for control of the bus based on the expansion slot in which they reside. The sequence is as follows: 1, 2, 3, 4, 1, 2, 3, 4, and so on. This is because slots 1 and 2 and slots 3 and 4 share arbitration requests. The arbitration level for each PCI master is contained in a software file.

Both buses have *fairness features* to ensure that all devices get a turn to own the bus, even those 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 device so that it controls the bus more than other devices. However, if you disable the fairness feature for a device, that device 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.

IBM data storage devices are available for use with your server. You can use these storage devices to expand the data storage capacity of your server.

Hard Disk Subsystem

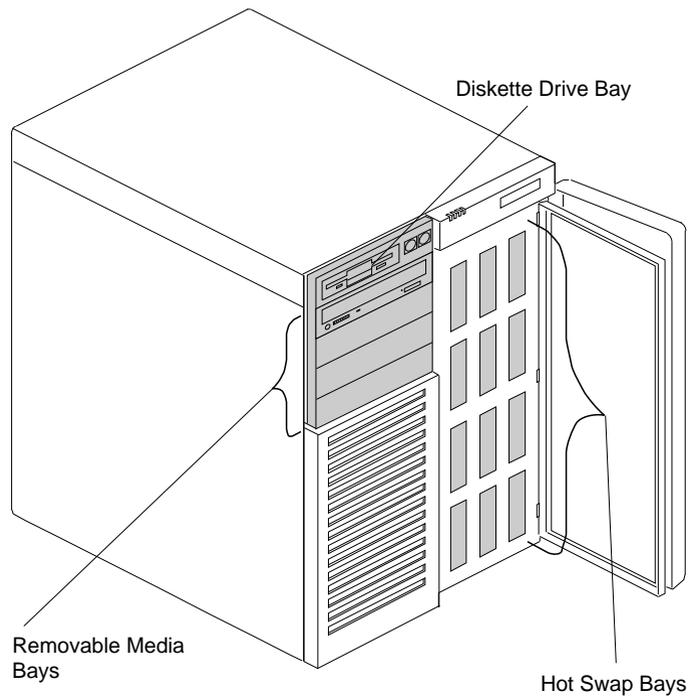
All IBM PC Server 704 models come with two bus-master UltraSCSI controllers on the system board. These controllers enable the system to communicate with devices, such as hard disk drives, CD-ROM drives, tape drives, scanners, and printers. For more information about these controllers, see Chapter 6, “SCSI Subsystem” on page 49.

Some PC Server 704 models come with a PCI RAID adapter preinstalled as a standard feature. If your server is a non-disk-array model, you can install a PCI RAID adapter and hard-disk drives to create a disk-array subsystem. See Chapter 7, “Disk Arrays” on page 55 for more information.

The hot-swap capability for hard disk drive replacement, provided in all PC Server 704 models, enables you to correct failures without turning off the server.

Internal Bays

Your server comes with one 3.5-inch, 1.44 MB diskette drive and one IDE CD-ROM drive. The number of hard disk drives installed varies by model.



Removable Media Bays: Your PC Server 704 has four drive bays for removable-media devices such as tape-backup drives, compact disc-read only memory (CD-ROM) drives, optical drives, or additional diskette drives. One of these drive bays contains an IDE CD-ROM drive, which comes standard in all PC Server 704 models. For information about inserting CDs and using the CD-ROM drive, refer to your *User's Handbook*.

Diskette Drive Bay: This dedicated bay contains a 3.5-inch, 1.44 MB diskette drive, which comes standard in all PC Server 704 models. This drive uses 1 MB and 2 MB diskettes. For optimum use, format 1 MB diskettes to 720 KB and format 2 MB diskettes to 1.44 MB.

Hot-Swap Bays: These drive bays support up to 12 SCSI hard disk drives. The number of preinstalled drives and their capacities vary by model. The hot-swap feature enables you to remove and replace hard disk drives without turning off the server.

Direct Access Storage Devices

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

SCSI Hard Disk Drives

The two UltraSCSI controllers on the system board each support up to 15 SCSI devices, including hard disk drives, CD-ROM drives, rewritable optical drives, and tape drives.

The PCI RAID adapter, that comes preinstalled as a standard feature in disk-array models, has three independent channels. Each channel supports up to 15 SCSI devices.

For a current list of the SCSI hard disk drives available for your PC Server 704, contact your IBM marketing representative or your IBM authorized reseller.

CD-ROM Drive

The *compact disc-read only memory* (CD-ROM) drive uses a compact disc (CD) that the server can read from but not write to. You can view the data, but you cannot change or update the data.

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

The IDE CD-ROM drive installed in your server supports CD-ROM Extended Architecture (XA), which is a standard for interleaving different types of data such as video, audio, and text. This enables your server to support advanced multimedia applications. Your CD-ROM drive also supports the Kodak Photo CD standard for recording pictures on CDs.

Handle a CD by the edges only. Periodically clean your CDs by wiping them with a soft, lint-free cloth, from the center of the CD to the outer edge. Do not clean CDs in a circular pattern, because this can cause loss of data.

Tape Drives

Backing up the entire contents of a hard disk to diskettes can be very time-consuming and, in the case of a network server 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 IBM 24/48 GB Internal Tape Autoloader is available for use with your server.

The IBM 24/48 GB Internal Tape Autoloader

You can install this autoloading tape-backup drive in one of the removable-media drive bays. This drive 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.

The *small computer system interface* (SCSI) enables the system to communicate with devices such as hard disk drives, CD-ROM drives, tape drives, scanners, and printers.

All IBM PC Server 704 models come with two bus-master UltraSCSI controllers on the system board. Disk-array models also come with a PCI RAID adapter preinstalled as a standard feature.

These UltraSCSI controllers support:

- A variety of devices that conform to ANSI SCSI Standards:
 - X3.131-1986 (SCSI)
 - X3.131-1994 (SCSI-2)
 - X3.253-1995 (SCSI-3 parallel-interface)
 - X3.277-1996 (SCSI-3 Fast-20 parallel-interface)
- Devices compatible with the SCSI Common Command Set.
- Removing and replacing hard disk drives without turning off the server (hot-swapping).
- Multiple physical devices.

The two UltraSCSI bus-master controllers on the system board each support a maximum of 15, 8-bit or 16-bit physical SCSI devices. The 16-bit devices support SCSI IDs 0 to 15, while 8-bit devices support only IDs 0 to 7. See your *SCSI Software User's Guide* for more information.

The PCI RAID adapter, that comes preinstalled as a standard feature in disk-array models, has three independent channels. Each channel supports up to 15 SCSI devices.

- Functions unique to a particular SCSI device (when the device manufacturer provides the appropriate device driver that supports the unique function).

Advantages of the SCSI Subsystem

The SCSI standards define the interface through which the system communicates with different types of devices. Before the SCSI standards were developed, most physical devices required their own controllers (usually adapters). These adapters 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 subsystem produces significant improvements in system performance and flexibility over other I/O interfaces. Each SCSI device has a built-in *device controller* that interprets the SCSI command set. The SCSI standards define the commands and messages used in the interface.

Several devices can operate concurrently and share the interface. This means significantly higher data-transfer rates and faster, more efficient, I/O management for network servers. The SCSI controller is the key component in the SCSI subsystem. You can connect internal and external physical devices to the same SCSI controller.

Types of Interfaces

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

Each device connected to a device-level interface has a controller that is mounted on a separate adapter, away from the physical device. Usually, there are limitations on the types of devices you can attach to the same controller. 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.

With a bus-level interface, each device has its own device controller that provides the details of how to accomplish its tasks. This type of interface uses a standardized set of commands, the Common Command Set, to issue tasks to the devices, and each device controller determines how to accomplish these tasks.

The UltraSCSI controllers impose few restrictions on the number and types of devices that you can attach to them. As new SCSI devices become available, there is no need to install adapters. You can attach many different devices, such as hard disk drives, CD-ROM drives, scanners, and printers, to the same UltraSCSI controller.

The SCSI standards allow 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. Device drivers supplied by the device manufacturer provide support for the functions that are unique to the device.

SCSI ID

The SCSI ID is also known as the *SCSI address* or *physical unit number* (PUN). This address enables the SCSI controller to distinguish one device from another.

The SCSI controller can service only one device at a time. When more than one physical device is connected to the same SCSI controller, the devices compete for support from the SCSI controller. When this occurs, the SCSI controller gives each device a priority based on its SCSI ID.

The UltraSCSI controllers in your server are preset to ID 7, and the backplanes are preset to SCSI ID 6. If your model came with one or more hard disk drives preinstalled, their SCSI IDs depend on where they are installed. See “Assigning SCSI IDs” on page 52 for information about how the backplane assigns SCSI IDs.

