


RS/6000 43P 7043 Models 150 and 260 Handbook

Volker Haug, Adnan Ikram, Heather Osbourne



International Technical Support Organization

<http://www.redbooks.ibm.com>

SG24-5144-00



International Technical Support Organization

RS/6000 43P 7043 Models 150 and 260 Handbook

October 1998

Take Note!

Before using this information and the products it supports, be sure to read the general information in Appendix B, "Special Notices" on page 247.

First Edition (October 1998)

This edition applies to the RS/6000 43P 7043 Models 150 and 260 workstations and workgroup servers for use with AIX Version 4.2.1, AIX Version 4.3.2, and subsequent releases of the operating system.

Comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. JN9B Building 045 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1998. All rights reserved

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

Contents

Figuresix
Tablesxi
Prefacexiii
The Team That Wrote This Redbookxiii
Comments Welcomexv
Chapter 1. RS/6000 Introduction	1
1.1 RS/6000 Workstation Offerings	1
1.2 RS/6000 43P 7043 Model 140 Overview	3
1.3 RS/6000 43P 7043 Model 150 Overview	3
1.3.1 RS/6000 43P 7043 Model 150 Key Attributes	4
1.3.2 RS/6000 43P 7043 Model 150 Feature Summary	5
1.4 RS/6000 43P 7043 Model 260 Overview	6
1.4.1 RS/6000 43P 7043 Model 260 Key Attributes	7
1.4.2 RS/6000 43P 7043 Model 260 Feature Summary	8
1.5 Performance Results	9
1.5.1 SPEC95 Software Benchmark	10
1.5.2 Performance Summary	13
1.6 Industry Trends and Directions	13
1.6.1 Microprocessor Future	13
1.6.2 Universal Serial Bus	14
1.6.3 The Accelerated Graphics Port	14
1.6.4 IEEE 1394	15
1.7 Workgroup Conferencing	15
1.8 Year 2000	18
Chapter 2. Hardware Overview	21
2.1 The RS/6000 Platform Architecture	21
2.1.1 The RS/6000 Platform Architecture Introduction	21
2.1.2 Platform Topology	22
2.1.3 RS/6000 Platform Architecture Achievements	24
2.1.4 RPA Minimum System Requirements	25
2.2 The Hardware Design	26
2.2.1 Hardware Design for the RS/6000 43P 7043 Model 150	26
2.2.2 Hardware Design for the RS/6000 43P 7043 Model 260	29
2.3 Bus Architecture	32
2.3.1 The PCI Bus Architecture	32
2.3.2 The ISA Bus Architecture	33
2.3.3 The 60X and 6XX Bus	34

2.3.4	The 6XX-MX I/O Bus	34
2.3.5	The X-Bus	35
2.4	The Processor Subsystem	35
2.4.1	The POWER3 Microprocessor	35
2.4.2	The PowerPC 604e Microprocessor	37
2.5	SMP Hardware Characteristics	39
2.5.1	Memory Hierarchy	39
Chapter 3. Detailed Product Descriptions		49
3.1	IBM RS/6000 43P 7043 Model 150 Product Description	49
3.1.1	Standard Features	50
3.1.2	Reliability, Availability, and Serviceability (RAS)	56
3.1.3	Supported Optional Features	58
3.1.4	RS/6000 43P 7043 Model 150 Adapter Placement Guidelines	66
3.1.5	Service Package	67
3.2	RS/6000 43P 7043 Model 260 Product Description	68
3.2.1	Model 260 System Architecture	69
3.2.2	Standard Features	69
3.2.3	Reliability, Availability, and Serviceability (RAS)	79
3.2.4	Supported Optional Features	82
3.2.5	RS/6000 43P 7043 Model 260 Adapter Placement Guidelines	92
3.2.6	Service Package	94
3.3	IBM Multimedia Kit for RS/6000	95
Chapter 4. Storage Architectures and Devices		97
4.1	Storage Boot Devices	97
4.2	SCSI Overview	97
4.2.1	SCSI-I	97
4.2.2	SCSI-II	98
4.2.3	SCSI-III	99
4.2.4	SCSI Terminology	102
4.2.5	SCSI Repeaters	104
4.2.6	Summary of SCSI Specifications	105
4.3	SCSI Cabling	105
4.3.1	General SCSI Cabling Considerations	106
4.3.2	Cabling the PCI Single-Ended Ultra SCSI Adapter	109
4.3.3	Cabling the PCI Differential-Ended Ultra SCSI Adapter	113
4.4	Serial Storage Architecture Overview	118
4.5	Comparison between SCSI and SSA Architecture	119
4.6	Internal SCSI Devices	120
4.6.1	Disk Drives for the Model 150 and Model 260	120
4.6.2	Internal CD-ROM Drive	121
4.6.3	Internal Tape Drives	122

4.7	Internal SSA devices	123
4.8	PCI Storage Adapters	124
4.8.1	SCSI Adapters	124
4.8.2	SSA Adapter	124
4.8.3	Storage Adapter Placement Guidelines for Model 150 and 260	126
4.9	External SCSI and SSA Storage Devices	126
Chapter 5. Adapters, Overview and Configuration		127
5.1	Communication Adapters	127
5.1.1	Token-Ring Adapter	127
5.1.2	Ethernet Adapters	128
5.1.3	ISDN Adapters	128
5.1.4	FDDI Adapters	129
5.1.5	ATM Adapters	131
5.1.6	Multiple Port Adapters	133
5.1.7	WAN Adapters	133
5.1.8	Miscellaneous Adapters	134
5.2	Graphics Adapters	135
5.2.1	POWER GXT120P (# 2838)	135
5.2.2	POWER GXT250P (# 2851) and POWER GXT255P (# 2852)	136
5.2.3	POWER GXT550P (# 2845)	137
5.2.4	POWER GXT3000P Graphics Accelerator (# 2825)	138
5.2.5	Graphics Adapter Feature Comparison	140
5.2.6	Graphics Adapters Attachment Cables	142
5.2.7	Multiple Adapter Support	143
5.2.8	Accessing the Second Display	145
5.3	Adapter and Device Configuration on PCI-Based RS/6000 Systems	146
5.3.1	Device Types	146
5.3.2	PCI Adapter Configuration	150
5.4	Summary Chart of Adapters	152
Chapter 6. Diagnostics and System Management Services		155
6.1	Diagnostics Operating Considerations	155
6.2	Stand-Alone Diagnostics	157
6.2.1	Booting the Stand-Alone Diagnostics CD-ROM	157
6.2.2	Running Stand-Alone Diagnostics in Concurrent Mode	158
6.3	On-line Diagnostics	158
6.3.1	Service Mode	159
6.3.2	Concurrent Mode	160
6.3.3	Single-User Mode	161
6.4	Diagnostic Tasks and Service Aids	162
6.5	Understanding the Firmware	166
6.5.1	General RS/6000 Boot Sequence	166

6.5.2	Typical Boot Sequence for RS/6000 43P 7043 Model 150	167
6.5.3	Typical Boot Sequence for RS/6000 43P 7043 Model 260	168
6.6	Power-On Self-Test	170
6.6.1	Processor POST	170
6.6.2	Memory DIMM and L2 Cache POST	170
6.6.3	Graphics Adapter POST	171
6.6.4	Keyboard and Mouse Controller POST	172
6.6.5	SCSI Subsystem POST	172
6.7	System Management Services (SMS)	172
6.7.1	SMS Start-Up	173
6.7.2	SMS Graphical Main Menu	174
6.7.3	SMS ASCII Main Menu	180
6.7.4	SMS Firmware Update	180
6.7.5	SMS Firmware Checkpoints	184
Chapter 7.	Service Processor	185
7.1	Introduction to the Service Processor	185
7.2	Customize and Access the Service Processor Menus	186
7.2.1	How to Customize the Service Processor	186
7.2.2	How to Access Service Processor Menus Locally	187
7.2.3	How to Access Service Processor Menus Remotely	187
7.3	Security Within the Service Processor	187
7.4	Service Processor Operational Phases	188
7.4.1	Pre-Standby Phase	189
7.4.2	Standby Phase	189
7.4.3	Bring-Up Phase	190
7.4.4	Run-time Phase	190
7.5	Main Menu	191
7.5.1	Service Processor Setup Menu	192
7.5.2	System Power Control Menu	195
7.5.3	System Information Menu	197
7.5.4	Language Selection Menu	202
7.5.5	CALL-IN/CALL-OUT Setup Menu	202
7.5.6	Set System Name	209
7.6	General User Menu	209
7.6.1	Power-On System	210
7.6.2	Read VPD Image from Last System Boot	210
7.6.3	Read Progress Indicators from Last System Boot	210
7.6.4	Read Service Processor Error Logs	210
7.6.5	Read System POST Errors	210
7.6.6	View System Environmental Conditions	210
7.7	Service Processor Firmware Update	210
7.7.1	Updating Firmware from the Service Processor Menus	211

7.7.2	Updating Firmware from the SMS Utilities	212
7.7.3	Updating Firmware from the Service Aids	212
7.7.4	Updating Firmware from AIX	212
7.8	Service Processor Checkpoints	213
Chapter 8.	Hints and Tips	215
8.1	Electronic Key-Switch Function	215
8.1.1	<code>bootlist</code> Command	215
8.2	Entering Debug Mode	216
8.3	Backup and Cloning	216
8.3.1	The <code>bootinfo</code> Command	216
8.3.2	Creating a System Backup	217
8.3.3	Restoring Your System Backup	217
8.3.4	Cloning	218
8.3.5	Cloning on AIX Version 4.2 and 4.3	223
8.4	Configuration Information	227
8.4.1	Information about the Processor	234
8.4.2	Information about the Service Processor	235
8.4.3	Information about the Environmental Sensors	235
8.5	Network Boot Support	236
8.6	Using the Error Logging Facility	237
8.7	Booting Problems	238
8.7.1	Booting from CD-ROM	238
8.7.2	System Hangs during AIX Boot Process	238
8.7.3	Recovering a System with No Bootable Media	239
Appendix A.	RS/6000 43P 7043 Model 140	241
A.1	RS/6000 43P 7043 Model 140 Key Attributes	241
A.2	RS/6000 43P 7043 Model 140 Feature Summary	242
A.3	Standard Features	243
A.4	Supported Optional Features	244
A.4.1	Storage Devices	244
A.4.2	AIX Operating System	244
A.4.3	Year 2000	245
Appendix B.	Special Notices	247
Appendix C.	Related Publications	251
C.1	International Technical Support Organization Publications	251
C.2	Redbooks on CD-ROMs	251
C.3	Other Publications	251
C.4	Web Resources	252

How to Get ITSO Redbooks	253
How IBM Employees Can Get ITSO Redbooks.....	253
How Customers Can Get ITSO Redbooks.....	254
IBM Redbook Order Form	255
List of Abbreviations	257
Index	259
ITSO Redbook Evaluation	267

Figures

1. RS/6000 PCI Workstation/Server Family	2
2. RS/6000 43P 7043 Model 150 (without Peripherals)	4
3. RS/6000 43P 7043 Model 260 (without Peripherals)	7
4. IEEE 1394 Connections	15
5. RS/6000 Hardware for Workgroup Conferencing	16
6. RS/6000 Models 150 and 260 are Year 2000 Ready	19
7. General Platform Topology	23
8. Planar Block Diagram of the RS/6000 Model 150	27
9. RS/6000 High Level System Block Diagram (Model 260)	30
10. POWER3 Microprocessor Logical Block Diagram	36
11. PowerPC 604e Microprocessor Logical Block Diagram	38
12. Memory Hierarchy	40
13. The 4-Way Set Associative PowerPC/POWER2 L1 Data Cache	41
14. The 128-Way Set Associative POWER3 L1 Data Cache	42
15. Loading Instructions from Memory to a Floating-Point Register	43
16. SMP Cache Coherency Problem	44
17. False Sharing	46
18. RS/6000 43P 7043 Model 150 (with Peripherals)	49
19. RS/6000 43P 7043 Model 150 - System Board Locations	51
20. RS/6000 43P 7043 Model 150 - Front View	53
21. RS/6000 43P 7043 Model 150 - Rear View	54
22. RS/6000 43P 7043 Model 260 (with Peripherals, and Spaceball)	68
23. RS/6000 43P 7043 Model 260 - System Planar & I/O Board	71
24. RS/6000 43P 7043 Model 260 - Operator Panel	73
25. RS/6000 43P 7043 Model 260 - Front View	75
26. RS/6000 43P 7043 Model 260 - Rear View	77
27. Slot Configuration on Memory Card	84
28. Model 150 with Multimedia Kit for RS/6000	95
29. SCSI III Standards Overview	100
30. SCSI Differential Example	104
31. Piggy Back Connectors	107
32. Differential External Narrow Bus	116
33. Differential External Wide Bus	117
34. Comparison between SCSI and SSA Subsystem	119
35. Devices Location Codes	149
36. System Management Services - Main Menu	174
37. SMS Submenu Multiboot	176
38. SMS Submenu Utilities	176
39. SMS Submenu - Firmware Passwords	178
40. Service Processor - Operational Phase Status Flow	188

41. RS/6000 43P 7043 Model 140..... 242

x RS/6000 43P 7043 Models 150 and 260 Handbook

Tables

1. SPEC and LINPACK Performance	13
2. Multiuser Performance	13
3. RPA Minimum System Requirements	25
4. PowerPC and Bus Specification	33
5. RS/6000 43P 7043 Model 150 - System Board Components Location	52
6. Publications Shipped with the Model 150	55
7. Optional Memory for Model 150	59
8. List of Orderable Monitors for Model 150	61
9. Adapter Placement Guidelines for Model 150	66
10. RS/6000 43P 7043 Model 260 - System Planar & I/O Board Components	71
11. RS/6000 43P 7043 Model 260 - Operator Panel Components	73
12. Storage Device Features	74
13. RS/6000 43P 7043 Model 260 - Rear View Components	76
14. Publications Shipped with the Model 260	78
15. Available Publications for Model 260	78
16. List of Orderable Monitors for Model 260	86
17. RS/6000 43P 7043 Model 260 - Adapter Placement Guidelines	93
18. Overview of SCSI-III Standards	101
19. Differential SCSI Modes	105
20. Single-Ended Ultra SCSI Adapter-to-First Device Cables	111
21. Device-to-Device Cables for Single-Ended Installations	112
22. 16-Bit SCSI-II System-To-System Cable	113
23. Terminators for Single-Ended Installations	113
24. Differential-Ended Ultra SCSI Adapter-to-First Device Cables	114
25. Device-to-Device Cables for Differential-Ended Installations	114
26. Terminator For Description-Ended Installations	116
27. Cabling for the PCI Differential Ultra SCSI Adapter	117
28. Comparing SSA and SCSI	120
29. Internal Disk Drives for the Models 150 and 260	121
30. Internal CD-ROM Drives for Models 150 and 260	122
31. Internal Tape Drive for Model 150	122
32. Internal Tape Drives for Model 260	123
33. Placement of PCI Storage Adapters into the Models 150 and 260	126
34. Graphics Adapters	140
35. Display Cable Feature Code Matrix	142
36. Accessing a Second Display	145
37. Communication, Graphics, and Storage Adapters Summary	152
38. Settings of Privileged and General Access Password	188
39. Valid Serial Port Speeds	206
40. Results of bootinfo -p	217

41. RS/6000 43P 7043 Model 140 Standard Configuration.....	243
42. RS/6000 43P 7043 Model 140 System Expansion	243

Preface

In an effort to provide customers with suitable products and solutions to help their business grow, IBM announced in October of 1998 two new additions to its RS/6000 workstation family. They are the RS/6000 43P 7043 Models 150 and 260. These products are an enhancement to the current line of entry workstations that are ready for your mission-critical business applications.