Logical Units

In addition to the physical unit number (PUN), each SCSI device also has at least one *logical unit number* (LUN). If a physical device supports no logical devices, its LUN is 0. Most SCSI devices support only LUN 0. However, some types of physical SCSI devices, such as *bridge controllers*, can support more than one logical device. The LUN identifies each device that is connected to the SCSI through its own device controller.

For example, a printer controller might support up to eight printers. The printer controller is a physical device, and each of the printers is a logical device. If the printer controller is set to ID 2, the ID of the first printer appears in the SCSI device configuration as PUN 2, LUN 0. The IDs of the other printers are PUN 2, LUN 1; PUN 2, LUN 2; and so on.

Assigning SCSI IDs

For the SCSI subsystem to function properly, each SCSI device installed in or attached to your server must have a unique SCSI ID.

Note: Devices attached to different controllers can use the same SCSI ID.

When you install additional hot-swap drives in bays 1 to 12, the *backplanes* (the printed circuit boards behind the bays) automatically set the SCSI IDs for the drives. Your server contains two backplanes, and each backplane supports up to six, SCSI hard disk drives.

The backplanes behind bays 1 to 12 assign the SCSI IDs as follows:

Bay	1	2	3	4	5	6	7	8	9	10	11	12
ID	8	9	10	11	12	13	8	9	10	11	12	13

Table 1. Automatically Assigned SCSI IDs

If the default SCSI ID settings assigned by the backplanes do not support your requirements, you can move the jumpers on the backplanes to change the SCSI ID assignments. See “Changing the Backplane SCSI ID Jumpers” on page 70 for more information.

Before you attach external SCSI devices or install a SCSI device in a removable-media bay, you must choose an ID for the device; then manually set the SCSI ID by following the instructions supplied with the device.

Your *User's Handbook* contains more information about setting SCSI IDs, and it also contains detailed information about the types of drives that you can install in your server.

Cabling

The SCSI subsystem in your server consists of UltraSCSI controllers, SCSI devices, and cabling. The cabling (interface) conforms to the SCSI standards. SCSI cables for 16-bit devices have a set of 16 data lines, two data-parity lines, and various control lines. SCSI cables for 8-bit devices have a set of eight data lines, one data-parity line, and various control lines.

The server comes with two 16-bit SCSI cables.

In non-disk-array models, one cable goes from the UltraSCSI 1 connector on the system board to the top backplane and the other cable goes from the UltraSCSI 2 connector to the bottom backplane.

In disk-array models, one cable goes from the channel 1 connector on the PCI RAID adapter to the top backplane and the other cable goes from the channel 2 connector to the bottom backplane.

The interface uses *daisy-chain cabling*, where each device except the last one (internal or external) has one cable going into it and another coming out of it. There are some limitations on the length of cabling that you can use. Information about these limitations is provided in “Termination,” but more detailed information is provided in your *User's Handbook*.

Termination

The SCSI standard defines the signals carried on the interface (cabling). To ensure the quality of these signals, a *terminator* must be present at both ends of the SCSI cabling. 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.

The interface uses daisy-chain cabling, where each device except the last one (internal or external) has a cable going into it and another coming out of it. The last device has an incoming cable and a terminator.

In the PC Server 704, the hot-swap backplane automatically sets the termination of the drives. Therefore, you do not need to install or remove terminators when you install or remove hot-swap drives. However, if you install a SCSI device in one of the removable-media bays, or you attach SCSI devices to an external connector on a RAID adapter, you might need to adjust the termination so that the last device in the chain is terminated.

To ensure proper signal quality, maximum cable length and termination should conform to the ANSI SCSI Standards for your SCSI devices. The following general rules apply to SCSI-device termination:

- If you install a SCSI device in a removable-media bay and you attach one or more external SCSI devices, you must terminate the devices at the ends of the internal and external cables.
- If you have only one external device, you must terminate that device.
- For specific information about termination, use the instructions supplied with the device.

For more information, refer to the “Installing Options” chapter of the *User's Handbook*.

This chapter provides an overview of disk array technology and RAID levels, as well as the advantages of each level.

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

The amount of security, performance, and improved reliability depends on the application programs that are running and the type of disk array you configure. The type of disk array you can configure depends on the data-storage patterns supported by your hard disk controller, operating system, and application programs.

Some PC Server 704 models come with a redundant array of independent disks (RAID) adapter preinstalled as a standard feature. This PCI RAID adapter enables you to configure RAID levels 0, 1, and 5 disk arrays.

If your server is a non-disk-array model, you can install a PCI RAID adapter and hard-disk drives to create a disk-array subsystem.

If your server came with a RAID adapter preinstalled, refer to the documentation that came with the adapter for detailed instructions about creating, deleting, and maintaining disk arrays.

Improved System Performance

When you group hard disk drives into a single logical drive, the RAID adapter can transfer data 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.

RAID Technology

RAID is the technology of grouping several hard disk drives in a server into an *array* that you can define as a single *logical* drive. This logical drive then appears to the operating system as a single physical drive.

RAID technology provides access and data transfer rates beyond the physical limitations of existing hard disk drives, which greatly enhances the system-drive capacity and performance.

In addition, if one of the physical drives in a RAID level-1 or RAID level-5 disk array fails, the system continues to run, with no operator intervention required. You can replace the failed drive without turning off the server (hot-swap), and the RAID adapter rebuilds the new drive contents 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: Data stored in a failed level-0 logical drive is lost.

Logical Drives

The following information will help you understand how the server manages logical drives and how many you can define.

When you create an array, you group hard disk drives into one storage area. You can define this storage area as a single logical drive, or you can subdivide it into several logical drives. A logical drive in an array can be any size you choose within the size limitations of the array. The number of logical drives and arrays that your server supports depends upon the RAID adapter installed.

If you have only one array, you can define it as a single logical drive, or you can divide it into several logical drives. If you have two or more arrays, each array can be one logical drive, or you can divide each array into multiple logical drives, as long as the total number of logical drives for all of the arrays does not exceed the number of logical drives supported by the RAID adapter.

The operating system considers each of these logical drives just as it does a physical hard disk drive. You can use the FDISK program (or its equivalent) to partition the logical drives in the same way that the operating system partitions a hard disk drive.

If you install an operating system, you can allow the installation program to determine how the FDISK program or its equivalent allots the space within a logical drive, or you can manipulate the program yourself to partition the available space. For more information, refer to your RAID adapter documentation and your operating system documentation.

Additional Storage Capacity

When you add hard disk drives to your server, you must configure a new array before you can use the drives. You can reconfigure an existing array to include the capacity offered with the new drives, or you can group the new drives into their own array. You also can create an array with only one drive. Your server has space for up to 12 hot-swappable hard disk drives. You can install additional SCSI devices in the four removable-media bays, but these bays do not support hard disk drives or hot-swapping. Depending upon your configuration, you might need to install an additional cable.

Notes:

1. 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.
2. The RAID adapter automatically assigns RAID level 0 to all logical drives in an array containing only one hard disk drive.

Hard Disk Drive Capacities

Hard disk drive capacities influence the way you create arrays. Drives in the array can be of different capacities (1 GB* or 2 GB, for example), but the RAID adapter treats them as if they all have the capacity of the *smallest* disk drive. For example, if you have four drives of 1 GB, 1 GB, 1 GB, and 2 GB grouped into one array, the total capacity of the array is 1 GB times 4, or 4 GB, not the 5 GB physically available.

If you add a smaller drive to an array of larger drives, such as a 1 GB drive to a group containing three 2 GB drives, the total capacity of that array is 4 GB, not the 7 GB physically available.

Disk-Array Classifications

With RAID technology, data is *striped* across an array of hard disk drives. This data-distribution technique complements the way the operating system requests data.

The six basic disk-array classifications are RAID-0 through RAID-5. The PCI RAID adapters available for your PC Server 704 support three, industry-recognized RAID levels: 0, 1, and 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 referring to hard drive capacity, MB stands for 1 000 000 000 bytes and GB stands for 1 000 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

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 requires two drives. When you group more than two drives into a RAID level-1 array, the RAID adapter will either automatically assign the enhanced RAID level 1 to the array, or it will allow you to assign the enhanced (0+1) level to the array.

The enhanced RAID level 1 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 the enhanced level 1. Notice that the mirror of the first data stripe (XXX YYY ZZZ) is shifted one drive. The other data stripes in level 1 follow the same pattern.

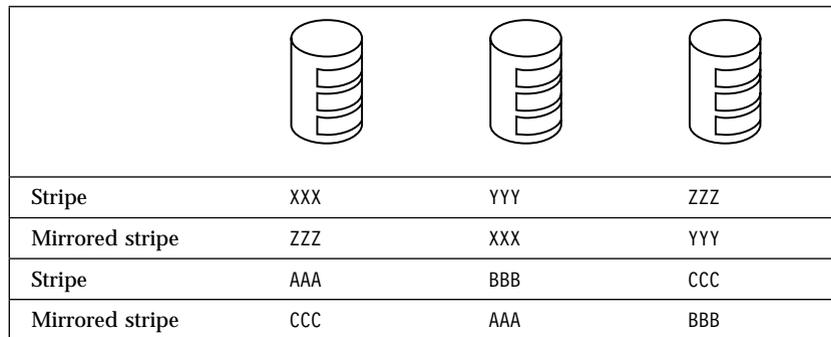


Table 2. Enhanced RAID Level 1

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.

RAID Adapter Features

The PCI RAID adapters available for your server support:

- A 32-bit PCI data bus
- A 4MB adapter cache, which significantly improves server performance by keeping frequently used SCSI-device data available for immediate transfer
- SCSI Tagged Command Queuing
- *Hot-spare drives*, which enable the disk-array subsystem to perform data recovery and data reconstruction automatically, in the event of a single disk failure in a level-1 or level-5 array

See the documentation that came with your RAID adapter for more specific information.

Overlapped Input/Output Operation

The RAID adapter available for your server provides multiple data paths to and from arrayed drives. This enables your server to respond to requests from several users simultaneously. If one user requests data that resides on the first drive in the array and a second user requests data that resides on the second drive, the RAID adapter can deliver the requested data to both users simultaneously.

Interleave Depth/Stripe Unit

The granularity at which data from one file is stored on one drive in the array before subsequent data is stored on the next drive in the array is called the *interleave depth*. For optimum system performance, you can set the interleave depth to a unit size of from 8 KB to 64 KB of data. The collection of these units in logical order, from the first drive in the array to the last drive in the array, is called a *stripe*.

Queue Depth

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

Cache

The RAID adapter available for your server has cache memory. You can configure this cache memory to operate in write-through or write-back mode on a logical-drive basis. Cache memory has parity to detect memory errors and retry algorithms to recover from errors that appear sporadically. For more information, see “Cache Memory” on page 34 and your RAID adapter documentation.

Hot-Spare Drive Replacement

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

If a hard disk drive fails in a RAID level-1 or RAID level-5 array, the RAID adapter automatically re-creates the data from the failed drive on the hot-spare drive. When you replace the failed drive and run the operating-system-specific utility program, the RAID adapter will automatically define the replacement drive as your new hot spare.

Data Protection

The RAID adapter provides the server with a data redundancy technique that distributes data across all the drives in the array. In the event of a single drive failure in a RAID level-1 or RAID level-5 array, the remaining drives in the array service read and write requests, without loss of data. The failed hard disk drive *appears* to be operating correctly, although at reduced performance.

When you replace the failed drive, the data in logical drives assigned RAID levels 1 and 5 is rebuilt on the new hard disk drive, using the parity information stored on the other hard disk drives in the array. You must replace and rebuild the failed drive in a timely manner to avoid multiple-drive failures and reduced performance. Once the RAID adapter rebuilds the drive contents, full performance and fault-tolerance functions are restored.

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 704 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)
- Enhanced industry standard architecture (EISA)
- Industry standard architecture (ISA)
- Small computer system interface (SCSI)

The System Configuration Utility is PCI and Plug-and-Play aware, and it complies with the EISA Specification (version 3.12) and the ISA Plug-and-Play Specification (version 1.1).

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.

Power-On Self-Test

Each time you turn on your server, it performs a series of tests that check the operation of the server and some options. This series of tests is called the power-on self-test (POST). The POST checks the system board, microprocessors, memory, video, keyboard, and most of the installed devices. During the memory test, the screen displays a number that represents the amount of available server memory.

Note: When you use the Ctrl+Alt+Del key sequence to restart the server, the POST does not run the memory test.

For PCI, EISA, and ISA devices, your server uses configuration programs and configuration files to assign interrupt levels,

input/output addresses, memory addresses, and other server resources. However, your system automatically recognizes and configures other options, such as memory and plug-and-play adapters.

Configuration Programs

All PC Server 704 models come with menu-driven configuration programs. These programs enable you to view or change:

- System-board settings and assignments
- PCI, EISA, and ISA device assignments
- SCSI device assignments
- Disk array assignments (some models)

The types of configuration programs provided with your server depend on your model. Step-by-step instructions for using most of the PC Server 704 setup, configuration, and utility programs are provided in your *User's Handbook*. Instructions for using the SCSI configuration programs are provided in your *SCSI Software User's Guide*. If your server came with a RAID adapter installed, see the RAID adapter documentation for configuration information.

The PC Server 704 comes with the following configuration programs:

- **System Configuration Utility Programs**

These menu-driven programs are the primary method for configuring your server and optional devices. The programs on the System Configuration Utility Diskette enable you to view or change your server's configuration information. You should run these programs when you first set up your server; when you get a configuration error during startup; or when you install, remove, or change the position of a device in the server.

Before you install a new device or program, or make changes to your existing configuration, be sure to start the System Configuration Utility programs and record the current settings. In the unlikely event that a conflict occurs, you will be able to change the settings back to their original values.

- **Built-in Setup Program**

All PC Server 704 models come with a built-in Setup program, which you can access during system startup. This program is stored in nonvolatile RAM on the system board.

You can use this program to set the date and time, set passwords, or to enable or disable the diskette drive.

The built-in Setup program also enables you to view or change system-board settings. However, you should use the System Configuration Utility programs to accomplish these tasks because they provide conflict-resolution support and access to information about ISA, EISA, and PCI devices.

- **SCSISelect Program**

The SCSISelect configuration program comes with all PC Server 704 models. It enables you to view or change the assignments of SCSI devices attached to the two SCSI controllers on the system board. 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 connected to the integrated UltraSCSI controllers. See “The SCSISelect Utility Program” on page 68 and your *SCSI Software User's Guide* for more information.

- **RAID Configuration and Monitoring Programs**

The PCI RAID adapter diskettes come with disk-array models. These diskettes contain device drivers, monitoring programs, and menu-driven configuration programs that enable you to define, create, modify, or delete your disk arrays. The programs on these diskettes also enable you to view or change the SCSI device assignments for the drives attached to the RAID adapter. See your RAID adapter documentation for more information about defining or changing your disk-array configuration.

The SCSISelect Utility Program

The UltraSCSI controllers on the system board come with a menu-driven configuration program that enables you to view or change SCSI settings. You can use the SCSISelect 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 UltraSCSI controllers from their default value of 7.
- SCSI Parity Checking
You can enable or disable parity checking on the UltraSCSI controllers. The default setting is *Enabled*.
- Host Adapter SCSI Termination
You can change the termination of the UltraSCSI controllers. 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.

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, the SCSI ID for the drive is assigned by the backplane behind the bay. 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 a removable-media bay, you will need to manually set the SCSI ID and termination for the device. Check the documentation that came with the device for instructions to set the SCSI ID and termination for the device. See “Assigning SCSI IDs” on page 52 and “Termination” on page 53 for more information.

The backplanes behind bays 1 to 12 assign the SCSI IDs as follows:

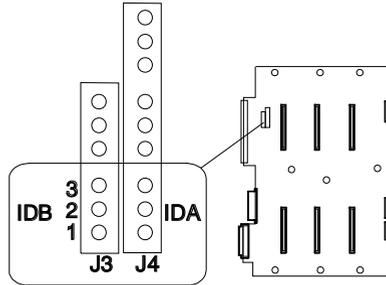
Bay	1	2	3	4	5	6	7	8	9	10	11	12
ID	8	9	10	11	12	13	8	9	10	11	12	13

Table 3. Automatically Assigned SCSI IDs

Changing the Backplane SCSI ID Jumpers

If the default SCSI ID backplane assignments do not support your requirements, you can move the J3 and J4 jumpers on the backplanes to change the SCSI ID assignments.

The top and bottom backplanes are identical, and each backplane has its own J3 and J4 jumpers. The jumpers do not need to be set the same way for each backplane. However, keep in mind that the backplanes are set to SCSI ID 6. The top backplane supports Bays 1 through 6, and the bottom backplane supports Bays 7 through 12. The following illustration and table show the location of the J3 and J4 jumpers and the alternative settings.



Jumpers		SCSI ID of Drive					
J4 IDA	J3 IDB	Bay 1 or 7	Bay 2 or 8	Bay 3 or 9	Bay 4 or 10	Bay 5 or 11	Bay 6 or 12
1-2	1-2	0	1	10	3	4	13
2-3	1-2	0	1	2	3	4	5
2-3	2-3	8	9	2	11	12	5
1-2	2-3	8	9	10	11	12	13

Table 4. Changing the Backplane SCSI ID Jumpers

Configuring EISA, ISA, and PCI Adapters

During startup, the server's basic input/output system (BIOS) code detects EISA, ISA, and PCI devices and automatically assigns system resources as follows:

1. ISA devices: If found and recognized, the server initializes and assigns resource priority over other devices of the same type to the ISA devices.
2. EISA devices: If found and recognized, the server uses predefined parameters to initialize and assign resources.
3. Optional PCI devices: If found and recognized, the server uses predefined parameters to initialize and assign resources.
4. Built-in PCI devices: The server uses predefined parameters to initialize and assign resources.