This redbook gives a detailed understanding of the RS/6000 43P 7043 Models 150 and 260, providing all the technical information that is generally requested by decision-makers during the pre-sale cycle. It is suitable as a single-source of reference for both technical and non-technical professionals and may therefore be useful to:

- Customers
- System administrators
- System engineers
- Customer engineers

The reader will find:

- An overview of the RS/6000 43P 7043 Models 150 and 260
- Discussion of hardware architecture
- Detailed product description
- Information on storage and communication features
- Discussion on diagnostics and the service processor

Availability, reliability, and performance features are also described.

The introduction of the RS/6000 43P 7043 Models 150 and 260 will expand the range of possible applications available on entry level workstations, providing industry-leading performance at an affordable price.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization Austin Center.

Volker Haug is an Advisory I/T Specialist in Stuttgart, Germany. He has eight years of experience in the RS/6000 and AIX field. He holds a degree in Business Management from the Berufsakademie in Stuttgart. Volker has

worked with UNIX systems for ten years. His areas of expertise include RS/6000 workstations and workgroup servers, graphics, and AIX systems management.

Adnan Ikram is a Product Specialist (AIX, RS/6000, Networking) in Karachi, Pakistan. He has 3 and 1/2 years of experience in the RS/6000 and AIX field. He holds a degree in Electronics (B.E.) from NED University of Engineering and Technology, Karachi. Adnan has worked with UNIX systems for eight years. His areas of expertise include technical marketing and support of the RS/6000 product family, along with second level AIX support. He is also involved in conducting AIX courses from Basic to Advanced levels at IBM's Education Center.

Heather Osbourne is an I/T Availability Specialist in Kingston, Jamaica. She has six years of experience in the RS/6000 and AIX field. She holds a degree in Computer Science and Physics (B.Sc.) from the University of the West Indies. Her areas of expertise include AIX support, RS/6000 project services implementation, networking and communications. She is also involved in conducting AIX courses from Basic to Advanced levels at IBM's Education Center.

The project that produced this publication was coordinated by:

Scott Vetter International Technical Support Organization, Austin Center

We would also like to acknowledge the professionals who took the time to provide invaluable advice and guidance during this project:

Aranda, Ricardo	IBM Austin
Arroyo, Ron	IBM Austin
Baldwin, Wayne	IBM Austin
Bluethman, Robert	IBM Austin
Bringol, Ron	IBM Austin
Capps, Louis	IBM Austin
Gschell, Verena	IBM Munich
Goodman, Marvin	IBM Austin
Henson, Tracey	IBM Austin
James, Joey V.	IBM Austin
Lehmann, Norma	IBM Austin
Maule, Warren	IBM Austin

Nguyen, Thoi	IBM Austin
Patel, Raj	IBM Austin
Peterson, Earl	IBM Austin
Randall, Dave	IBM Austin
Thurber, Steve	IBM Austin
Walton, Scott	IBM Austin

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 267 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

For Internet users <http://www.redbooks.ibm.com>

For IBM Intranet users <http://w3.itso.ibm.com>

- Send us a note at the following address:

redbook@us.ibm.com

Chapter 1. RS/6000 Introduction

Throughout 1997 IBM greatly simplified its RS/6000 product line by eliminating redundant offerings and introducing new models that cover a broad range of price and performance. In 1998, IBM continued this mission by enhancing its server line with better performance and by offering bundled solutions that target high availability markets and customers requiring outstanding graphics performance.

The RS/6000 family is a scalable, compatible line of RISC UNIX workstations, servers, and supercomputers that are reshaping the way companies operate. This economical family of systems fulfills the computing requirements of many small businesses, workgroups, and large enterprises. Powered by IBM's award winning AIX operating system, the RS/6000 has the advanced technology and architecture needed to grow and adapt to your customer's ever changing workloads.

This publication covers the IBM RS/6000 43P 7043 Models 150 and 260 entry-level and technical workstations. The RS/6000 Models 150 and 260 are workstations that are intended to enhance the RS/6000 family's competitive product line for the price-conscious UNIX market. The RS/6000 Models 150 and 260 are the systems targeted toward those customers looking for a reliable, highly expandable system.

1.1 RS/6000 Workstation Offerings

Figure 1 on page 2 shows the current range of IBM RS/6000 PCI-based workstation and server offerings. The workstations that are the topic of this publication are shown in the shaded area.

The selected workstations share the following characteristics:

- Powered by the PowerPC 604e and POWER3
- Contain PCI adapter slots
- Offer more performance than the Model 140

RS/6000 PCI Workstation/Server Family

Reliability, Affordability, Growth Options

43P Model 260

- 200 MHz POWER3
- 1-2 way SMP

7017-S7A

262 MHz RS64 II
4-16-Way SMP

43P Model 150

- 375 MHz PowerPC 604e
- Powerful and expandable

7025-F50

166/332 MHz
PowerPC 604e
1-4-Way SMP

7026-H50

332 MHz
PowerPC 604e
1-4-Way SMP

7043-140

200/233/332 MHz
PowerPC 604e
Uni-Processor

7024-E30

233 MHz
PowerPC 604e
Uni-Processor

7025-F40

233 MHz
PowerPC 604e
1-2-Way SMP

Figure 1. RS/6000 PCI Workstation/Server Family

The RS/6000 naming convention is patterned similarly to many of IBM's product lines. A product is represented as a four digit machine code followed by a model number. A unique machine code is assigned to each series of machines that fit together using similar designs or customer solutions. The first character of the model number groups machines with a specific enclosure together, and the remaining digits are typically used to indicate performance within similar models. This convention is not without variations. However, with this in mind, you can determine that the Models 140 and 150 are related, with the 150 being more powerful. This is also valid for the Models 240 and 260.

The following sections provide an overview of the models that are the subject of this publication.

1.2 RS/6000 43P 7043 Model 140 Overview

The main purpose of summarizing the IBM RS/6000 43P 7043 Model 140 in this publication is to provide you with some background of this system. The Model 140 is one of the most successful RS/6000 products that IBM has announced in the workstation marketplace. Over 10,000 applications are supported on this platform.

On October 6, 1997, IBM announced the PowerPC 604e 332 MHz microprocessor option for the RS/6000 43P 7043 Model 140 as a new member of the RS/6000 workstation and server family. The PowerPC 604e 332 MHz microprocessor is an easy upgrade from the current microprocessor available with 200 MHz and 233 MHz processor options. The PowerPC 604e 332 MHz microprocessor option shows IBM's commitment to improve the value of its products and to protect existing investment.

More detailed information about the Model 140 can be found in Appendix A, "RS/6000 43P 7043 Model 140" on page 241.

1.3 RS/6000 43P 7043 Model 150 Overview

The IBM RS/6000 43P 7043 Model 150 is an entry-level desktop RS/6000 workstation or workgroup server offered at an affordable price. The Model 150 provides a continuation of the successful line of entry workstations, offering enhanced performance over the Model 140.

The Model 150 is a uni-processor system that provides enhanced performance by utilizing a 375 MHz PowerPC 604e processor and an enhanced memory controller. With this memory controller, the Model 150 uses SDRAM memory and an 83 MHz memory bus speed. The system memory can be expanded up to 1 GB.

With Ethernet and Ultra SCSI controllers integrated on the planar, five PCI slots and bays, the Model 150 is ready for expansion and growth.

The Model 150 is shown in Figure 2 on page 4.



Figure 2. RS/6000 43P 7043 Model 150 (without Peripherals)

The Model 150 supports a variety of 2D and 3D graphics adapters including the advanced 3D GXT3000P graphics adapter, providing excellent graphics price performance. In addition, a robust set of disk drive and communications features are available. The AIX Version 4.2.1 or 4.3.2 operating system for one- to two-users is included in the cost of the Model 150 and can be pre-installed, if desired.

1.3.1 RS/6000 43P 7043 Model 150 Key Attributes

The following is a list of the key attributes of the RS/6000 Model 150:

- Excellent graphics workstation that is ideal for running MCAD and other technical applications.
- Powerful server that is ideal for running small-business and departmental applications.
- Small footprint to fit into your department or work area.

- Ultra SCSI and SSA storage options for improved data transfer performance and high availability.
- Exceptional AIX operating system with reliability, availability, and system-management features.

The Model 150 is designed to help customers run many mission-critical business applications and networks twenty-four hours a day, seven days a week.

1.3.2 RS/6000 43P 7043 Model 150 Feature Summary

The following is a summary of the key RS/6000 Model 150 features:

- Standard configuration:
 - Microprocessor: 375 MHz PowerPC 604e processor
 - Level 2 (L2) cache: 1 MB
 - RAM (memory): 128 MB of ECC SDRAM memory
 - Memory slots: Four
 - Ports: One parallel and two serial
 - Internal disk drive: 4.5 GB Ultra SCSI disk drive
 - Disk/media bays: Five
 - Expansion slots: 5 PCI slots
 - PCI bus speed: 33 MHz
- Standard features:
 - 32X max speed CD-ROM
 - 1.44 MB 3.5-inch diskette drive
 - Integrated 10/100 Mbps Ethernet controller (IEEE 802.3 compliant)
 - Integrated Ultra SCSI controller
 - Tablet port
- AIX operating system:
 - Version 4.2.1 or Version 4.3.2 (1-2 user server license is standard)

- System expansion:
 - RAM: Up to 1 GB
 - Internal storage: 27.3 GB max
 - External disk storage:
 - Maximum external SCSI with expansion units: Up to 254.8 GB
 - Maximum external SSA with expansion units: Up to 873.6 GB
- System dimensions:
 - 6.5" H x 16.5" W x 18.1" D (165 mm x 420 mm x 460 mm)
 - Weight: 14.5 kg (32 lbs)-Base configuration; 18.0 kg (40 lbs)-Full featured configuration
- Operating environment
 - Temperature: 16 to 32 degrees C (60 to 90 F)
 - Relative Humidity: 8 to 80 (percent)
 - Operating Voltage: 100 to 125 V AC or 200 to 245 V AC 50/60 Hz (manual switch)
- Warranty:
 - On-site for one year

1.4 RS/6000 43P 7043 Model 260 Overview

The RS/6000 43P 7043 Model 260 is a one-way or two-way symmetric multiprocessing (SMP) system with the power to run complex commercial, business and engineering applications.

The Model 260 provides enhanced performance over its predecessor, the Model 240, by utilizing a 200 MHz POWER3 processor and an enhanced memory controller. With this memory controller, the Model 260 uses ECC SDRAM memory and a 100 MHz memory bus speed. The system memory can be expanded up to 4 GB.

With Ethernet and Ultra SCSI controllers integrated on the planar, the Model 260 also contains five PCI slots and bays for expansion and growth capability.

Designed with reliability, availability, and serviceability (RAS) features typically found in higher-priced systems, such as Error Checking and Correcting (ECC) memory and an integrated service processor, the Model 260 can help keep your mission-critical business, technical applications, and networks running twenty-four hours a day, seven days a week.



Figure 3. RS/6000 43P 7043 Model 260 (without Peripherals)

The Model 260 supports a variety of 2D and 3D graphics adapters including the advanced 3D GXT3000P graphics adapter, offering excellent graphics price and performance. In addition, a robust set of disk drive and communications features are available. The AIX Version 4.2.1 or 4.3.2 operating system for one- to two-users is included in the cost of the Model 260 and can be available pre-installed, if desired.

The Model 260 is an ideal product for customers with a need for either an entry or mid-range 2D or 3D graphics workstations or as an entry workgroup server.

1.4.1 RS/6000 43P 7043 Model 260 Key Attributes

The following is a list of the key attributes of the Model 260.

- Powerful one- to two-way symmetric multiprocessor (SMP) desk-side server with superior commercial and technical application price/performance.

- Outstanding reliability with ECC SDRAM memory and ECC Level 2 (L2) cache, plus a standard service processor.
- Ultra SCSI and SSA storage options for improved data transfer performance and high availability.
- Enhanced disk bays for improved system availability.
- Exceptional AIX operating system with reliability, availability, and system-management features.

The Model 260 is specially designed to help customers run many mission-critical business applications and networks twenty-four hours a day, seven days a week.

1.4.2 RS/6000 43P 7043 Model 260 Feature Summary

The following is a summary of the key Model 260 features:

- Standard configuration:
 - Microprocessor: 200 MHz POWER3 processor
 - Level 1 (L1) cache: 64 KB data / 32 KB instruction
 - Level 2 (L2) cache on each processor: 4 MB
 - RAM (memory): 260 MB of ECC Synchronous DRAM memory
 - Memory slots: Two (each memory card can contain up to 16 DIMMs)
 - Ports: One parallel and two serial
 - Internal disk drive: 4.5 GB Ultra SCSI disk drive
 - Disk/media bays: Five
 - Expansion slots: Five PCI slots (Two 64-bit PCI slots and three 32-bit slots)
 - PCI bus speed: 33 MHz for the 32-bit and 50 MHz for the 64-bit slots
- Standard features:
 - 32X max speed CD-ROM
 - 1.44 MB 3.5-inch diskette drive
 - Integrated 10/100 Mbps Ethernet controller (IEEE 802.3 compliant)
 - Integrated Ultra SCSI controller
 - Tablet port
- AIX operating system:
 - Version 4.2.1 or Version 4.3.2 (1-2 user server license is standard)

- System expansion:
 - RAM: Up to 4 GB
 - Internal storage: 27.3 GB max
 - External disk storage:
 - Maximum external SCSI with expansion units: Up to 254.8 GB
 - Maximum external SSA with expansion units: Up to 873.6 GB
- System dimensions:
 - 24.2" H x 13.4" W x 26.8" D (610 mm x 340 mm x 713 mm)
 - Weight: 36.2 kg (80 lbs.) - minimum configuration; 43.9 kg (97 lbs.) - maximum configuration
- Operating environment
 - Temperature: 16 to 32 degrees C (60 to 90 F)
 - Relative Humidity: 8 to 80 percent
 - Operating Voltage: 100 to 127 V AC or 200 to 240 V AC 50/60 Hz (autoranging)
- Warranty:
 - On-site for one year

1.5 Performance Results

This section contains the results of several industry standard performance measurements for the RS/6000 Models 150 and 260.

The performance benchmarks and the values shown here were derived using particular, well configured development-level computer systems. Unless otherwise indicated for a system, the values were derived using 32-bit applications and external cache, if external cache is supported on the system. All performance benchmark values are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM. Actual system performance may vary and is dependent upon many factors including system hardware configuration, and software design and configuration. Buyers should consult other sources of information to evaluate the performance of systems they are considering buying and should consider conducting application-oriented testing. For additional information about the performance benchmarks, values and systems tested, please contact your IBM local Branch Office or IBM Authorized Reseller or access the following on the Web:

SPEC <http://www.specbench.org>

Linpack <http://www.netlib.no/netlib/benchmark/performance.ps>

Unless otherwise indicated for a system, the performance benchmarks were conducted using AIX Version 4.2 or Version 4.3. IBM C for AIX Version 4.1.0.1 and XL Fortran Version 5.1.1.1 were the compilers used in the benchmark tests. The preprocessors used in the benchmark tests include KAP 3.2 for Fortran and KAP/C 1.4.2 from Kuck & Associates and VAST-2 Version 4.01X8 from Pacific-Sierra Research. The preprocessors were purchased separately from these vendors.

1.5.1 SPEC95 Software Benchmark

SPEC95 is the forward step in the performance measurement of the core of the system. It covers the CPU, caches, memory, and compiler. The programs and data sets that make up the suite cannot load entirely into cache, making the benchmark more representative of real workloads. SPEC has also standardized the compiler settings so that the results for *base* measurements are more comparable between suppliers.

SPEC95 is a software benchmark produced by the Standard Performance Evaluation Corp. (SPEC), a non-profit group of computer vendors, systems integrators, universities, research organizations, publishers, and consultants throughout the world. It was designed to provide measures of performance for comparing computational-intensive workloads on different computers systems.

SPEC95 contains two suites of benchmarks:

CINT95 Measures and compares computational-intensive integer performance

CFP95 Measures and compares computational-intensive floating point performance

The two groups of programs are referred to as component-level benchmark suites because they test the core of the system, CPU, caches, memory, and compiler, but not the I/O sub-system.

One of the goals of SPEC95 is increased portability; the current offering from SPEC is for UNIX only, although the member companies have indicated that the benchmark programs are portable to various flavors of UNIX, Windows NT, and Open VMS.

SPEC95 introduces a new reference platform against which other systems are measured, changing from the out-dated VAX 11/780 to a SPARCstation 10/40 with 64 MB memory but without Level 2 cache. This is more representative of the types of systems being sold today, but it is also a machine that will beat few, if any, of the machines being benchmarked.

The rules have also changed. Each benchmark must be run a minimum of three times to get a valid result, with the median time for all runs being used as the benchmark time.

The SPEC base metric (for example, SPECint_base95) are required for all reported results and have set guidelines for compilation (for example, the same four flags must be used in the same order for all benchmarks). The non-base metrics (for example, SPECint95) are optional and have less restrictive requirements (for example, different compiler options may be used on each benchmark).

There are several different ways to measure computer performance. One way is to measure how fast the computer completes a single task. This is a speed measure. Another way is to measure how many tasks a computer can accomplish in a certain amount of time. This is called a throughput, capacity, or rate measure. The SPEC speed metrics (for example, SPECint95) are used for comparing the ability of a computer to complete single tasks. The SPEC rate metrics (for example, SPECint_rate95) measure the throughput or rate of a machine carrying out a number of tasks.

The following SPEC and Linpack benchmarks reflect the performance of the microprocessor, memory architecture, and compiler of the tested system.

SPECint95	SPEC component-level benchmark that measures integer performance. Result is the geometric mean of eight tests that comprise the CINT95 benchmark suite. All of these are written in C language.
SPECint_base95	The result of the same tests in CINT95 with a maximum of four compiler flags that must be used in all eight tests.
SPECint_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95).
SPECint_base_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95) with the restrictive compiler options.

SPECfp95	SPEC component-level benchmark that measures floating point performance. Result is the geometric mean of ten tests that comprise the CFP95 benchmark suite. All of these are written in FORTRAN.
SPECfp_base95	Result of the same tests in CFP95 with a maximum of four compiler flags that must be used in all ten tests.
SPECfp_rate95	Geometric average of the ten SPEC rates from SPEC floating point tests (CFP95).
SPECfp_base_rate95	Geometric average of the ten SPEC rates from the SPEC floating-point tests (CFP95) with the restrictive compiler options.
SPECweb96	Maximum number of Hypertext Transfer Protocol (HTTP) operations per second achieved on the SPECweb96 benchmark without significant degradation of response time. The Web server software is ZEUS v1.1 from Zeus Technology Ltd.
LINPACK SP	Single precision, n=100 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS).
LINPACK DP	Double precision, n=100 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS).
LINPACK TPP	Toward Peak Performance, n=1000 results with AIX XL FORTRAN compiler with optimization. Units are megaflops (MFLOPS). ESSL Version 3.1.1 was used in this test.

Relative on-line transaction processing (ROLTP) is an estimate of commercial processing performance derived from an IBM analytical model. The model simulates some of the system's operations such as CPU, cache, and memory. However, the model does not simulate disk or network I/O operations. Although general database and operating systems parameters are used, the model does not reflect specific databases or AIX version or releases. Unless otherwise indicated for a system, the model assumes the use of 32-bit applications. ROLTP is estimated only at the time the system is introduced, unless otherwise indicated for a system. An IBM RS/6000 Model 250 is the baseline reference system and has a value of 1.0.

Although ROLTP may be used to compare estimated RS/6000 commercial processing performance, actual system performance may vary and is dependent upon many factors including system hardware configuration, and software design and configuration. All performance estimates are provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

1.5.2 Performance Summary

Table 1 and Table 2 provide, at time of publication, the related performance values.

Table 1. SPEC and LINPACK Performance.

Model	CPU	Clock Rate (MHz)	L1 (KB)	L2 (MB)	SPEC int95	SPEC int_base95	SPEC fp95	SPEC fp_base95	LINPACK DP	LINPACK SP	LINPACK TPP
150	604e	375	32/32	1	15.1	14.5	10.1	9.85	64.8	151.2	255.7
260	POWER3	200	32/64	4	13.2	12.5	30.1	27.6	236.5	248.1	630.0

Table 2. Multiuser Performance

Model	CPU	Clock Rate (MHz)	L1 Cache (KB)	L2 Cache (MB)	Relative OLTP Perf.	SPEC int_rate95	SPEC int_base_rate95	SPEC fp_rate95	SPEC fp_base_rate95
150	604e	375	32/32	1	6.0	-	-	-	-
260	POWER3 (1-way)	200	32/64	4	10.5	116	111	266	243
260	POWER3 (2-way)	200	32/64	4	21.0	232	222	509	468

1.6 Industry Trends and Directions

There are many emerging trends and directions in the UNIX hardware marketplace. Some will be determined as frivolous, and others will become industry standards. In this section there is a discussion of some key emerging technologies that may, in the opinion of the authors, become pervasive in the marketplace.

1.6.1 Microprocessor Future

As applications grow more complex, more processor power is required to give you acceptable performance. It is easy to predict the following developments pertaining to the growth of computing power available for the RS/6000.

- Moore's law will remain true. 1000 MHz processors are already running in the lab. More importantly, IBM's use of copper technology and complementary metal oxide semiconductor (CMOS) fabrication

techniques will continue to bring significant enhancements to the PowerPC.

- The POWER3 architecture provides a significant impact to the market by:
 - Providing SMP scalability not found with POWER2
 - Redefining the versatility of a departmental computing solution by providing outstanding float-point performance with solid integer performance
 - Providing an affordable entry-level 32-bit and 64-bit solution

1.6.2 Universal Serial Bus

The Universal Serial Bus (USB) is a plug and play peripheral connection for devices such as keyboards, mice, joysticks, scanners, tablets, printers, and digital cameras, to name a few. USB could replace the serial and parallel ports with a single bus.

USB supports many peripherals at one time. Special USB hubs will allow devices to be daisy-chained together. USB also distributes power to the attached devices, eliminating the need for dedicated peripheral power supplies.

Data flow in a USB is bidirectional. Devices can make use of this two-way communication to report status, or control other peripherals.

1.6.3 The Accelerated Graphics Port

The Accelerated Graphics Port (AGP) is a high-bandwidth 32-bit PC bus architecture introduced in 1997 by Intel. It provides up to 528 MB/s memory access to a graphics controller, yielding the bandwidth necessary for complex tasks such as texturing directly from system memory. AGP is physically different from the PCI bus and does not replace PCI. AGP uses a combination of frame-buffer memory local to the graphics controller, as well as system memory for graphics data storage.

Although the AGP interface is very popular in the PC world, it has not become a standard. However, the importance of images, video, and graphics is growing. The next step in this evolutionary path is visual computing.

Visual computing is the convergence of high-performance 3D graphics, video, and digital imaging technologies to deliver a new class of interactive, intuitive, and life-like computing experiences to users.

1.6.4 IEEE 1394

IEEE 1394 is a transport protocol standard for a high performance serial bus (also known as FireWire) that is a bus technology for the digital age. More than 50 companies are supporting this standard. FireWire was originally developed by Apple Computer. It was started as a specification for the transmission of digitized video and audio quickly, reliably, and at low cost. The architecture is scalable and hot-pluggable. Industry consortiums are moving toward a 1 gigabit data transfer rate.

Like USB, 1394 enables plug-and-play peripheral connectivity and supports isochronous data transfers. The major difference is data rate. USB ideally supports slow-speed devices like keyboards. IEEE 1394 is able to handle all I/O, including printers, DASD, PCI, and IDE devices, as shown in Figure 4.




Figure 4. IEEE 1394 Connections

1.7 Workgroup Conferencing

Workgroup Conferencing Version 1.1.0 for AIX (commonly referred to as Workgroup) is IBM's first interoperable and industry standards-based desktop conferencing product that unleashes the real-time videoconferencing

capabilities of IP-based intranet or Internet network environments on RS/6000 and AIX systems. This full-function AIX videoconferencing application features an easy-to-use user interface to setup and manage spontaneous videoconferencing sessions between RS/6000 and other desktop conferencing products, such as Microsoft NetMeeting, on Windows and other UNIX systems.



Figure 5. RS/6000 Hardware for Workgroup Conferencing

Customer benefits of this application include:

- Flexibility
Workgroup provides full-function desktop videoconferencing for a variety of different application scenarios.
- Features
Workgroup was developed using a feature-rich implementation of the dominant International Telecommunications Union (ITU) Recommendation H.323 Version 2 industry standards for Packet-Based Multimedia Communications Systems.
- Interoperability

Workgroup was designed to interoperate with AIX, other UNIX, and Windows H.323-based conferencing systems including other desktop videoconferencing systems such as Microsoft NetMeeting Version 2.1, H.323/PSTN Gateways and Gatekeepers.

- Scalability

Workgroup supports point-to-point and multiparty calls in a peer-to-peer mode. WorkGroup was also designed to participate in large group, server-based conferencing sessions.

- Ease of use

Workgroup features simple point-and-click dialing access to other desktop conferencing users with personalized speed-dial or popular directory server applications.

Workgroup breaks down the barriers of distance:

- Workgroup delivers full-function desktop videoconferencing. By simply configuring an RS/6000 system with a video capture card, camera, and connection to an IP network, Workgroup enables AIX Version 4.3.2 customers to host small group videoconferencing sessions without the expense or complexity of sophisticated videoconferencing equipment or conferencing servers.
- This new IBM application provides easy access to a variety of videoconferencing call setup and management features. Using a high-level user interface, initiating a person-to-person or multi-party Workgroup videoconference call is almost as easy as dialing a telephone number. Workgroup features speed-dialing support for frequently called locations and an automated call history log that contains an ongoing record of all incoming calls. All necessary speaker, microphone, do-not-disturb, and video controls are available in an easy-to-use Call Center workspace. Local and remote video windows are easily detached from the Call Center workspace and can be resized using certain RS/6000 graphics accelerators, from original size up to full screen video window.
- By exploiting a feature-rich implementation of the ITU-T Recommendation H.323 Version 2, Workgroup is designed to provide full-function videoconferencing and interoperates with other H.323 standards-based videoconferencing systems on AIX, other UNIX, and Windows systems. In addition, IBM participates in worldwide H.323 interoperability events to promote Workgroup compatibility with other videoconferencing systems based on this dominant standard for highly interactive, multimedia teleconferencing products and services. These H.323 standards, among others, are also promoted by the International Multimedia

Teleconferencing Consortium, Inc. (IMTC), a non profit corporation composed of more than 150 members and affiliates from North America, Europe, and Asia/Pacific. The IMTC organization is actively involved in promoting open international standards in the development of interoperable, full-function multimedia teleconferencing solutions.

- Workgroup contains support for popular Lightweight Directory Access Protocol (LDAP) and Microsoft Internet Locator Service (ILS) database server applications that reveal other users available for on-line, videoconference sessions within a corporate intranet or public Internet network environment.
- WorkGroup Conferencing for AIX, Version 1.1.0 features a full-function user interface designed for both videoconferencing and the addition of future desktop conferencing features.

Note

For more information about WorkGroup, visit the following Web site:

<http://www.rs6000.ibm.com/solutions/interactive>

Workgroup Conferencing for AIX, Version 1.1.0 is included in the *Bonus Pack for AIX 4.3.2* CD, which is shipped with the base AIX operating system.

1.8 Year 2000

As a global organization, IBM has been supporting customers around the world in all the major phases of Year 2000 transition period. The reach of information systems has continued to expand at a rapidly increasing pace and the breadth of the information that can be distributed is growing at an equally rapid rate. Keeping in mind the Year 2000 challenges, the pressure on the IT organization is increasing. As business must achieve unmatched levels of customer satisfaction and competitive advantage, the IT organization is challenged to continue to deliver advanced uses of information technology.

The IBM RS/6000 43P 7043 Models 150 and 260 do not have date dependencies and are therefore Year 2000 ready. When used in accordance with its associated documentation, it is capable of correctly processing, providing, and/or receiving date data within and between the 20th and 21st centuries, provided all other products (for example, software, hardware and firmware) used with the product properly exchange accurate date data with it.

Further information can be found on IBM's Year 2000 webpage at:

<http://www.ibm.com/year2000>.



Figure 6. RS/6000 Models 150 and 260 are Year 2000 Ready

Chapter 2. Hardware Overview

This chapter introduces the standards and architectures that are shared between the RS/6000 43P Series systems. The discussion encompasses the following topics:

- PowerPC microprocessor RS/6000 Platform Architecture (RPA)
- Hardware design
- Processor subsystem
- SMP hardware characteristics

The general architectures discussed in these sections help build an overall appreciation for RS/6000 hardware.