For PCI, EISA, and ISA devices, your server uses configuration programs or configuration files to assign interrupt levels, input/output addresses, memory addresses, and other server resources. In most cases, the server will operate using the default settings. However, in the event that you install a EISA or ISA device that is not recognized, you should refer to the instructions in “Configuring EISA or ISA Adapters.”

Configuring EISA or ISA Adapters

Before you install an EISA or ISA adapter in your server, be sure to check the documentation that came 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 System Configuration Utility 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.

5. Set the jumpers or switches on the new adapter.
6. Install the new adapter in an available expansion slot.
7. Use the System Configuration Utility programs to create a custom configuration file (.CFG) for the device, or to assign or verify the adapter's configuration.

Configuring PCI Adapters

Programmable configuration hardware enables your server to recognize and configure PCI adapters automatically. This programmable 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 PCI 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 passwords.

Each PCI adapter has a unique device ID. The POST routine uses the device ID to identify the type and location of each PCI adapter in the server.

The programmable configuration hardware, together with the configuration programs and configuration files, simplifies the process of configuring PCI devices by:

- Eliminating the need for setting jumpers and switches on PCI 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

Option Diskettes

The configuration files for an optional device normally are provided on an Option Diskette that comes with the device. The server uses these configuration files to assign settings that define how the device will operate. Sometimes Option Diskettes also contain device drivers and diagnostic programs for the device.

Whenever you install a device that comes with an Option Diskette, be sure to read the instructions that come with the device. Refer to the information provided with the option, the instructions provided in your *User's Handbook*, and your operating system documentation to install the device-specific files provided on the Option Diskette.

Chapter 9. Input/Output Ports and Connectors

The input/output (I/O) connectors are for attaching external devices, such as printers, keyboards, and displays, to your server.

The I/O connectors on the PC Server 704 include:

- Two serial-port connectors
- One parallel-port connector
- One video-port connector
- One keyboard-port connector
- One auxiliary-device-port connector
- SCSI-port connectors (vary by model)

Serial Ports

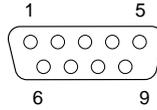
Your server comes with two serial ports. These ports are used to communicate with printers, plotters, external modems, scanners, and auxiliary terminals. These ports also enable you to transfer data between computers.

Serial ports transfer data one bit at a time, using direct memory access (DMA). DMA is a method of transferring data between I/O devices and system memory without intervention by the system microprocessor. See “Direct Memory Access” on page 31 for more information.

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

The serial ports can transmit and receive data and commands at rates of from 300 bits per second up to 345,600 bits per second. To use a serial port at 345,600 bits per second, you need a special shielded cable. For information about this cable, contact your IBM marketing representative or your IBM authorized reseller.

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



The following 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		

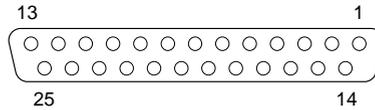
Table 5. Serial Port Pin-Number Assignments

When you turn on your server, the POST routine assigns the serial ports to specific communication port addresses.

Some application programs use only certain ports, and some modems are designed for use only at certain communication port addresses. You might need to change communication port address assignments to resolve conflicts by using the System Configuration Utility programs.

Parallel Port

The parallel port usually is used to communicate with printers, and transfers data one byte at a time using DMA. The parallel port has a 25-pin, female D-shell connector on the back of your server.



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

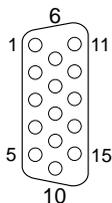
Pin	Signal	Pin	Signal
1	STROBE	14	-AUTO FEED XT
2	Data 0	15	-ERROR
3	Data 1	16	-INIT
4	Data 2	17	-SLCT IN
5	Data 3	18	Ground
6	Data 4	19	Ground
7	Data 5	20	Ground
8	Data 6	21	Ground
9	Data 7	22	Ground
10	-ACK	23	Ground
11	BUSY	24	Ground
12	PE (paper end)	25	Ground
13	SLCT (select)		

Table 6. Parallel Port Pin-Number Assignments

When you turn on your server, the POST routine assigns the parallel port a specific port address. You can change the parallel-port assignments by using the System Configuration Utility programs.

Video Port

The system board in your server has one SVGA video port. This port is used to attach a video monitor. The video port has a 15-pin analog connector on the back of the server.



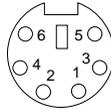
The following table shows the pin-number assignments for the video connector.

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

Table 7. Video Port Pin-Number Assignments

Keyboard and Auxiliary-Device Ports

The IBM PC Server 704 system board has one keyboard port and one auxiliary-device port that supports a mouse or other pointing device. See “Expansion Slots and Input/Output Ports” on page 12 for the locations of these connectors.



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

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

Table 8. Keyboard and Auxiliary-Device Port Pin-Number Assignments

UltraSCSI Ports

All PC Server 704 models come with two UltraSCSI bus-master controllers on the system board. Each of these controllers supports up to 15 SCSI devices. The 68-pin SCSI connectors for these controllers allow you to expand the capabilities of your server by attaching different types of SCSI devices, such as drives or printers.

The PCI RAID adapter, that comes preinstalled as a standard feature in disk-array models, has three independent channels. Each channel supports up to 15 SCSI devices.

The following table shows the pin-number assignments for the 68-pin SCSI 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 9. 68-Pin SCSI Port Pin-Number Assignments

Chapter 10. Video Subsystem

The IBM PC Server 704 comes with a super video graphics array (SVGA) controller on the system board. This controller uses a 15-pin, analog video connector. For information about the video connector, see “Video Port” on page 78.

This video controller supports high resolution graphics and fast refresh rates that meet the specifications of the 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, see “How Monitors Operate” on page 82.

The SVGA controller includes a video display buffer with 512 KB of video dynamic random access memory (DRAM). You can increase the amount of video memory in your server to 1 MB by installing an additional 512 KB (256 K by 16-bit), 70 ns CMOS DRAM module with Fast Page Mode, such as Samsung KM416C256AJ-7, Hitachi HM514260AJ7R, or equivalent.

The amount of video memory in your server determines the resolution and the number of colors supported. Additional features of the SVGA controller include:

- Integrated VGA emulation
- Compatible with CGA, EGA, and Hercules Graphics
- 132-column text mode support (with 1 MB of video DRAM)
- Screen resolutions with 512 KB of video DRAM:
 - 640 x 480 pels with up to 256 colors
 - 800 x 600 pels with up to 256 colors
 - 1024 x 768 pels with up to 16 colors
- Screen resolution with 1 MB of video DRAM:
 - 1280 x 1024 with up to 16 colors
 - 1024 x 768 with up to 256 colors
- Support for analog monitors (single and multiple frequency, interlaced and noninterlaced) with a maximum frequency of 87 Hz

Depending upon your operating system, you might need to install some video device drivers. Refer to the documentation that came with your operating system for more information.

How Monitors Operate

The most common distinguishing characteristics of monitors are resolution, dot pitch, vertical-refresh rate, 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, 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 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 might do 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. Therefore, the image has to be redrawn many times each second to refresh the phosphors. The vertical-refresh rate is the speed at which the image on the screen is redrawn.

The image is redrawn from top to bottom. By the time the bottom rows are drawn, the top rows are starting to fade. The two ways of redrawing the image are *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 provide a very stable image with little or no flicker detectable. Monitors with slower refresh rates tend to have more flicker, which is particularly noticeable with white and other light-colored backgrounds. You can minimize the flicker by changing the colors you work with. Many programs allow you to select colors.

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

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

Selecting a Monitor

Monitors come in many sizes and types. The video controller in your server supports a wide range of monochrome and color monitors with various capabilities.

IBM provides many monitors that you can use with your system. To take full advantage of the video capabilities provided by your server, use a monitor with a fast noninterlaced refresh rate, such as the IBM G40, G50, or P50 monitor.

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
G40	0.28	800x600	60 Hz noninterlaced
G40	0.28	640x480	75 Hz noninterlaced
G50	0.28	1024x768	72 Hz noninterlaced
G50	0.28	800x600	85 Hz interlaced
G50	0.28	640x480	85 Hz noninterlaced
P50	0.26	1024x768	85 Hz noninterlaced
P50	0.26	800x600	85 Hz noninterlaced
P50	0.26	640x480	85 Hz noninterlaced

Table 10. Monitor Types and Characteristics

For more detailed information about the IBM monitors available for the PC Server 704, contact your IBM marketing representative or your IBM authorized reseller.

Using Your Monitor

The following information will help you use your monitor comfortably and correctly.

Brightness and Contrast Controls

Use the brightness and contrast controls 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 the monitor.