2.1 The RS/6000 Platform Architecture

The RS/6000 Platform Architecture (RPA) introduces a comprehensive computer system hardware-to-software interface definition, combined with minimum system requirements, that enables the development and porting of software to a range of compatible computer systems from portables through servers. These systems are based on the PowerPC microprocessor, as defined in the *PowerPC Architecture: A Specification for a New Family of RISC Processors*. Both the RS/6000 43P 7043 Model 150 and Model 260 are RPA compliant.

2.1.1 The RS/6000 Platform Architecture Introduction

The RPA is essentially a combination of its predecessor, the PowerPC Microprocessor Common Hardware Reference Platform (shortened by the industry to be known as the PowerPC Platform) and some IBM architecture extensions. This platform architecture officially came into being in August of 1997.

Prior to IBM's RPA, an alliance existed between Apple, IBM, and Motorola that developed a common hardware platform with the purpose of defining a system that would become the pervasive open industry standard, from single user to server configurations. As a result of this alliance, the PowerPC Platform emerged. The PowerPC Platform, previously known as CHRP, is a superset of the earlier PowerPC Platform Specification (PRep) system architecture. The Apple, IBM, and Motorola alliance is no longer functional.

A key benefit of the RPA specification is the ability of hardware platform developers to have degrees of freedom of implementation below the level of

architected interfaces and therefore have the opportunity for adding unique value. This flexibility is achieved through architecture facilities including:

- Device drivers
- Open Firmware (OF)
- Run-Time Abstraction Services (RTAS)
- Hardware abstraction layers

2.1.2 Platform Topology

Figure 7 on page 23 shows a general platform topology of the RS/6000 Platform Architecture. All platforms (from notebooks to high-end servers) consist of one or more PowerPC microprocessors, a volatile system memory separate from other subsystems, and a number of I/O devices, that may initiate transactions to system memory.

The processors are linked over the primary processor bus/switch to each other, to the system memory, and to one or more host bridges (host bridge 0 must be a PCI host bridge).

In general, I/O devices do not connect to the primary processor bus/switch. The host bridges connect to secondary buses that have I/O devices connected to them. In turn, one or more bus bridges may be employed to tertiary buses (for instance ISA or PCI) with additional I/O devices connected to them. Typically, the bus speeds and throughput decrease and the number of supportable loads increases as one progresses from the primary processor bus to more remote buses.

There are variations to this topology that are likely to occur and are therefore worth describing in the following list. The architecture describes interfaces not implementation. The logical software model must remain the same, even if the physical topology is different.

- In a smaller platform, the host bridge, memory, or an I/O device may be integrated into a single chip. In this case, the topology would not look like Figure 7 on page 23, from a chip point of view, but, instead, would be integrated onto the single chip.
- In a larger platform, secondary buses may be implemented, with two or more host bridges, as two or more parallel expansion buses for performance reasons. Similarly, tertiary buses may be two or more parallel expansion buses off each secondary bus. This is indicated by the dots near the host bridge and the bus bridge.

- In a high performance platform, with multiple processors and multiple memories, a switch may be employed to allow multiple parallel accesses by the processors to memory. The path through the switches would be decided by the addressing of memory.

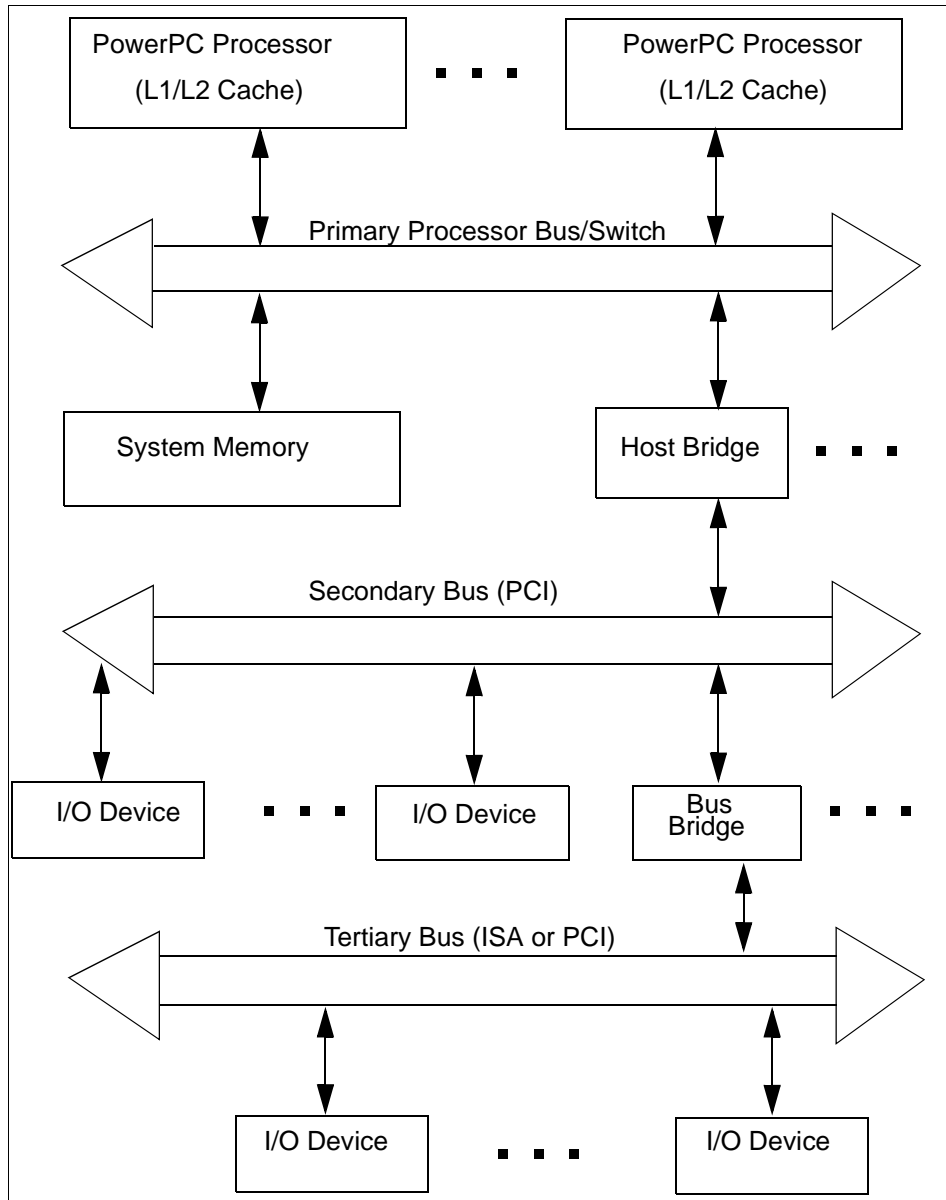


Figure 7. General Platform Topology

2.1.3 RS/6000 Platform Architecture Achievements

The RS/6000 Platform Architecture specification achieves the following:

- Creates a stable platform architecture to be used by all RS/6000 platforms based on PowerPC processors and PCI I/O bus (and may contain other I/O buses in addition to the PCI bus).
- Creates an architecture that allows platforms to operate with a previous version of AIX (n-1 capability).
- Leverages existing and future industry-standard buses and interfaces. Existing bus architectures have a proven level of performance and function. Established industry-standard interfaces (for example, SCSI, IDE, LocalTalk, Ethernet) and newer bus architectures, interfaces, and protocols (PCI, PC Card, IrDA, to name a few) provide higher levels of performance or utility that are not achievable by the older standards. The architecture allows platform and system designers to determine which buses, interfaces, and protocols best suit their target environment.
- Provides a flexible address map. Another key attribute of this specification is the relocatability of devices and subsystems within the PowerPC address space. Subsystem address information, which defines where I/O devices reside, is detected by the Open Firmware and passed to the operating systems in the device tree. The architecture accommodates the use of multiple identical buses and adapters in the same platform without address conflicts.
- Builds upon the Open Firmware boot environment defined in *IEEE 1275, IEEE Standards for Boot (Initialization Configuration) Firmware, Core Requirements and Practices*, IEEE part number DS02683, ISBN 1-55937-426-8. (To order copies of IEEE standards, load the following Web site:
<http://standards.ieee.org/index.html>).

In the past, the abstraction approach for the operating system used platform description information discovered by a legacy boot process and passed it to the operating system in data structures.

- Architects the control of power management by the operating system. It is important that the combination of hardware and software be allowed to minimize power consumption through automatic or programmed power-saving methods. Power management of systems will reduce the operational cost for the user and reduce the impact of the system on the environment.
- Provides an architecture that can evolve as technology changes.
- Minimizes the support cost for multiple AIX versions through the definition of common platform abstraction techniques. Common and compatible

approaches to the abstraction of hardware will reduce the burden on hardware developers who produce differentiated machines.

- Architects a mechanism for error handling, error reporting, and fault isolation. The architecture provides for the implementation of more robust systems if desired by the system developers.
- Architects a mechanism for dynamic reconfiguration of the hardware.

2.1.4 RPA Minimum System Requirements

Table 3 contains the RPA compliance, minimum system requirements for the workstations addressed in this publication.

Table 3. RPA Minimum System Requirements

SUBSYSTEM	SPECIFICATION	REQUIRED	SUPPORT
Processor	PowerPC microprocessor	R	Yes
Minimum system memory	8 MB expandable to at least 32 MB	R	Yes (128 MB min)
Firmware storage	Sized as needed	R	Yes (1 MB Flash)
Non-volatile memory	8 KB	R	Yes (64 KB)
External cache	Recommended for performance	O	Yes (1 MB)
Hard disk ¹	SCSI	R	Yes
Floppy ¹	3.5" 1.44 MB MFM Media sense Manual eject	O R R R	Yes Yes Not supported Yes
CD ROM	4X speed ISO9660 Multi-session	O R R R	Yes Yes 32X Speed Yes Yes
Alphanumeric input device ¹	PS/2 keyboard interface Terminal	R R O	Yes Yes Yes
Pointing device ¹	2 buttons PS/2 Interface	R R R	Yes Yes (3 button) Yes
<u>Required definition</u> R - Required O - Optional			

SUBSYSTEM	SPECIFICATION	REQUIRED	SUPPORT
Audio	16-bit stereo, 22.05 and 44.1 KHz, full duplex	R	Yes
Tone		R	Yes
Graphics	default 1024x768 Bi-Endian 640x480x8 LFB VGA	R R R R O	Yes Yes Yes Yes Yes
Real-time clock		R	Yes
Serial port ¹	16550	R R	Yes Yes
Parallel port	P1284 + ECP mode	R	Yes
Network		O	Yes (Ethernet 10/100)
Interrupt controller	Open PIC 8259 tree	R R	Yes Yes
Direct memory access (DMA)	ISA	R	Yes
Power management	States as defined	O	NO
Infrared	IrDA	O	NO
<u>Required definition</u> R - Required O - Optional			

¹See the published CHRP document for the original value of this row for platforms that need to support Apple Mac OS.

2.2 The Hardware Design

The designs of the Model 150 and the Model 260 differ slightly and will therefore be addressed in separate sections.

2.2.1 Hardware Design for the RS/6000 43P 7043 Model 150

The RS/6000 Model 150 is a uni-processor system that utilizes the PowerPC 604e processor running at 375 MHz. It supports an integrated 1 MB

parity-checked synchronous L2 cache. The Model 150 also features an enhanced memory controller that uses ECC SDRAM memory.

The local system bus on the RS/6000 Model 150 is a 60X bus running at 83 MHz. This bus is 64 bits wide (with an additional 8 bits for parity) and is synchronous. The processor clock and regulator card reside as one unit, namely the OVdd/Vdd/PLL card. This is illustrated in the planar design of the Model 150 in Figure 8.




Figure 8. Planar Block Diagram of the RS/6000 Model 150

In the Model 150, the 60X bus runs from the PowerPC 604e processor and connects to an enhanced memory controller chip, which has a processor

system interface of a 64-bit data bus and a 32-bit address bus. This highly integrated chip acts as both a PCI bridge and a memory controller. It contains the system bus arbitration, provides support for full memory coherency, and pipelining of processor accesses.

Information may be routed from the memory chip to main memory, using the memory bus, or to I/O devices using the I/O bus.

In the Model 150, the memory bus runs at 83 MHz. The memory subsystem supports up to four industry-standard DIMM sockets with supported capacities of 64, 128 and 256 MB. The memory DIMM can be populated in any order. An ECC controller/buffer, running at 83 MHz, handles high performance ECC operation with SDRAM memory, and parity operation on the system bus. Single bit errors are corrected, while double bit errors are detected.

For access to I/O devices, the PCI bus is used, running at 33 MHz. Devices such as the integrated Ethernet adapter, the integrated Ultra SCSI controller and the IBM Multiprocessor Interrupt Controller (MPIC) are all attached to the PCI bus. This bus also leads to the PCI-to-ISA bridge, which the Model 150 uses as the system I/O bridge. The PCI-to-ISA bridge supports a PCI bus controller interface to enhanced IDE drives, an ISA bus bridge, and an XD-bus interface (for support of flash EPROM).

The ISA bus runs at 8 MHz and provides support for the following:

- Business audio controller
- Tablet port
- Native I/O controller (for diskette controller, serial ports, parallel port)
- 1 MB Flash ROM (to contain the IPLROS Open Firmware)

2.2.2 Hardware Design for the RS/6000 43P 7043 Model 260

Figure 9 on page 30 shows the system block diagram for the Model 260. This workstation is a 1- or 2-way symmetric multiprocessing system using the 64-bit 200 MHz POWER3 processor. The processor subsystem is implemented on pluggable processor cards. If two cards are installed in the system, they must both be the same type or the system will not power on.

The POWER3 processor card has two main components:

- POWER3 chip itself
- 4MB L2 cache

The controller for the L2 cache is integrated on the POWER3 chip, so there is no need for the separate X5 cache that exists on some RS/6000 systems, such as the Model F50.

Three main buses are used in this system design, namely:

- System bus
- Memory bus
- I/O bus

The system bus within the RS/6000 Model 260 is an SMP 6XX high performance bus. The POWER3 processor delivers the SMP server bus directly, and therefore can have direct connectivity to the memory controller chip. Each POWER3 CPU has two paths leading from the processor, namely an address bus (64-bit) and a data bus (128-bit), highlighting the true split transaction capability of the 6XX bus. The address bus of one CPU merges with the address bus of the other to create a single input to the Address Controller unit. Similarly, the data bus of one CPU merges with the data bus of the other, creating a single input to the Data Controller unit.

The Address Controller unit and the Data Controller unit, both semi-custom CMOS5S chips, are separate partitions of a recently developed enhanced memory controller unit. This chipset supports running the bus at 100 MHz. A clock distribution chip was designed and placed on each memory card to achieve 100 MHz performance.