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

Screen-Saver Programs

Leaving the monitor turned on for prolonged periods without changing what appears on the screen can have an adverse affect on the monitor. 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. This turns off the video signal when the keyboard has not been used for a defined period. When you are ready to work with the system again, pressing a predefined key restores the image to the screen. Several companies produce screen-saver programs. Contact your IBM marketing representative or your IBM authorized reseller.

Avoiding Fatigue

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

Similarly, sore muscles or joints can result from any job or activity during which you stay in a fixed position or repeat the same movement over a long period. This is not unique to viewing a monitor screen. Eye fatigue and sore muscles and joints can be reduced by modifying work practices. Take breaks, rotate tasks, and adjust the monitor, lighting, and chair height.

See your *User's Handbook* for additional information.

Security features deter unauthorized use of your server and data. If you use your server in a public environment, such as an office, you might want to protect it and the data stored on it by using one or more of the PC Server 704 security features.

Begin by evaluating your security needs. Where will the system be located? Should access to the server's internal physical devices, such as hot-swap drives, adapters, and memory be limited? Should use of the system or configuration programs be limited? When you have identified your security needs, you can activate or implement the appropriate security features.

This chapter contains information about the hardware and software security features that come with the PC Server 704. This chapter also provides information about the methods you can use to secure the data stored on your server.

Hardware Security Features

One method of deterring access is to locate critical resources, such as servers, in a limited-access area behind locked doors. Keylocks and padlocks also help prevent unauthorized access.

This section provides information about the built-in hardware security features on your PC Server 704. These features consist of a keyed door and three padlock loops.

Door Lock and Key

A door lock mechanically deters access to the drives. It controls access to the small door that covers the removable-media bays as well as the entire front cover. Always store your key in a safe place.

Padlock Loops

If your server is in an area accessible to the public, you might need a method of limiting access to the server's internal physical devices, such as hot-swap drives, adapters, and memory. Your PC Server 704 comes with three padlock loops. See "Padlock Loops and Power Supplies" on page 14 for the location of the padlock loops.

The two side cover padlock loops enable you to insert padlocks (not provided) to secure the top and side covers to the system unit.

The EMI panel on the front of the server also has a padlock loop. You can insert a small padlock (not included) through this loop to secure the EMI panel to the system unit and help prevent unauthorized access to the hard disk drives in the hot-swap bays.

Software Security Features

This section provides information about the software security features provided with your PC Server 704. These features consist of:

- User password
- Administrative password
- Features provided in the Security Subsystem

User Password

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

You do not need to set a user password to use your server, but a password helps protect the data you store on your server. Keep a record of your password in a secure place.

Your password can be up to seven alphanumeric characters. You must run the built-in Setup program to define your password; then, start the System Configuration Utility programs and enable the user-password feature. See your *User's Handbook* for instructions.

After you define and enable the user password, you must enter your password during system startup or when the keyboard inactivity timer expires. The keyboard and mouse (if attached to the mouse port) will remain locked until you enter your password.

Administrative Password

Using this security feature allows you to control who has access to the built-in Setup program and the System Configuration Utility programs. If an administrative password is set, you must enter the password before you can use the configuration programs.

After an administrative password is set and the covers are secured with padlocks, only those who know the password can alter the server's configuration settings and features.

Note: The padlocks prevent access to the system configuration switches on the system board.

Your administrative password can be up to seven alphanumeric characters. You must run the built-in Setup program to define your password; then, start the System Configuration Utility programs and enable the administrative-password feature. See your *User's Handbook* for instructions.

After you define and enable the administrative password, you must enter your password to access the built-in Setup program or the System Configuration Utility programs.

Security Subsystem Features

The PC Server 704 comes with utility programs that enable you to enable or disable some of the PC Server 704 security features. You can access most of these features through the System Configuration Utility programs. They are located in the Security Subsystem or the Boot Subsystem sections in the Peripheral Configuration Group.

In addition to the user and administrative password features, the System Configuration Utility programs provide access to the following security features:

- Keyboard-lockout timer
- Secure mode
- Drive-startup sequence

Keyboard Lockout Timer

This feature is useful if you leave your system on and unattended. With the keyboard and mouse ports locked, no one can use your system. When the keyboard and mouse ports are locked, the system automatically write protects the diskette drive.

The keyboard-lockout timer allows you to specify an inactivity time-out period of 1 to 128 minutes (in one minute intervals). If you enable the timer and there is no keyboard or mouse activity during the specified period, the following occurs:

- Keyboard and mouse input is inhibited
- A blank screen appears on the monitor
- The diskette drive (if enabled) becomes write-protected

You must enter your user password to unlock the keyboard and mouse.

Secure Mode

After you set a user password, you can enable the secure mode. This mode locks the power-on button, reset button, keyboard, and mouse; but it allows the system to start the operating system and carry out the instructions in the CONFIG.SYS and AUTOEXEC.BAT files. The keyboard and mouse remain locked until you enter the correct user password.

This mode is ideal for systems that operate unattended. If a power failure occurs, the system automatically restarts when power returns, and resumes normal operation without operator intervention. By disabling the power-on and reset buttons on the front of the server, this mode also protects against an accidental power-off or reset of the server.

When enabled, this feature prevents tampering with the configuration settings and stored data. If both user and administrative passwords are set, the secure mode prevents access to the configuration programs provided in ROM and on the System Configuration Utility Diskette. When you enable the secure mode, you can choose to enable or disable writing to a diskette. You also can activate a *hot-key* combination that allows you to put the system

into the secure mode immediately, instead of waiting for the keyboard-inactivity timer to expire.

See your *User's Handbook* for more information.

Drive-Startup Sequence

The Boot Subsystem section in the Peripheral Configuration Group 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*.

The default drive-startup sequence checks the primary diskette drive for a self-starting 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.

Note: If you change your startup sequence, be extremely careful when doing *write* operations (such as copying, saving, or formatting). You can accidentally overwrite data or programs if you select the wrong drive.

In most cases, you do not need to change the default drive-startup sequence. However, if you are working with multiple hard disk drives or multiple operating systems, you might want to change the default drive-startup sequence. The default drive-startup sequence includes the diskette drive. To increase security, you can remove the diskette drive from the startup sequence. To further increase

security after you configure your system, you can disable the diskette drive completely using the built-in Setup program.

Securing Data

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

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

To secure extremely sensitive data, you can use a commercially available data-encryption tool. These tools encode the data files so that they are unintelligible, and thus useless if stolen. You can encrypt data by using a program or using an encryption device. The software programs are usually less expensive than the hardware devices, but they are slower.

Protecting Data from Loss

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

The three primary ways to protect your data from loss are:

- Wait until the data accumulates on the server, and then make backup copies of all the hard disks.
- Configure a RAID level 1 disk array (disk mirroring) to duplicate the data as it is entered.

- Configure a RAID level 5 disk array and store the parity information about the data on the array as the data is being entered.

Backup Copies

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

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

There are two problems with making backup copies. First, someone has to remember to do it. And second, 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 (such as with a RAID level-1 disk array). 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 space needed. See “RAID Level 1” on page 59 for more information.

Parity Information

Storing parity information about the data as the data is being entered (such as with a RAID level-5 disk array) 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. See “RAID Level 5” on page 60 for more information.

Viruses

Computer viruses are a problem that exists within the personal computer community. A virus is a program (or instruction 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 or command. 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.

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

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

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

These tips will help you 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 unexpected changes in file size.
- Use password protection to limit access to your system or network.

Erased Files

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

To help ensure that no readable information is left on a hard disk, you can do a low-level format. An operating-system format operation does not remove all information from a hard disk.

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

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

Backup Power Supply

To prevent the loss of data caused by a power fluctuation, you can install a backup power supply. The two basic types of commercially available backup power supplies are *uninterruptible* and *standby*.

Uninterruptible Power Supply: This type of power supply works as an interface between the main power source and the server. It automatically takes over when the main power source is lost.

An uninterruptible power supply connects directly to the main power source (electrical outlet). The server then connects to the

power supply. The power supply transfers the power from the main power source to its internal battery, which then provides the power to the server.

Uninterruptible power supplies require no switching time (because they are always on), and they protect the server from power surges. However, they usually cost more than standby power supplies.

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

Appendix A. Product Warranties and Notices

Warranties

The following is the statement of limited warranty for the United States, Canada, and Puerto Rico.

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 your reseller. The term "Machine" means an IBM machine, its features, conversions, upgrades, elements, or accessories, or any combination of them. Unless IBM specifies otherwise, the following warranties apply only in the country where you acquire the Machine. If you have any questions, contact IBM or your reseller.

Machine: IBM PC Server 704

Warranty Period*: Three Years

**Contact your place of purchase for warranty service information.*

Production Status

Each Machine is manufactured from new parts, or new and used 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 for Machines

IBM warrants that each Machine 1) is free from defects in materials and workmanship and 2) conforms to IBM's Official Published Specifications. The warranty period for a Machine is a specified, fixed period commencing on its 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, if authorized by IBM, 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.