Figure 9. RS/6000 High Level System Block Diagram (Model 260)

The memory controller chip can act as the system interrupt controller, or can pass interrupts through to an external device. Two distinct interrupt protocols

are supported, namely the RISC System/6000, or the industry standard OpenPIC.

Special attention was given to providing support for the memory controller unit in the Model 260:

- Benchmark traces were used to fine-tune the buffer structure and queue depths.
- The traces were also utilized to remove any wasted or unused cycles internal to the memory controller.

The output from the memory controller function may be routed to main memory, using the memory bus, or to I/O devices, using the I/O bus.

To transport addressing information from the address controller to main memory, the memory address bus is used. Data is passed from the data controller to main memory using the memory data bus. This bus supports sustainable memory bursts across a 128-bit data bus, which translates to a read bandwidth of 1.6 GB/s at 100 MHz. The memory bus drives up to four memory cards populated with non-buffered synchronous dynamic random access memory (DRAM) standard dual in-line memory modules (DIMMs). The SDRAM sizes supported are 16 MB, 64 MB, 128 MB, and 256 MB, with 64 MB being the minimum total system memory supported. Two error checking and correction (ECC) codes are supported by this memory controller unit. The first supports single bit correction, double bit detection, and four bit packet detection. The second supports single bit correction, two bit packet correction, and double packet detection.

For access to I/O devices, the 6XX-MX I/O bridge bus is used. This bus is a 64-bit, time-multiplexed address and data bus, and can attach a variety of I/O bridge chips. It runs asynchronously to the system and memory bus, and will support speeds up to 75 MHz, translating to a sustainable read bandwidth of 533 MB/s. From the 6XX-MX bus, information passes through the PCI bridge chip, where it is converted to PCI format for use with PCI devices. Each PCI controller is able to communicate with a limited number of slots, so additional PCI controllers are required to support a large number of PCI cards. Also, specifications for 64-bit slots and standard 32-bit slots will often require specialized PCI controllers.

The ISA bridge located on the I/O backplane allows conversion from PCI to ISA, enabling the use of native ISA devices. The integrated service processor, also located on the I/O backplane, allows for improved reliability, availability, and serviceability features.

2.3 Bus Architecture

The following sections cover the different bus architectures that exist on the RS/6000 43P 7043 Models 150 and 260.

2.3.1 The PCI Bus Architecture

The Peripheral Component Interconnect (PCI) is a specification standard for computer bus implementation developed by the PCI Special Interest Group (PCI-SIG), led by a group of companies including Compaq, Digital, IBM, Intel, and NCR. There are now over 300 companies in the PCI-SIG supporting the architecture and currently producing PCI products.

The goal is to provide a common system-board bus that could be used in personal computers, from laptops to servers. It is envisioned as a local system-board bus that would serve as a common design point, supporting different system processors as the various processors evolve over time. This is much like operating systems that have defined application binary interfaces (ABIs) so that applications need not change with each generation of the operating system. The PCI local bus would serve as a common hardware interface that would not change with different versions of microprocessors.

2.3.1.1 PCI Specification Revisions

According to PCI Specification Revision 2.0, the PCI bus operates on 32- or 64-bits of data at a clock speed of 33 MHz. However, as the speed of processors continues to increase and devices such as video adapters require more bandwidth, bus speeds must also increase to avoid data bottlenecks. Consequently, PCI Specification Revision 2.0 has been superseded by PCI Local Bus Specification Revision 2.1, which introduces support for a 66 MHz bus operation speed, doubling the throughput of the former 33 MHz standard. The RS/6000 43P Model 260 complies with PCI Local Bus Specification Revision 2.1. IBM's implementation of the 64-bit bus slot uses a clock speed of 50 MHz. The Model 150 is PCI 2.0 compliant.

Under PCI Specification Revision 2.1, the 66 MHz bus operation speed is an optional feature and is backward-compatible with 33 MHz PCI devices and buses. If a 66 MHz capable PCI device is installed in a 33 MHz PCI bus, the device must operate at 33 MHz. Likewise, if any 33 MHz PCI devices are installed into a 66 MHz PCI bus, the PCI bus must operate at 33 MHz.

2.3.1.2 PCI Features and Benefits

The PCI bus architecture has many advantages including the following:

- High data transfer speed
- Processor independence
- Cross-platform compatibility
- Multi-bus support
- Plug and Play
- Investment protection

2.3.2 The ISA Bus Architecture

The Industry Standard Architecture (ISA) is the most widely used system bus in the PC industry. Initially, it had no official definitions or standards. However, its specifications have since been defined by the Institute of Electrical and Electronics Engineers (IEEE) standards group.

The ISA bus allows a transfer rate of up to 8.3 MB/s. Transfers over the ISA bus are synchronized around 8 MHz, and they usually take a minimum of two cycles of the bus clock to perform a data transfer. Since the data path of an ISA bus is 16 bits wide, up to 2 bytes may be transferred during each transaction.

On the RS/6000 43P Series Model 150 and Model 260 computers, the ISA bus is used to attach internal devices such as the diskette drive. No ISA slots are provided for the addition of non-native ISA devices.

There is a disadvantage in trying to connect the processor to the ISA bus directly. The processor's speed is reduced to match the slow ISA bus speed. Thus, the systems cannot take advantage of a fast processor.

To avoid this disadvantage, the PCI local bus is used as the primary system bus and the ISA bus as an expansion bus. This way, the system can take advantage of the high-speed data transfer provided by the PCI bus when communicating with the processor and memory. On the other side, through the PCI-ISA bridge, the bus clock can be reduced to match the ISA bus requirements.

Table 4 provides the bus specification for different architectures and compares them to the PowerPC processor's speed.

Table 4. PowerPC and Bus Specification

Specification	PowerPC	PCI	ISA(8)	ISA(16)	EISA
Processor Speed	166, 200 233, 375				

Specification	PowerPC	PCI	ISA(8)	ISA(16)	EISA
Data Bus	64	64	8	16	16/32
Address Bus	32	64	20	24	24/43
Bus Clock	66 MHz 83 MHz 100 MHz	33 MHz (32-bit)	4.7 MHz	8.3 MHz	8.3 MHz
Interrupts		4	6	11	11
DMA Channel		busmaster	3	7	7

2.3.3 The 60X and 6XX Bus

The 60X bus was designed to provide a common interface that supplies communication between processors, memory controllers, and high-speed I/O devices. Arbitration on the bus is performed by a central, off-chip arbitrator. The 60X bus provides the capability to build an SMP system with the MESI protocol. This bus is also known as the MP System Bus.

The 60X bus uses separate address and data buses and a variety of control and status signals. The address bus is 32 bits wide (with 4 bits parity) and the data bus is 64 bits wide (with 8 bits parity). These separate address and data buses run independently, allowing split transactions. The bus is synchronous, with all timing relative to the rising edge of the bus clock. Inputs are sampled, and outputs are driven from this edge.

It is through this bus that the two-way SMP's processors are attached, memory is accessed, and cache snooping is achieved.

2.3.4 The 6XX-MX I/O Bus

The 6XX-MX I/O bus is a high-performance bus designed to be used as a high-end I/O bus for mid- and high-range systems. It will accommodate both 32- and 64-bit processors with real memory addressing extending beyond 32 bits.

The requests and replies are tagged, allowing for out-of-order replies, which are common for programmed I/O to another bus or in a non-uniform memory access (NUMA) environment.

This bus must have a bridge chip to translate a 6XX or 60X transaction with the corresponding operation. Due to the 66 MHz PCI capability as defined in the PCI Local Bus Specification (Revision 2.1), local bus performance has increased from 264 MB/s to 528 MB/s for 64-bit transfers.

2.3.5 The X-Bus

The X-Bus is an 8-bit ISA subset bus used to attach several common subsystems to the PCI-based RS/6000 servers. The following components are attached to this bus:

- The Keyboard/Mouse controller
- Real-Time Clock and Non-volatile RAM

The functions of Real-Time Clock (RTC) and Non-volatile RAM (NVRAM) are integrated into a single component in the PCI-based RS/6000 servers. This component also supplies the logic required to perform the function of powering the system on at a designated time. The component is operated from a lithium battery on the board so that all time-keeping functions continue to work while system power is turned off.

- Mini-Support Processor (I²C Controller)

The Mini-Support Processor is a minicontroller that is imbedded into the I/O planar of the PCI-based RS/6000 servers. It allows the PowerPC processor access to VPD, operator panels and other I²C bus-attached devices.

2.4 The Processor Subsystem

Since its inception, the RS/6000 product line has used a number of different processor designs, the original implementation being the POWER architecture. The most recent development, however, is the POWER3 processor that is implemented in the RS/6000 43P 7043 Model 260. The Model 150 uses the PowerPC 604e chip.

2.4.1 The POWER3 Microprocessor

The POWER3 microprocessor introduces a new generation of 64-bit processors especially designed for high performance and visual computing applications. POWER3 processors will replace the POWER2 and the POWER2 Super Chips (P2SC) in high-end RS/6000 workstations and technical servers. The RS/6000 43P 7043 Model 260 workstation features the POWER3 microprocessor.

2.4.1.1 Processor Overview

The POWER3 implementation of the PowerPC architecture provides significant enhancements compared to the POWER2 architecture. The SMP capable POWER3 design allows for concurrent operation of fixed-point instructions, load/store instructions, branch instructions, and floating-point instructions. Compared to the P2SC, which reaches its design limits at a

clock frequency of 160 MHz, POWER3 is targeting up to 600 MHz by exploiting more advanced chip manufacturing processes, such as copper technology. The first POWER3 based system, RS/6000 43P 7043 Model 260, runs at 200 MHz.

Features of the POWER3, exceeding its predecessor (P2SC), include:

- A second load-store unit
- Improved memory access speed
- Speculative execution

Figure 10 shows the POWER3 microprocessor architecture.

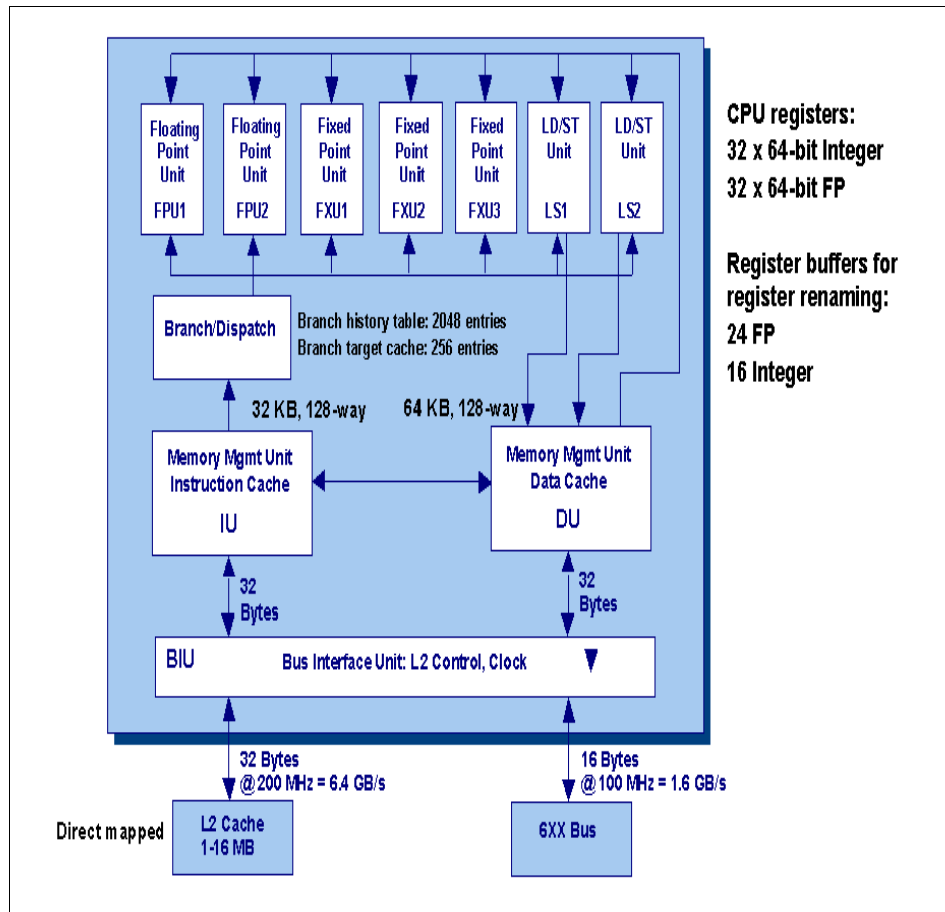


Figure 10. POWER3 Microprocessor Logical Block Diagram

The two additional execution units give POWER3 a peak instruction rate of eight instructions per cycle (two floating-point, two load/store, two single-cycle integer, a multi-cycle integer, and a branch instruction).

Significant investments in the chip's memory interface have been made in order for the POWER3 to have a sustainable execution rate of four instructions per cycle (two load instructions and two floating-point). Although its 64 KB data cache is only half the size of the P2SC's, its advanced core, a dedicated second level (L2) cache, and aggressive prefetching mechanisms improve the memory access speed. The initial POWER3 implementation applies 4 MB (L2) cache.

POWER3 is defined by the following specifications:

- POWER3 microprocessor running at:
 - 200 MHz on RS/6000 Model 260
- Superscalar design with integrated integer, floating-point, and branch units
- 32 KB instruction cache
- 64 KB 128-way set associative data cache
- 64-bit memory interface with 64-bit addressing
- Real memory support for up to 4 GB (3 GB on AIX 4.2.1)

2.4.2 The PowerPC 604e Microprocessor

The RS/6000 43P 7043 Model 150 workstation features the PowerPC 604e microprocessor. The 604e is an equivalent, but enhanced derivative of the 604. Enhancements to the PowerPC 604e, exceeding its predecessor (PowerPC 604), include:

- Doubled L1 cache (instruction and data)
- Higher clock frequencies
- Built-in performance monitor

The superscalar design of the 604e provides up to four instructions to be dispatched per cycle and four to complete at once, with one store and one branch per cycle.

Figure 11 on page 38 shows the PowerPC 604e microprocessor architecture.