For IBM or your reseller to provide warranty service for a feature, conversion, or upgrade, IBM or your reseller may require that the Machine on which it is installed be 1) for certain Machines, the designated, serial-numbered Machine and 2) at an engineering-change level compatible with the feature, conversion, or upgrade. Many of these transactions involve the removal of parts and their return to IBM. You represent that all removed parts are genuine and unaltered. A part that replaces a removed part will assume the warranty service status of the replaced part.

If a Machine does not function as warranted during the warranty period, IBM or your reseller will repair it or replace it with one that is at least functionally

equivalent, without charge. The replacement may not be new, but will be in good working order. If IBM or your reseller is unable to repair or replace the Machine, 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. However, for Machines which have a life-time warranty, this warranty is not transferable.

Warranty Service

To obtain warranty service for the Machine, you should contact your reseller or call IBM. In the United States, call IBM at **1-800-772-2227**. In Canada, call IBM at **1-800-565-3344**. You may be required to present proof of purchase.

IBM or your reseller will provide certain types of repair and exchange service, either at your location or at IBM's or your reseller's service center, to restore a Machine to good working order.

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. You represent that all removed items are genuine and unaltered. The replacement may not be new, but will be in good working order and at least functionally equivalent to the item replaced. The replacement assumes the warranty service status of the replaced item. Before IBM or your reseller exchanges a Machine or part, you agree to remove all features, parts, options, alterations, and attachments not under warranty service. You also agree to ensure that the Machine is free of any legal obligations or restrictions that prevent its exchange.

You agree to:

1. obtain authorization from the owner to have IBM or your reseller service a Machine that you do not own; and
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, and
 - c. inform IBM or your reseller of changes in a Machine's location.

IBM is responsible for loss of, or damage to, your Machine while it is 1) in IBM's possession or 2) in transit in those cases where IBM is responsible for the transportation charges.

Extent of Warranty

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

The warranties may be voided by misuse, accident, modification, unsuitable physical or operating environment, improper maintenance by you, removal or alteration of Machine or parts identification labels, or failure caused by a product for which IBM is not responsible.

THESE WARRANTIES REPLACE ALL OTHER WARRANTIES OR CONDITIONS, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OR CONDITIONS OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THESE WARRANTIES GIVE YOU SPECIFIC LEGAL RIGHTS AND YOU MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM JURISDICTION TO JURISDICTION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF EXPRESS OR IMPLIED WARRANTIES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU. IN THAT EVENT SUCH WARRANTIES ARE LIMITED IN DURATION TO THE WARRANTY PERIOD. NO WARRANTIES APPLY AFTER THAT PERIOD.

Limitation of Liability

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

1. damages for bodily injury (including death) and damage to real property and tangible personal property; and
2. the amount of any other actual direct damages or loss, up to the greater of U.S. \$100,000 or the charges (if recurring, 12 months' charges apply) 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) SPECIAL, INCIDENTAL, OR INDIRECT DAMAGES OR FOR ANY ECONOMIC CONSEQUENTIAL DAMAGES (INCLUDING LOST PROFITS OR SAVINGS), EVEN IF IBM OR YOUR RESELLER IS INFORMED OF THEIR POSSIBILITY. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO YOU.

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	NetFinity
OS/2	Operating System/2
ServerGuide	

The following terms are trademarks of other companies:

Adaptec	Adaptec, Inc.
Hercules	Hercules Computer Technology
Hitachi	Hitachi Ltd.
Kodak	Eastman Kodak Company
Samsung	Samsung Electronics Co., Ltd.
SCSISelect	Adaptec, Inc.

TME 10 is a trademark of Tivoli Systems, Inc., an IBM company.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Intel and Pentium are trademarks or registered trademarks of Intel Corporation.

Other company, product, and service names may be trademarks or service marks of others.

Important Notices

Processor speeds shown indicate the internal clock speed of the microprocessor; other factors also affect application performance.

When referring to hard drive capacity, MB stands for 1 000 000 000 bytes and GB stands for 1 000 000 000 000 bytes. Total user-accessible capacity may vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives available from IBM.

Unless otherwise stated, IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Applications included with IBM products may vary from retail versions and may not include all documentation or functions. Not all products sold separately.

Electronic Emission Notices

Industry Canada Class B Emission Compliance Statement

This Class B 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 B respecte toutes les exigences du Règlement sur le matériel brouilleur du Canada.

Federal Communications Commission (FCC) Statement

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

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

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

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

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 following publications.

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, NY 10036. Definitions are identified by the symbol (A).

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).

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). These definitions are identified by the symbol (I). Definitions from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T), 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.

B

backplane. In personal computers, a printed circuit board that sets the SCSI ID and termination for hot-swap hard disk drives.

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

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

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

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

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

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

C

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

CD-ROM. Compact disc-read only memory.

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)

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.

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

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. Sometimes referred to as a pack.

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.

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

DMA. Direct memory access.

Dynamic Execution. The term given to the micro-architecture design of the Pentium Pro microprocessor. This design enables the Pentium Pro microprocessor to speculatively execute instructions in any order, but always complete the instructions in their original program order.

E

EEPROM. Electrically erasable programmable read-only memory.

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

EPROM. Erasable programmable read-only memory.

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

execution unit. In microprocessors, a logical unit that performs the arithmetic and logical operations specified by the microprocessor instruction set.

F

flash memory. See electrically erasable programmable ROM (EEPROM) modules

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

FPU. floating-point unit.

G

GB. Gigabyte.

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

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

hot swap. A hard disk subsystem feature of servers and storage enclosures that enables you to remove and replace hard disk drives without turning off the system.

I

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.

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

K

KB. Kilobyte.

KBps. Kilobytes per second.

kilobyte. 2^{10} or 1024 bytes.

L

LAN. Local area network.

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)

LUN. Logical unit number

M

math coprocessor. 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. Million bits per second.

MBps. Megabytes per second.

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

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

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

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.

multitasking. In operating systems, the running of two or more programs simultaneously.

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.

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

pack. Two or more hard disks interconnected to increase security, performance, or reliability. Commonly referred to as a disk array.

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.

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

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

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.

PUN. Physical unit number.

R

RAID. Redundant array of independent disks.

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-only memory (ROM).

Memory in which stored data cannot be modified by the user except under special conditions. See also EEPROM, EPROM, and PROM.

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

S

SCSI. Small computer system interface.

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

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

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

server. (1) A functional unit that provides shared services to

workstations over a network. (2) In a network, a data station that provides facilities to other stations.

SMP. symmetric multiprocessing.

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

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

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

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

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

T

TB. Terabyte.

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

V

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

volatile. (1) Pertaining to a storage

device whose contents are lost when power is cut off. (2) Contrast with nonvolatile.

W

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

Numerics

- 1-800 phone numbers ix, 4, 5, 17
- 101-key keyboard 2
- 102-key keyboard 2
- 115 V ac 2
- 16-bit
 - adapter 29
 - SCSI cable 53, 54
 - SCSI connector 53, 61
- 230 V ac 2
- 25-pin signal cable 13
- 32-bit
 - adapter 29
 - address 40
 - bus 37
 - data interface 29
- 512 KB level-2 cache 36
- 8-bit SCSI connector 53
- 9-pin-to-25-pin adapter 13

A

- about this book vii
- adapters
 - 16-bit or 32-bit 29
 - addition of 40
 - network 22
 - RAID 3, 55, 61
- adding storage capacity 57
- address
 - bus 40
 - COM port 76
 - interface of the
 - microprocessor 29, 34
 - SCSI 51
- administrative password 89
- advanced configuration options 68
- air temperature 19
- altitude, maximum 19

American National Standards
Institute (ANSI) 49

arbitration

- description 31
- EISA 41
- PCI 41

architecture

- EISA 3, 37
- PCI 3, 38
- SCSI 49

array

- capacity 58
- classifications 58
- number supported 56
- one drive 57
- redefining space 57, 60

autoloader, tape-backup drive 47

auxiliary-device connector 12

B

- backplane, SCSI ID 51, 52, 69
- backup
 - copies 92, 93
 - power supply 95
- basic input/output system (BIOS)
 - memory reserved for 33
 - upgradable 2
- bays, expansion 2, 10, 44
- board, memory 22
- board, processor 23
- boot subsystem 91
- bridge controller 51
- brightness and contrast controls 85
- built-in cache 36
- bulletin-board services 95
- bus
 - arbitration 39
 - control of 41
 - EISA 3, 37

- bus (*continued*)
 - ISA 37
 - level interface 50
 - memory 3
 - parity checking 38
 - PCI 3, 37, 38
 - types of 37
- bus master
 - capability 13
 - description 40
 - network adapters 22
 - SCSI controllers 23
- button, power-reset 7