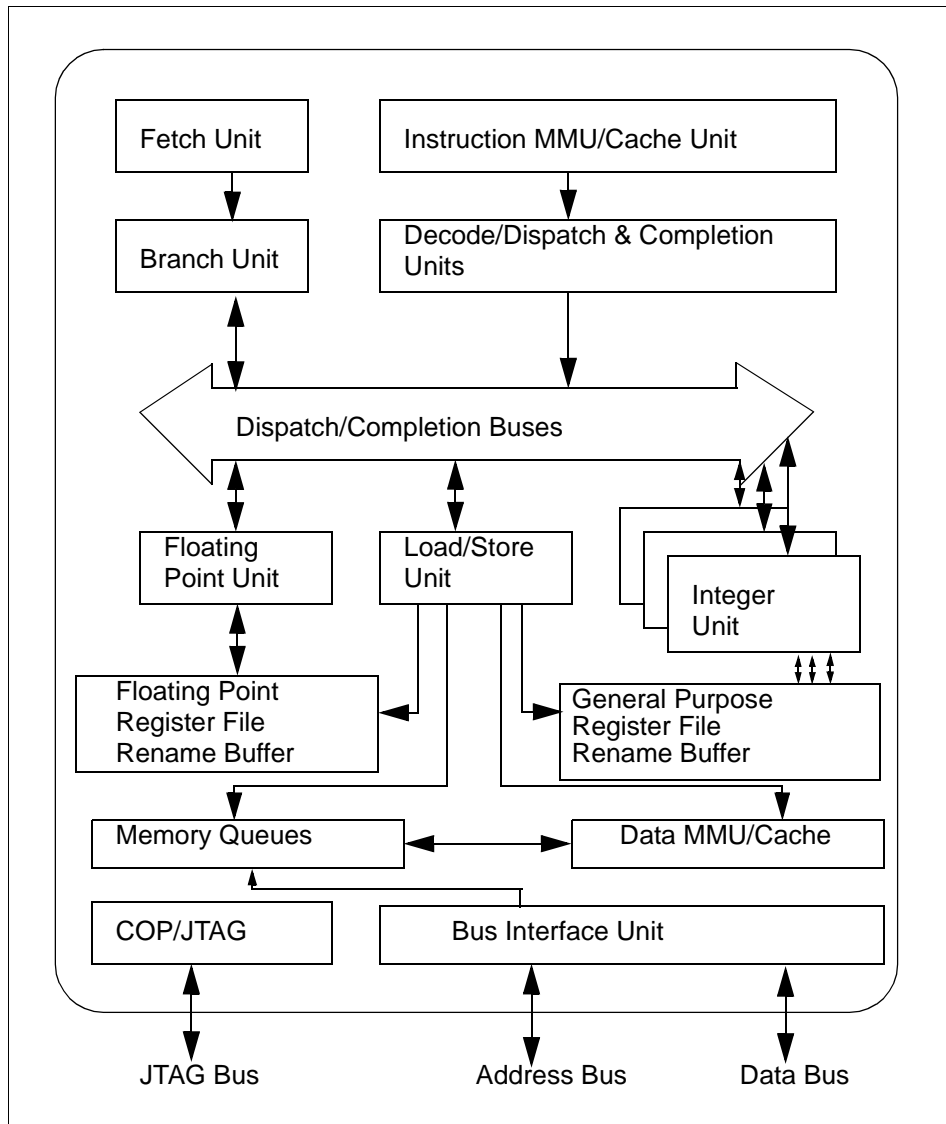


Figure 11. PowerPC 604e Microprocessor Logical Block Diagram

A 604e is defined by the following specifications:

- PowerPC 604e microprocessor running at:
 - 375 MHz on RS/6000 Model 150
- Superscalar design with integrated integer, floating-point, and branch units

- 32 KB four-way set-associative instruction cache
- 32 KB four-way set-associative data cache
- 64-bit memory interface with 32-bit addressing
- Virtual memory support for up to 4 petabytes (2^{52})
- Real memory support for up to 4 gigabytes
- ~5.1 million transistors
- Nap power management mode
- JTAG bus (IEEE 1149.1) for RAS

2.5 SMP Hardware Characteristics

Sharing resources is probably the main technical issue in the design of a symmetric multiprocessor (SMP) system. To support symmetric multiprocessing, specific techniques must be provided at the hardware level and the software level. This section introduces the memory hierarchy concept and some of the techniques used to solve resource-sharing and contention issues.

2.5.1 Memory Hierarchy

In order to improve the hardware performance of a system (uniprocessor (UP) or multiprocessor (MP)), different levels of memory are used. These different levels of memory can be ordered according to their access time and capacity.

If you look at the different types of memory available on a typical system, you will find the CPU registers at one end. They are fast but small, and have a high cost per bit. At the other end, you will find the disks, which are slow but have a very low cost per bit, allowing for high disk storage capacity.

In most UP or SMP implementations, a first level of cache (L1) lies between these two ends. L1 cache is a fast form of memory with a small capacity. The number of CPU cycles that are needed for the processor to load data from L1 depends on the L1 implementation. In the PowerPC implementation, L1 is on the CPU chip itself, so it takes only one cycle to load data from L1. When the L1 cache is outside the processor chip, several cycles are required to load data from L1. A typical L1 capacity is around 32 to 64 KB.

You might also find a second level of cache (L2), which is another form of high-speed memory. It takes about two to ten cycles to load data from L2.

The main memory is the third level of memory. Its access time is slow in comparison to L1 and L2, but much faster than disks. Twenty to fifty CPU clock cycles are needed to load data from the memory, and the capacity can reach several gigabytes.

Figure 12 shows the relationship between proximity of memory to the processor versus the number of cycles required to access it.

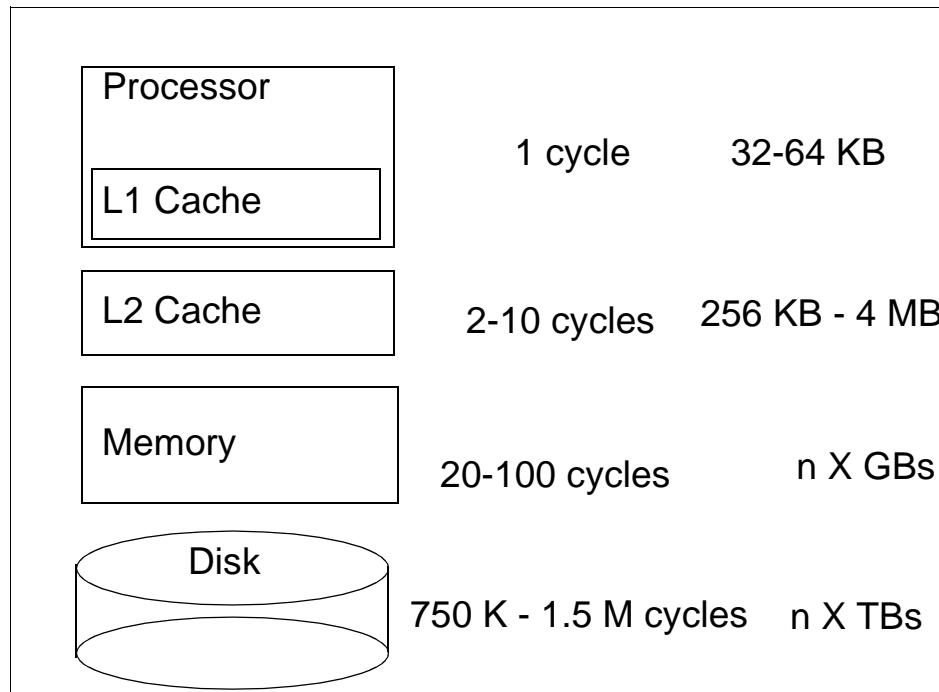


Figure 12. Memory Hierarchy

2.5.1.1 Cache Hit versus Cache Miss

When a CPU fetches a memory address, if the data is found in the cache, it is a cache hit. Otherwise, it is a cache miss. If a cache miss occurs, the data is loaded from main memory to the CPU and stored in the cache to take advantage of the higher speed of the cache for a future fetch of the same memory address. The hit ratio is the percentage of cache hits. Logically, the higher the hit ratio, the better the system performance. All RISC System/6000 computers use a scheme called *set associativity* to reduce the number of cache misses.

2.5.1.2 Set Associativity

Set associativity is a mechanism that allows for more efficient sharing of the L1 data cache among multiple programs running simultaneously. One of the main differences between the POWER3 and the POWER2 (also PowerPC) architectures is the way set associativity is implemented. PowerPC has a 4-way set-associative data cache, while POWER3 has a 128-way set-associative data cache.

To gain a better understanding of set associativity, the concept of cache lines needs to be introduced. A cache line is the basic unit of transfer between main memory of cache. It is 128 bytes long on POWER3.

Conceptually, memory is sectioned into contiguous 128-byte lines, each one starting on a cache-line boundary whose hardware address is a multiple of 128. The cache is similarly sectioned and all data transfer between cache and memory is in units of these lines. If, for example, a particular floating point number is required to be loaded into a floating point register, so that computation may be done with it, then the whole cache line containing that number is transferred from memory to cache.

Figure 13 shows the L1 data cache mapped onto memory, using POWER2, while Figure 14 on page 42 shows it using POWER3.




Figure 13. The 4-Way Set Associative PowerPC/POWER2 L1 Data Cache

The 128-Way Set-Associative POWER3 Data Cache




Figure 14. The 128-Way Set Associative POWER3 L1 Data Cache

Each cell represents a cache line; each column is called a *congruence class*. Figure 13 on page 41 contains 128 congruence classes. There are four rows in the cache area of the figure. Each of these rows contains 128 cache lines of 128 bytes each ($128 \times 128 = 16 \text{ KB}$).

Any particular line from memory may only be loaded into a cache line in the same congruence class: for POWER2, into one of only four locations; for POWER3, into one of 128 locations.

Generally, when a new line from memory is loaded into the cache, existing data must be displaced. If the previous contents of the line have been modified, the line must be stored back into memory. The algorithm used by the hardware for selecting which cache line to use is an approximation of Least Recently Used on POWER2, and is Round-Robin on POWER3.

The set associative structure of the cache can lead to a reduction in its effective size. Suppose successive data elements are being processed that are regularly spaced in memory (that is with a constant *stride*). With the POWER2 cache, the worst case is when the stride is exactly 16 KB (128 lines \times 128 bytes each) or a multiple of 16 KB. In this case, all elements will lie in the same congruence class and the effective cache size will be only four

cache lines (since any particular line from memory may only be loaded into a cache line in the same congruence class). This effect happens with strides that are a multiple of a power of two.

The POWER3 cache, with its much greater degree of set associativity, is much less susceptible to this problem than the POWER2 cache. Strides of multiples of 1024 bytes will cause all the data to be in the same congruence class but will only cause a reduction in apparent cache size of a factor of 4. Odd multiples of 512 will halve this effective size. This is, however, minor compared with the possible reduction by a factor of 128 on POWER2.

2.5.1.3 Effects of L1 Cache Misses on Performance

The effects of cache misses on performance for POWER3 (the Model 260) can be illustrated in Figure 15. A load instruction is to be performed to a floating-point register (assume data is already resident in memory).

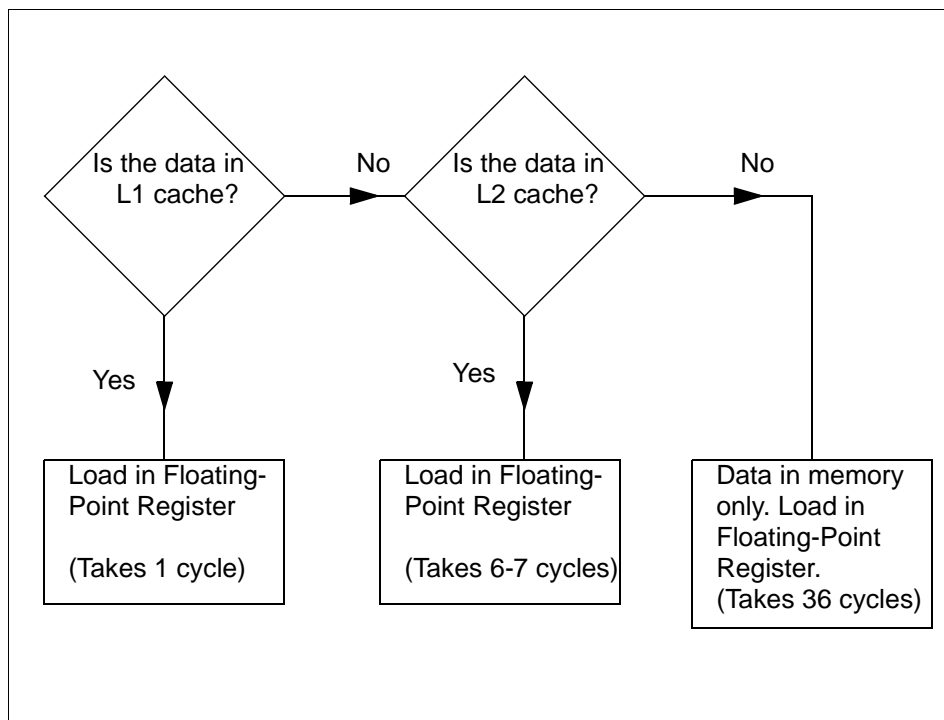


Figure 15. Loading Instructions from Memory to a Floating-Point Register

If the data is in memory only, it takes about 36 cycles; if it is in L1 cache, it takes one cycle. That is, the cost of a cache miss to memory is 35 cycles. The same timing applies to storing data from registers into memory. If the store is

into a previously unreferenced line, the complete line must be fetched from memory first before the new value can be stored into it. If a cache line is overwritten by newly accessed data, then, if the data from the old line is needed again, it must be reloaded and another cache miss taken.

2.5.1.4 Cache Line Prefetch

Because of the relatively large number of cycles needed for a cache miss, POWER3 has a mechanism for mitigating the performance impact for sequentially accessed data. For up to four streams of data, the hardware attempts to detect sequential access and initiates the loading of subsequent lines in parallel, so they stream into cache behind the first line without waiting for the miss to occur. This important feature of POWER3 is not present on POWER2.

2.5.1.5 Cache Coherency Problem

In an SMP, all of the processors have their own cache to improve data locality. Only the main memory is shared. Since caches are not shared, it is necessary to keep all the processors' caches coherent. The coordination of current versus stale data between processors is called *cache coherency* and is one of the most important issues when designing an SMP system.

Consider an application that runs on two processors, processor 1 and processor 2, as shown in Figure 16.




Figure 16. SMP Cache Coherency Problem

Assume that processor 1 loads into its cache memory address 0x0123, which happens to contain the character A. Then processor 2 writes B into address 0x0123. If processor 1 wants to load address 0x0123 again, what will happen? In a naive implementation, processor 1 will see the value A in its cache and load that value because it does not know that processor 2 has already changed the same memory address in its cache. This is called the cache coherency problem.

2.5.1.6 Snooping

One solution to the cache coherency problem is snooping. Snooping is hardware logic that is added to the processor and is affiliated with normal memory reads. While a memory operation is in process, the other caches in the system are interrogated (snooped) to see if the data currently resides there. If one processor needs to write into a cache, a message is broadcast which causes that entry to be invalidated in all other caches. This is called a cross invalidate. Cross invalidate reminds the processor that the value in the cache is not valid. In this case, there is a cache miss. The processor must then look for the correct value in another cache or in the main memory.

Since cross invalidate increases cache misses and the snooping protocol adds to the bus traffic, solving the cache consistency problem reduces the performance and scalability of all SMP systems. In other words, because latency time for a request can be widely variable due to the location of data and snooping activity, adding more processors is not always the best thing to do in an attempt to improve response time for a given request.

The POWER3 processor has this extra logic. In fact, all PowerPC processors except the 603, POWER and POWER2 have this function.

Bus snooping is used to drive a MESI four-state protocol as it is described in the following section.

2.5.1.7 MESI Protocol

The unit of storage in the cache is the cache line. The size of the cache line is implementation dependent. The PowerPC has a cache line size that is 64 bytes. This cache line is divided into two 32-byte sectors. The POWER3 has a 128-byte cache line and a single sector.

The PowerPC maintains cache coherency on a cache sector basis by using the four-state MESI protocol. Each sector has two state bits. The four states are:

- M (modified)** The addressed sector is valid in this cache only. The value in this sector has been changed in the cache, but the change is not yet reflected in memory.
- E (exclusive)** The addressed sector is valid in this cache only. The data is consistent with system memory.
- S (shared)** The addressed sector is valid in this cache and at least one other cache. It is still consistent with system memory.
- I (invalid)** The addressed sector is not valid in the cache.

2.5.1.8 False Sharing

Figure 17 illustrates false sharing. Here, a whole cache unit of storage (a cache sector in the case of the IBM SMP systems) is represented.




Figure 17. False Sharing

Suppose processor 1 and processor 2 loaded the same memory address in their cache. If processor 1 changes only a portion of the cache sector, d1 for example, the cache consistency logic will invalidate all the sectors in the processor 2 cache. Then, if processor 2 tries to modify another portion of its cache sector, d2, for example, which is still invalid since the whole sector is invalid, a cache miss will occur. This is called false sharing.

Thus, false sharing increases cache misses and bus traffic, and this may cause the SMP throughput to be reduced. The bigger the size of the cache line, the higher the miss rate. Some implementations have a 256-byte cache line. Both 64-byte and 128-byte cache line sizes are supported by the IBM SMP design, and these cache line sizes are equivalent to the coherency size of the system.

Chapter 3. Detailed Product Descriptions

This chapter takes an in-depth look at the hardware packaging and features that comprise the IBM RS/6000 43P 7043 Model 150 and Model 260 workstations. These workstations use the PowerPC and POWER3 architectures and offer large memory capacities, better I/O performance, powerful graphics, and flexible configurations.

3.1 IBM RS/6000 43P 7043 Model 150 Product Description

The IBM RS/6000 43P 7043 Model 150 is a powerful PowerPC 604e uni-processor with a 32-bit implementation of the PowerPC architecture. The Model 150 is well-suited for demanding 3-D applications including MCAD engineering analysis, petroleum, and seismic interpretation. It can also be deployed as a workgroup server or stand-alone business system, providing function that includes internet, database, application, and print/file serving.

Figure 18 shows a Model 150 with available monitor, keyboard, and mouse.




Figure 18. RS/6000 43P 7043 Model 150 (with Peripherals)

The RS/6000 43P 7043 Model 150 is designed to allow for system growth in memory, disk, and media devices as the customer's needs grow. By combining AIX reliability and function with industry standard bus architectures, the Model 150 offers exceptional value for multi-user applications and provides an excellent solution in the entry to mid-range workstation marketplace.

3.1.1 Standard Features

In this section, the standard features of the Model 150 and their most important characteristics are described.

3.1.1.1 Processor Subsystem

The PCI-based IBM RS/6000 43P Model 150 features the PowerPC 604e microprocessor with the internal clock speed of 375 MHz. With the internal processor clock to system bus clock frequency ratio set at 9:2, the system bus is run at 83 MHz. The processor 604e can support bus clock ratios of 1:1, 3:2, 2:1, 5:2, 3:1, 7:2, 4:1, 9:2 and 5:1.

The 604e integrated L1 cache is split into separate instruction and data caches. Each of them contains 32 KB four-way set associative data arrays organized as 32-byte cache lines with no sectoring.

The Model 150 L2 cache subsystem consists of a L2 cache controller and L2 memory. L2 cache memory consists of on-board Tag-Ram and 1 MB parity synchronous pipelined burst SRAM.

The Model 150 has two cache copy-back write buffers.

3.1.1.2 System Planar Board

One planar is installed in every 7043-150. It provides the system bus and card connections for the riser card, a voltage regulator and clock module, four memory DIMM slots. Power supply connections, I/O connections, and a pin-grid array socket for the PowerPC processor. Only single processor boards are available.

Figure 19 on page 51 shows the system board location for a Model 150, Table 5 on page 52 lists the purpose of the various connections.

System Board Locations (for Model 150)




Figure 19. RS/6000 43P 7043 Model 150 - System Board Locations

Table 5. RS/6000 43P 7043 Model 150 - System Board Components Location

Location	Description	Location	Description
B1	Battery connector	J27, J30	System fan connector
J2	5x5 Auxiliary 5v connector	J28	Memory connector #2
J3	Audio input connector	J32	Privileged-access password jumper
J4	Internal SCSI Conn. #2	J33	Internal SCSI 16-bit conn.
J5	CD-ROM connector	J35, J36	SCSI security jumper
J6	CD-ROM audio connector	J38	External SCSI jumper
J7, J8	Power connector	J39	Riser card connector
J10	Media fan connector	J40	Power switch connector
J11	Memory connector	J41	Op panel connector
J13	Audio output connector	J42	Power indicator LED conn.
J15	Remote power-up jumper	J43	HDD activity LED connector
J16	Diskette drive connector	J44	Voltage regulator connector
J17	Microphone jack	J45	Ethernet twisted pair conn.
J18	Reset connector	J46	Internal speaker connector
J19	Memory connector #3	J49	Parallel port connector
J20	Power-on password jumper	J50	Ethernet AUI
J22	Headphone jack	J53	Ethernet EPROM connector
J23	Serial port connector #2	J56	Mouse port connector
J24	Serial port connector #1	J59	Keyboard connector
J25	Table port connector	J139	Memory connector #4
J26	CPU fansink connector		

3.1.1.3 Memory Subsystem

The Model 150 comes standard with a 128 MB of ECC SDRAM memory. It supports four industry standard DIMM sockets and can support up to 1 GB of system memory. Memory DIMM size includes 64 MB, 128 MB, or 256 MB. The memory DIMMs can be populated in any order.

3.1.1.4 Bays

The system contains five bays, three of which accommodate the following standard devices:

- 3.5-inch diskette drive bay
It accommodates the standard 1.44 MB diskette drive.
- 5.25-inch CD-ROM drive bay
It accommodates the standard 32X CD-ROM drive with the convenience of loading a CD without a caddy.
- 5.25-inch disk drive bay.

It accommodates the standard 4.5 GB Ultra SCSI disk drive that can be upgraded in the initial machine order to 9.1 GB Ultra SCSI disk. For more detailed information about the disk drives see Table 29 on page 121.




Figure 20. RS/6000 43P 7043 Model 150 - Front View

The empty bays can accommodate a second CD-ROM, tape drive or other media device. A 4.5 GB (# 2900) or 9.1 GB (# 2908) hard disk can be installed in bay 3 with a mounting kit (# 6509).

3.1.1.5 Internal CD-ROM Drive

The 32X SCSI-2 Internal CD-ROM Drive is a tray loading CD-ROM drive providing up to 4800 KB/s maximum media data transfer rate. It is a 5.25 inch half-high form factor, single ended, 8 bit, multi-session capable, CD-ROM drive that provides state-of-the-art performance and supports existing 600 MB CD-ROM discs. It supports either the standard 12 cm or the smaller 8 cm CD-ROM disk.

For more detailed technical information about the CD-ROM drive, refer to Table 30 on page 122.

3.1.1.6 Standard I/O Ports

The system provides the following standard I/O ports:

Keyboard port	The optional Quiet Touch Keyboard may plug into this port.
Mouse port	A port is provided to support an optional three-button mouse.
Tablet port	If required, a tablet can be connected.
Serial port	Two 9-pin D-shell serial connectors are available for external RS232 access.
Parallel port	One 25-pin parallel port is provided that maintains complete compatibility with machines containing ISA, EISA, and Micro Channel parallel ports.

3.1.1.7 PCI Adapter Slots

The Model 150 contains five 32-bit PCI slots. Following is a brief description according to their slots number.

- Slots 1 and 5 are short slots
- Slots 2, 3, and 4 are long slots
- Slots 2, and 3 are connected to the primary PCI bus on the planar
- Slots 1, 4, and 5 are on the secondary PCI bus provided by the bridge chip on the riser card.




Figure 21. RS/6000 43P 7043 Model 150 - Rear View

3.1.1.8 Integrated Ethernet Adapter

An integrated 10/100 Mbps Ethernet AUI/RJ45 port (10Base5/100BaseT) is part of the base system.

3.1.1.9 Integrated Ultra SCSI Controller

The Model 150 has an integrated Ultra SCSI controller, which is operated at a speed of 40 MB/s.

3.1.1.10 AIX Operating System

The RS/6000 Model 150 includes one to two user license, for a full-function AIX Version 4.2.1 or Version 4.3.2 package at no additional charge.

3.1.1.11 Year2000 and EuroReady

The IBM RS/6000 43P 7043 Model 150 is Year 2000 ready. Please refer also to 1.8, "Year 2000" on page 18 for details.

When used with AIX 4.3.2, this product is EuroReady.

3.1.1.12 Publications

The following publications will be shipped with the Model 150.

Table 6. Publications Shipped with the Model 150

Order Number	Title
SA38-0510	7043 43P Series Setup Instructions
SA38-0511	7043 43P Series Users Guide
SA23-2690 (USA customers only)	Customer Support Information
SA23-2652	System Unit Safety Information
SA38-0538	PCI Adapter Placement Reference Guide
Customer Installable Options Library CD-ROM	
Warranty Booklet (USA customers only) ¹	

¹ The CD-ROM is not orderable, it is shipped only at manufacturing. No form number is available.

Only the ship-pack publications appear in the sales manual. A complete list of publications is provided in the PUBS section of HONE.

3.1.2 Reliability, Availability, and Serviceability (RAS)

Following are the features that provide the IBM RS/6000 Model 150 reliability, availability, and serviceability.

3.1.2.1 Reliability, Fault Tolerance, and Data Integrity

The reliability of the Model 150 system starts with reliable components, devices, and subsystems. During the design and the development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing, systems go through a testing process to ensure the highest product quality level.

The Model 150 system memory offers ECC (Error-Checking and Correcting) fault-tolerant features. ECC corrects environment-induced single-bit intermittent memory failures. As well as single hard failures. With ECC, the majority of memory failures will not impact system operation. ECC also provides double-bit memory error detection which protects data integrity in the event of the double-bit memory failures. The system bus and PCI buses are designed with parity error detection.

Disk mirroring and disk controller duplexing capability are provided by the AIX operating system.

The journaled file system (JFS) of AIX operating system maintains file system consistency and prevents data loss when the system is abnormally halted due to a power failures.

An available RAID hardware feature external to the system provides data integrity and fault tolerance in the event of the disk failure.

3.1.2.2 Fault Monitoring Functions

Following are the functions used to monitor faults during the boot process.

- POST (Power-on-Self Test) that checks the processor, L2 cache, memory and associated hardware that are required for proper booting of the operating system every time the system is powered on. If a non-critical error is detected, or if the error(s) occur in the resources that can be removed from the system configuration, the booting process will proceed to completion. The error(s) are logged in the system Non Volatile RAM.
- Disk drive fault tracking is a facility that can alert the system administrator of an impending disk failure before it impacts customer operation.
- AIX log facility where hardware and software failures are recorded and analyzed (by the Error Log Analysis routine) to provide warning to the system administrator on the causes of system problems. This also

enables IBM service representatives to bring along needed replacement hardware components during a service, thus minimize system repair time.

3.1.2.3 Environment Monitoring Functions

Following is a list of the provided environment monitoring functions.

- Fan speed monitoring to provide warning and orderly system shutdown when the speed is out of operational specification.
- AC power loss sensing to provide orderly system shutdown.

3.1.2.4 Error Handling and Reporting

In an unlikely event of system hardware or environmental induced failure, the system Run-Time Error Capture capability will systematically analyze the hardware error signature to determine the cause(s) of failure. The analysis result will be stored in system Non Volatile RAM (NVRAM). When the system can be successfully rebooted either manually or automatically, the error will be automatically reported to AIX operating system. The Error Log Analysis (ELA) can be used to display the failure cause and the physical location of failure hardware.

3.1.2.5 Availability Enhancement Functions

On-line (concurrent) Diagnostics with Error Log Analysis and Service Aids allow administrators or IBM service representatives to diagnose potential system malfunction without interrupting end-user operation.

The auto-restart (reboot) option, when enabled, automatically reboots the system following an unrecoverable software error or hardware failure.

3.1.2.6 Serviceability

The Model 150 is designed for customer setup of the machine and for subsequent addition of most features (adapters/devices). The Model 150 supports personnel to remotely log into a system to review error logs and perform remote maintenance.

The diagnostics consists of Stand-Alone Diagnostics (system dedicated to service) and on-line diagnostics. On-line diagnostics, when installed, are resident with AIX on the disk or system. They can be booted in single user mode (service mode), run in maintenance mode, or run concurrently (concurrent mode) with other applications. They have access to the AIX Error Log and the AIX configuration data.

Service Mode allows checking of system devices and features.

Concurrent Mode allows the normal system functions to continue while selected resources are being checked.

Maintenance Mode allows checking of most system resources.

The System Management Service (SMS) error log is accessible from the SMS menu for tests performed through SMS programs.

3.1.2.7 Service Director for RS/6000

The Service Director for RS/6000 is available at no additional charge if the RS/6000 processor is covered by an IBM Warranty or IBM Maintenance Service Agreement. When installed on the RS/6000, the Service Director for RS/6000 can enhance IBM's ability to provide the system with maintenance service. The Service Director for RS/6000 is responsible for the following operations:

- Monitors and analyzes system errors, and if needed, can automatically place a service call to IBM without customer intervention.
- Can reduce the effect of business disruptions due to unplanned system outages and failures.
- Performs problem analysis on a subset of hardware-related problems and, with customer authorization, can automatically report the results to IBM service.

3.1.2.8 On-line Customer Support (OCS)

On-line Customer Support for hardware problem reporting may be performed through remote login by RS/6000 specialists using normal AIX facilities. The Service Director for RS/6000 software can also be used for this capability. AIX support offerings will be under AIXSERV and Service Director for RS/6000.

3.1.3 Supported Optional Features

In this section, *optional features* are defined as internal optional devices, external adapters, and external subsystems that can be configured on the RS/6000 Model 150.

3.1.3.1 Memory

RS/6000 Model 150 comes standard with 128 MB of ECC SDRAM memory. The Model 150 can be expanded up to 1 GB of memory. This allows the customer to easily upgrade the processor cards by selecting one to four

memory DIMMs (64 MB, 128 MB, and 256 MB). The available memory specifications are shown in Table 7.