C

- cable
 - power cords 104
 - SCSI 53, 54
- cache
 - description of 34
 - first-level miss 36
 - hit 35
 - level-1 34, 35
 - level-2 34, 35
 - memory 2, 34
 - miss 35
 - second-level miss 36
 - write-back mode 36
 - write-through mode 36
- calendar, battery-backed 2
- CD-ROM
 - drive 2, 46
 - eject button 7
 - extended architecture (XA) 46
 - in-use light 9
 - package, ServerGuide 4
- changing
 - port assignments 76

- chapter descriptions vii
- checking, bus parity 38
- clock
 - battery-backed 2
 - cycle 27
 - speed 27
- Common Command Set, SCSI 49
- communication port 76
- compatibility, video 2, 81
- computer viruses 94
- computing, numeric 30
- configuration
 - conflicts 67
 - description 65
 - EISA adapters 37
 - EISA devices 71
 - errors 67
 - ISA devices 71
 - PCI devices 72
 - peripheral, security 89
 - process 69
 - programs 66
 - RAID 69
 - restricting access 89
 - SCSISelect 67, 68, 69
- configuring your server 66
- connectors
 - 68-pin SCSI 79
 - auxiliary device 79
 - external SCSI 79
 - input/output 12, 75
 - internal SCSI 79
 - keyboard 79
 - parallel port 77
 - serial port 75
 - video 78
- contrast and brightness controls 85
- controller
 - bridge 51
 - bus 41

- controller (*continued*)
 - cache 34
 - memory 31
 - UltraSCSI 2, 46, 49
 - video 2, 81
- controls
 - brightness and contrast 85
 - server 6
- cooling-failure light 9
- cords, power 104
- CPU
 - See* microprocessor
- cycle, clock 27

D

- daisy-chain cabling, SCSI 53
- DASD (direct access storage device) 46
- data
 - bus 29, 40
 - Direct Memory Access (DMA) 31
 - encryption tool 92
 - interface, of the
 - microprocessor 29
 - mirroring 59
 - parallel transfer 55
 - parity 60
 - protection 62
 - redundancy technique 62
 - securing 92
 - security risk 95
 - storage devices 43
 - stripe 59
 - transfer 31
 - transfer rate, EISA 37
- data error detection 38
- default drive-startup sequence 91

- definitions
 - bus 37
 - cache 34
 - clock cycle 27
 - clock speed 27
 - disk array 23, 55
 - microprocessor 27
 - numeric computing 30
 - RAM (random access memory) 32
 - register 29
 - ROM (read-only memory) 33
 - system memory 33
 - terms, glossary of 107
- demonstration program 4
- descriptions
 - arbitration 31
 - bus master 40
 - cache memory 34
 - configuration 65
 - EISA and ISA bus 37
 - memory cycle 32
 - memory interleaving 32
 - memory subsystem 31
 - microprocessor 27
 - PC Server 704 1
 - PCI bus 37, 38
- device
 - CD-ROM drive 46
 - controller, SCSI 50
 - DASD (direct access storage device) 46
 - encryption, data 92
 - interface 50
 - physical and logical SCSI 51
 - priority, SCSI 51
 - SCSI hard disk drives 46
 - SCSI, physical and logical 51
 - tape drives 47

- direct access storage device
 - See DASD
- disable diskette drive 92
- disk array
 - adapter 61
 - classifications 58
 - configuring 23
 - internal subsystems 55
 - one drive 57
 - subsystem 43
- diskette
 - drive eject button 7
 - drive in-use light 9
 - drive, disabling 92
 - drives 2, 46
 - formatting 95
 - option 73
 - self-starting 91
- display
 - See monitor
- DMA controller 31
- door lock 2, 7, 87
- DOS, use of real mode by 30
- dot pitch, monitor 82
- drive
 - autoloader tape-backup 47
 - CD-ROM 2, 46
 - DASD 46
 - diskette 2
 - failed 56, 62
 - formatting 95
 - hard disk 2, 46
 - hot-spare 61
 - hot-swap 2, 3, 43, 55, 56
 - locations 10
 - logical 55
 - physical 55
 - SCSI hard disk 2, 46
 - sequential access 47
 - startup sequence 89, 91

- drive (*continued*)
 - tape 47
- drive-failure light 9
- drive-startup sequence 91
- duplicate data 59, 93
- dynamic execution 27

E

- EEPROM (electrically erasable programmable read-only memory) 33
- EISA (extended industry standard architecture)
 - adapters, configuring 37
 - bus architecture, overview 37
 - devices, configuring 71
- eject button
 - CD-ROM 7
 - diskette drive 7
- electrical input 19
- electrically erasable programmable read-only memory (EEPROM) 33
- electronic emission notices 102
- encryption, data tool 92
- enhanced RAID level-1 59
- environmental specifications 19
- erased files 95
- even parity 38
- execution, dynamic 27
- expansion
 - bays 2, 10
 - bus 40
 - EISA 40
 - PCI 40
 - slots 2, 12
- extended architecture
 - See XA
- extended industry standard architecture

- extended industry standard
 - architecture (*continued*)
 - See EISA (extended industry standard architecture)
- external
 - physical device 50
 - SCSI device 50

F

- fairness feature, arbitration 41
- fatigue, avoiding 86
- fault-tolerance 63
- FDISK command 56
- features
 - at a glance 2
 - availability 18
 - bus parity checking 39
 - disable the keyboard 90
 - disable the mouse 90
 - disable the power-on button 90
 - disable the reset button 90
 - diskette drive 2
 - door lock 87
 - drive startup sequence 89, 91
 - fairness 41
 - hard disk drive 2
 - hot-key, secure mode 90
 - hot-swap drives 43
 - keyboard lockout timer 90
 - level-2 cache 35
 - lockout timer, keyboard 89
 - memory 2, 27
 - microprocessor 2, 27
 - padlock loops 87
 - passwords 88
 - RAID 56, 61
 - RAS 18
 - reliability 18
 - secure mode 89, 90

- features (*continued*)
 - security 2, 14, 87
 - serviceability 18
 - software-controlled, security 89
 - special 3
 - SVGA subsystem 81
- files, erased 95
- first-level cache miss 36
- flicker, reduced monitor 81
- floating-point unit 30
- format 95
- FPU 30
- functions, integrated 2

G

- glossary 107
- guidelines, SCSI termination 54

H

- hard disk
 - adding 57
 - capacity, in an array 58
 - failure 62
 - fault light 9
 - formatting 95
 - hot-spare 62
 - in-use light 9
 - interface 50
 - number supported 2
 - power-on light 9
 - SCSI drives 46
 - subsystem 23
- hot swap
 - drive bays 2, 3
 - drives 52, 57
 - feature 43, 49
- hot-key, secure mode 90
- hot-spare drives 61, 62

humidity 19

I

ID, SCSI 51
identification number, SCSI 51
important notices 102
improving server performance 25
in-use lights 9
industry standard architecture
 See ISA (industry standard architecture)
information panel 9
input/output (I/O)
 auxiliary-device port 79
 keyboard port 79
 mouse port 79
 overlapped operation 61
 parallel port 77
 ports 12
 SCSI port 79
 serial port 75
 video port 78
integrated
 functions 2
 systems management 4
interface
 bus-level 50
 device-level 50
 SCSI 49
interlaced image 83
interleave depth/stripe unit 61
interleaved memory 2, 32
interleaving data 46
internal
 bays 44
 disk-array subsystem 55
 physical device 50
 SCSI device 50

international capabilities 16
International Organization for
 Standardization (ISO) 81
Internet 5
ISA (industry standard architecture)
 bus, description 37
 devices, configuring 71

K

keyboard
 connector 12, 79
 disable 90
 international support 16
 locking 88
 lockout timer 89, 90
 port 2, 12, 79
 standard 2
keyboard lock 2
keys, door lock 87
Kodak Photo CD 46

L

LAN (local area network)
 management program 4
language support 16
level-0 RAID 58
level-1 cache 34, 35
level-1 RAID 59
level-2 cache 35, 36
level-5 RAID 60
light-emitting diodes (LEDs) 8
lights, status indicators 8, 15
lock
 door 2, 7, 87
 keyboard 2, 88
 mouse 88
 padlock loops 2, 14, 15, 87
locking the keyboard 88

- lockout timer, keyboard 89
- logical
 - drives 55, 56
 - unit number (LUN) 51
- low-level format 95
- LPT assignments 77
- LUN (logical unit number) 51

M

- maintenance manual ix
- manual, maintenance ix
- maximum supported memory 22, 32
- memory
 - board 22
 - bus 3
 - cache 2, 34
 - controller 31
 - cycle 32
 - definition of 33
 - electrically erasable programmable read-only (EEPROM) 33
 - installing additional 22
 - interleaving 32
 - maximum supported 22, 32
 - module, single inline 32
 - nonvolatile 33
 - physical 33
 - single inline memory module 32
 - speed 32
 - standard 2
 - static random access (SRAM) 34
 - subsystem 31
 - test 65
 - type supported 32
 - video 81
 - virtual 33, 34
 - volatile 32

- microprocessor
 - address interface 29, 34
 - clock 27
 - data interface 29
 - description 2, 27
 - Pentium Pro 2, 27
 - performance 27
 - protected mode 29
 - real mode 29
 - register 29
- mirroring data 59
- mode, secure 89, 90
- model features 2
- modem 76
- monitor
 - blank screen 90
 - brightness and contrast controls 85
 - characteristics 84
 - connector 12
 - description 82
 - dot pitch 82
 - international 16
 - port 12
 - selecting 84
 - signal cable 13
 - supported 81
 - vertical-refresh rate 83
 - video port 2
- mouse
 - connector 79
 - disable 90
 - locking 88
 - port 2, 12, 79
 - standard 2
- multiprocessing
 - boards 3
 - operating systems 17
- multitasking 33

N

- NetFinity 4
- network
 - adapter, performance 22
 - backup copies 93
 - bus master adapters 22
 - performance 21
 - redundant copies, data 93
 - security 92
- network adapter 22
- noise emission values 19
- nominal operating current 19
- noninterlaced image 83
- noninterleaved memory 32
- nonvolatile
 - memory 33
 - RAM 33
- notices 100
- numeric computing 30

O

- odd parity 38
- on/off button 7
- one-drive array 57
- open bays 44
- operating system
 - international support 16
 - multiprocessing 17
 - support 17
 - use of protected mode by 30
- Operating System/2 30
- optional power supply 2
- options
 - direct access storage devices (DASD) 46
 - diskettes 73
 - SCSI hard disk drives 46
 - security features 87
 - tape drive 47

- overlapped operation 61
- overload protection 2

P

- padlock loops 2, 14, 15, 87
- panel, information 2, 9
- parallel
 - connector 77
 - port 2, 12, 77
 - transfer 55
- parity
 - checking 18, 38
 - information 93
- password
 - administrative 89
 - network programs 92
 - securing data 92
 - user 88
- PC Server Startup Support 1, 4
- PCI (peripheral component interconnect)
 - bus, description 37
 - devices, configuring 72
- pels 82
- performance 55
 - estimating 24
 - improving 25
 - microprocessor 27
 - server 21
 - system 50
- peripheral component interconnect
 - See PCI (peripheral component interconnect)
- peripheral configuration group 89, 91
- phone numbers, 1-800 ix, 4, 5, 17
- physical
 - devices 50
 - drives 51, 55

- physical (*continued*)
 - maximum supported 34
 - memory 29, 33, 34
 - SCSI devices 49
 - unit number (PUN) 51
 - pin assignments
 - auxiliary-device port 79
 - keyboard port 79
 - parallel port 77
 - SCSI port 80
 - serial port 76
 - video port 78
 - pipeline 28
 - pipestage 28
 - ports, input/output 12
 - POST (power-on self-test) 2, 33, 65
 - power
 - 420 watt supplies 2
 - connectors 12
 - cords 104
 - current good light 14, 15
 - on and off button 7
 - power good light 14, 15
 - reset button 7
 - supplies 14, 15
 - supplies, hot-swap 2
 - supply, backup 95
 - supply, standby 96
 - supply, uninterruptible 95
 - surge protection 2, 95
 - power-failure light 9
 - power-on button, disable 90
 - power-on light 9
 - power-on self-test (POST) 2, 33, 65
 - power-reset button 7
 - preface vii
 - priority, SCSI device 51
 - process, configuration 69
 - processor
 - See* microprocessor
 - processor board 3, 23
 - programmed I/O 31
 - programs
 - application, performance 21
 - configuration 66
 - network, performance 21
 - PC Server Startup Support 1, 4
 - RAID configuration 67, 69
 - restricting access 89
 - SCSISelect 67, 68, 69
 - SCU 66
 - security 88, 92
 - Setup 67
 - System Configuration Utility 66
 - protected mode 29
 - protection, data 62
 - public-domain software 95
 - publications, related ix
 - PUN (physical unit number) 51
- Q**
- queue depth 62
- R**
- RAID (redundant array of independent disks)
 - adapter 3, 55, 61, 62
 - classifications 58
 - configuration program 67
 - controller 61
 - critical status 58
 - level with single drive 59
 - levels 55, 58, 60, 61
 - maximum capacity 59
 - monitoring program 67
 - speed enhancement 58
 - storage capacity 57
 - technology 56

- RAM
 - definition of 32
 - maximum supported 32
 - nonvolatile 33
- random access memory
 - See* RAM
- RAS features 18
- real mode 29
- recording technique, CD 46
- redefining space in an array 57, 60
- redundant array of independent disks
 - See* RAID
- redundant copies, data 93
- refresh rates, video controller 81
- register
 - memory 31
 - used to store data 29
- related publications ix
- removable media drives 46
- reset button, disable 90
- resolution, monitor 82
- responses, simultaneous 55
- restricting access to the SCU 89
- ROM (read-only memory)
 - definition of 33
 - electronically erasable programmable (EEPROM) 33

S

- screen-saver programs 85
- SCSI
 - 68-pin port 79
 - adapters 61
 - address 51
 - advantages 49
 - bus-level interface 50
 - cabling system 53
 - CD-ROM drive 46

SCSI (*continued*)

- Common Command Set (CCS) 49
- connector 79
- controller 46
- daisy-chain cabling 53
- devices 49
- drive 10
- external device port 79
- guidelines 52
- hard disk drives 46
- identification number (ID) 51, 52
- interface 49
- internal device port 79
- logical unit number (LUN) 51
- port 79
- priority, device 51
- programs, configuration 67, 68, 69
- programs, RAID 69
- subsystem 49
- system performance 50
- termination 53, 54
- second-level cache miss 36
- secure mode 89, 90
- securing
 - data 92
 - hardware 87
- securing hardware 87
- security
 - blank monitor screen 90
 - disable the keyboard 90
 - disable the mouse 90
 - disable the power-on button 90
 - disable the reset button 90
 - door lock 2, 7, 87
 - drive startup sequence 89, 91
 - features 2, 14, 87
 - hot-key, secure mode 90

- security (*continued*)
 - keyboard lock 2
 - lockout timer, keyboard 89
 - monitor screen, blank 90
 - padlock loops 2, 87
 - passwords 88
 - secure mode 89, 90
 - software features 88, 89
 - subsystem group 88, 89
- self-starting diskette 91
- serial
 - connector 75
 - port 2, 12, 75
- server
 - backup 93
 - controls 6
 - data, securing 92
 - disk arrays 23, 55
 - features 2
 - performance 21, 24
 - power cords 104
 - redundant copies, data 93
 - size, server 19
 - specifications 19
- ServerGuide, description 4
- Setup program, built-in 67
- simultaneous responses 55
- sine-wave input 19
- single inline memory module 32
- size, server 19
- slots, expansion 2, 12, 38, 40
- small computer system interface
 - See* SCSI
- software
 - compatibility determined by
 - register size 29
 - security features 88
 - special features 3
 - specifications, server 19
 - speed, clock 27
 - standby power supply 96
 - startable diskette 91
 - startup sequence, drives 89, 91
 - startup support 4
 - static random access memory (SRAM) 34
 - status indicators 8
 - storage devices
 - data 23, 46, 55
 - direct access (DASD) 46
 - striping 58
 - subsystem, security 88, 89
 - subsystems
 - memory 31
 - SCSI 49
 - SVGA 81
 - video 81
 - superpipelined 28
 - supplies, power 14, 15
 - surge protection 2
 - SVGA video controller 2, 78, 81
 - symmetric multiprocessing
 - boards 3, 23
 - operating systems 17
 - system
 - architecture 37
 - configuration programs 66
 - information panel 9
 - microprocessor 23
 - performance 40, 50, 55
- System Configuration Utility 66

T

- tape drives 47, 93
- technology, RAID 56
- terms, glossary of 107
- test, memory 65

theft, deterring 87
throughput 55
tool, data-encryption 92
trademarks 101
transfer, parallel data 55
transmit and receive data 75

U

uninterruptible power supply 95
unit, floating-point 30
upgradable POST and BIOS 2
user password 88

V

vertical-refresh rate 83
video
 compatibility 2, 81
 connector 2, 12, 78, 81
 device drivers 81
 International Organization for
 Standardization (ISO) 81
 monitor 81
 port 12, 78
 SVGA 2, 81
virtual memory 33, 34
viruses, computer 94
volatile memory 32
voltage 16, 19

W

wait state 34
warranty Service 97
weight 19
World Wide Web 5
writable memory 33
write-back mode, cache 36
write-through mode, cache 36

X

XA (extended architecture),
 CD-ROM 46



Part Number: 30F6952

Printed in U.S.A.

March, 1997

30F6952