Table 7. Optional Memory for Model 150

Memory Density (SDRAM DIMM)	256 MB	128 MB	64 MB
Feature Code	# 4169	# 4150	# 4149
Average Seek Time	10 ns	10 ns	10 ns
Number of PINs	168	168	168
Type	Buffered SDRAM	Buffered SDRAM	Buffered SDRAM

Note

The DIMMs do not have to be installed in pairs.

3.1.3.2 Internal Disk Drives

Three internal Ultra SCSI disks drives can be accommodated in the internal bays of the system. For more detailed information about the disk drives see Table 29 on page 121.

3.1.3.3 Internal CD-ROM Drive

One additional 32X CD-ROM drive can be added to the system (# 2624) into media bay 3. For this the Media Bay Disk Drive Mounting Kit (# 6509) is required. You can find more detailed technical information about the CD-ROM drive in Table 30 on page 122.

3.1.3.4 Internal Tape Drive

One 12 / 24 GB 4 mm internal tape drives can be installed in the internal bay of the system. For detailed information about the internal tape drives, refer to the Table 31 on page 122.

3.1.3.5 Keyboard

New Stealth Black Quiet Touch Keyboards (# 87xx¹) are available in 37 languages. They have the same quiet and soft touch as the Pearl White Quiet Touch Keyboards (# 66xx¹). Like their predecessors, these keyboards have removable wrist rests for added comfort, and a 3-meter cable for added flexibility.

(¹ xx depends on the country)

The present Pearl White Quiet Touch Keyboards will also continue to be available with black systems giving users maximum flexibility in choosing their desktop color scheme.

There is Euro currency symbol support on appropriate national language versions.

3.1.3.6 Mouse

A Stealth Black 3-Button (# 8741) mouse is available. The new mouse has a 3-meter cable with a standard 6-pin mini-DIN connector providing the flexibility users want. It is also still available in pearl white (# 6041).

3.1.3.7 Graphics Adapters

For 2D graphics applications the Model 150 offers the POWER GXT120P, GXT250P, and GXT255P accelerator cards. The GXT120P meet the graphics needs for many commercial applications and provides an entry point for workstation users. It is also a good price/performance choice for server environments.

For advanced graphics functions, the GXT3000P graphics accelerator for 3D visualization marks a breakthrough in performance and functionality for design and visualization solutions.

Refer to 3.1.4, "RS/6000 43P 7043 Model 150 Adapter Placement Guidelines" on page 66 or to 5.2, "Graphics Adapters" on page 135 for more information.

Note

GXT3000P requires AIX 4.3.2. Support for AIX 4.2.1 is planned for a later date.

3.1.3.8 Graphic Peripheral Devices

The following graphic peripheral devices are officially supported:

- 6093-011 CursorPad Tablet
- 6093-012 Large Tablet
- 6093-021 IBM Tablet Model 021
- 6094-010 Dials
- 6094-020 Lightened Programming Function Keyboard (LPFK)
- 6094-030 Spaceball
- 6094-031 Spaceball

- 6094-040 Spaceball
- 6094-600 Magellan Enhanced Input Device

3.1.3.9 Monitors

The IBM P72, P92, and P202 color monitors are the premium graphics monitors for RS/6000 workstations. These monitors are for customers who demand image accuracy in a large screen format combined with performance, quality, reliability, and connectivity.

Table 8. List of Orderable Monitors for Model 150

Description	Feature Code	Screen Size (in inches)
P72 Color Monitor, Stealth Black	# 3623	17.0
P72 Color Monitor, Pearl White	# 3620	17.0
P92 Color Monitor, Stealth Black	# 3625	19.0
P92 Color monitor, Pearl White	# 3621	19.0
P202 Color Monitor, Stealth Black	# 3626	21.0
P202 Color Monitor, Pearl White	# 3622	21.0

3.1.3.10 Personal Microphone

The IBM personal microphone (# 6308) is a lightweight microphone and headset that conveniently mounts over and behind the user's ear. The microphone element is attached to the earpiece with a thin tubular boom. The IBM personal microphone allows adjustments to be made to position the microphone and headset for either the left or right ear.

The IBM personal microphone is suitable for a wide range of applications, but is especially suited for speech recognition applications including dictation and navigation in an office environment. The microphone allows consistent performance by providing for an optimum distance between the voice source and the microphone element and by providing concurrent feedback to the user through the personal headset.

3.1.3.11 Software Pre-Install

Licensed programs may be pre-installed on the media selected. To order this option, specify code (# 5005) on the initial system order. The pre-installed licensed programs will be at the same release level as programs being shipped from the program libraries at the time the system is manufactured.

3.1.3.12 PCI Storage Adapters

The following PCI SCSI adapters are supported in the Model 150

- PCI Single-Ended Ultra SCSI Adapter
- PCI Differential Ultra SCSI Adapter
- Serial Storage Architecture (SSA)

PCI Single-Ended Ultra SCSI Adapter (# 6206)

The PCI Single-Ended Ultra SCSI has a maximum data transfer rate of 40 MB/s which is two times the maximum data transfer rate of Fast/Wide adapters and requires one PCI slot

PCI SCSI-2 Differential Ultra SCSI Adapter (# 6207)

The PCI Differential Ultra SCSI adapter is the next generation of SCSI-2 performance with a maximum data transfer rate of 40 MB/s and occupies one PCI slot.

PCI SSA Multi-Initiator/RAID EL Adapter (# 6215)

The PCI SSA Multi-Initiator/RAID EL adapter allows up to eight RS/6000 to attach to an SSA loop in non-RAID mode and two RS/6000 to attach in RAID-5 mode and requires one PCI slot.

You can find more detailed information about the PCI storage adapters in 4.8, "PCI Storage Adapters" on page 124.

3.1.3.13 Communications Adapters

The following communications adapters are officially supported in Model 150.

- IBM PCI Token-Ring adapter (# 2920)
- IBM PCI 10/100 Mbps Ethernet adapter (# 2968)
- ATM adapters
 - Turboways 155 UTP PCI ATM Adapter (# 2963)
 - Turboways 155 MMF PCI ATM Adapter (# 2988)
 - Turboways 25 ATM PCI ATM Adapter (# 2998)
- FDDI adapters
 - SysKonnnect SK-NET FDDI-LP SAS PCI Adapter (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS PCI Adapter (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS PCI Adapter (# 2743)
- Asynchronous adapters
 - PCI 8-Port Asynchronous Adapter EIA-232/RS-422 (# 2943)
 - PCI 128-Port Asynchronous Controller (# 2944)

- 2-Port Multiprotocol PCI Adapter (# 2962)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)

3.1.3.14 Supported Devices

The following is a list of the external subsystems that are officially supported by the RS/6000 Model 150.

External Disk Subsystems

- 7131-105 Multi-Storage Tower
- 7131-405 SSA Multi-Storage Tower
- 7133-010 SSA Disk Subsystem (Rack-Mounted)
- 7133-020 SSA Disk Subsystem (Rack-Mounted)
- 7133-500 SSA Disk Subsystem
- 7133-600 SSA Disk Subsystem
- 7137-412 Disk Array Subsystem (1.97 - 6.91 GB)
- 7137-413 Disk Array Subsystem (3.94 - 13.77 GB)
- 7137-414 Disk Array Subsystem (8.39 - 29.36 GB)
- 7137-415 Disk Array Subsystem (17.7 - 70.83 GB)
- 7203-001 Portable Disk Drive
- 7204-010 1 GB Disk Drive
- 7204-112 1.1 GB Fast/Wide Disk Drive
- 7204-113 2.2 GB Fast/Wide Disk Drive
- 7204-114 4.5 GB Fast/Wide Disk Drive
- 7204-118 18 GB SE External Disk Drive
- 7204-139 9.1 GB Fast/Wide Disk Drive
- 7204-215 2 GB Differential Disk Drive
- 7204-315 2 GB Differential Fast/Wide Disk Drive
- 7204-317 2.2 GB Differential Fast/Wide Disk Drive
- 7204-325 4.5 GB Differential Fast/Wide Disk Drive
- 7204-339 9.1 GB External Disk Drive
- 7204-402 External 2.2 GB Disk Drive
- 7204-404 External 4.5 GB Disk Drive
- 7204-418 18 GB DE External Disk Drive

External SCSI Devices

- 3995-C60 Optical Library Dataserver (1 Drive)
- 3995-C62 Optical Library Dataserver (2 Drives)
- 3995-C64 Optical Library Dataserver (2 Drives)
- 3995-C66 Optical Library Dataserver (4 Drives)
- 3995-C68 Optical Library Dataserver (4 Drives)
- 7209-002 Optical Disk Drive
- 7209-003 Optical Disk Drive
- 7210-005 CD-ROM Drive
- 7210-010 CD-ROM Drive (Quad speed)
- 7210-015 CD-ROM Drive (8x speed)

Tape Drive Subsystems

- 3490-F00 IBM Magnetic Tape Subsystem
- 3490-F01 IBM Magnetic Tape Subsystem
- 3494-L12 Tape Library Dataserver
- 3570-B00 IBM Magstar MP Tape Subsystem
- 3570-B01 IBM Magstar MP Tape Subsystem (Floor Library, 1 Drive)
- 3570-B02 Tape Subsystem (2 Drives)
- 3570-B11 IBM Magstar MP Tape Subsystem (Rack Mounted, 1 Drive)
- 3570-B12 IBM Magstar MP Tape Subsystem (Rack Mounted, 2 Drives)
- 3570-B1A IBM Magstar MP Tape Drive Model B1A
- 3570-C00 Magstar Magnetic Tape Subs * 1 Tape Lib (1), SCSI
- 3570-C01 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C02 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3570-C11 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C12 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3575-L06 IBM Magstar MP Tape Library DataServer - 60 Cartridge
- 3575-L12 IBM Magstar MP Tape Library DataServer - 120 Cartridge
- 3575-L18 IBM Magstar MP Tape Library DataServer - 180 Cartridge
- 3575-L24 IBM Magstar MP Tape Library DataServer - 240 Cartridge
- 3575-L32 IBM Magstar MP Tape Library DataServer - 320 Cartridge

- 3590-B11 High Performance Tape Subsystem
- 3590-B1A High Performance Tape Subsystem
- 3590-C12 Silo Compatible Tape Subsystem
- 7205-311 DLT Tape Drive Bridge Box
- 7206-005 4 GB 4mm Tape Drive
- 7206-110 Cartridge Autoloader 4mm, DDS-3, SCSI-2, DAT
- 7207-011 525 MB 1/4" Tape Drive
- 7207-012 1.2 GB 1/4" Tape Drive
- 7207-122 4 GB External SLR5 QIC Tape Drive
- 7207-315 13 GB 1/4" Tape Drive
- 7208-011 5 GB 8mm Tape Drive
- 7208-341 20 GB External 8mm Tape Drive
- 7331-205 8mm Tape Library
- 7331-305 8mm Tape Library
- 7332-005 4mm DDS-2 Tape Autoloader
- 7332-310 4mm Tape Autoloader
- 7336-205 4mm Tape Library
- 7337-305 Digital Linear Tape Subsystem
- 7337-306 Digital Linear Tape Subsystem
- 9348-012 1/2" 9-Track Magnetic Tape Unit

Communications Subsystems

- 7318 Model P10 Serial Communication Network Server
- 7318 Model S20 Serial Communication Network Server

Printers, Plotters and ASCII Terminals

These devices are all supported through standard ports. Refer to the list of supported devices in the appropriate section of the AIX sales manual pages.

Modems

- IBM 5841, 5853, 5865
- Hayes compatible
- Recal-Vadic compatible

3.1.4 RS/6000 43P 7043 Model 150 Adapter Placement Guidelines

This machine is designed for customers to install adapters. Use the following table to identify specific slot location options for the following adapters in your 7043 Model 150 system.

When determining a slot in which to place a new adapter, you should start from the top of Table 9. The list of slot numbers represent the order in which the slots should be used.

Table 9. Adapter Placement Guidelines for Model 150

Adapter	Feature Code	max. Quantity	Slot Usage
Ultimedia Video Capture	2639	1	Any Slot
Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter	2708	1	Any Slot
PCI SysKonnnect SK-NET FDDI-LP SAS ²	2741	2	Any Slot
PCI SysKonnnect SK-NET FDDI-LP DAS ²	2742	2	Any Slot
PCI SysKonnnect SK-NET FDDI-UP SAS ²	2743	2	Any Slot
POWER GXT3000P	2825	1	Slot 3
POWER GXT120P ¹	2838	4	Any Slot
POWER GXT550P	2845	1	Slot 2, 3
POWER GXT255P ¹	2852	4	Any Slot
PCI Token Ring Adapter	2920	4	Any Slot
PCI 8-Port Asynchronous EIZ-232/RS-422A	2943	2	Any Slot
PCI 128-Port Asynchronous Controller	2944	2	Any Slot
PCI 2-Port Multiprotocol	2962	2	Slots 5, 4, 1
PCI 155 TURBOWAYS ATM UTP	2963	2	Any Slot
PCI 10/100 Mbps Ethernet Tx	2968	4	Any Slot
PCI 155 TURBOWAYS ATM MMF	2988	2	Any Slot
PCI 25 TURBOWAYS ATM	2998	4	Any Slot
PCI Single-Ended Ultra SCSI	6206	2	Any Slot
PCI Differential Ultra SCSI	6207	2	Any Slot
PCI SSA Multi-Initiator/RAID EL	6215	1	Slots 2, 3

¹ If both the POWER GXT120P and the POWER GXT255P are installed, the POWER GXT255P must be installed in slot 2 or 3.

² If two SysKonnnect SK-NET FDDI PCI adapters are installed, they cannot be installed in slots next to each other.

3.1.5 Service Package

Service package kit (# 3752) contains service publications, wrap plugs, and diagnostic CD-ROM. It also includes Maintenance Analysis Procedures (MAPs) that are not common to other RS/6000 systems. The *Diagnostics Information for Multiple Bus Systems*, SA38-0509 contains information common to all multiple bus units and is used primarily for problem isolation by a customer's service specialist trained on the unit being serviced.

These publications are intended to be used in the customer's enterprise by the customer's own service specialists, and not by IBM Service personnel.

Two wrap plugs are included for use in diagnostic testing of the following connections to the system units:

- RS-232 serial I/O port
- Parallel I/O port

This kit also contains a system diagnostic CD-ROM, along with a test diskette for use by the customer service specialist (not by IBM service personnel).

3.1.5.1 Options Library (Hard Copy)

A hard copy version of the options publications (# 3753) can be obtained. This hard copy is also available through PUBORDER (*CIO - Option Pubs Library English*, SK2T-7452).

3.2 RS/6000 43P 7043 Model 260 Product Description

The IBM RS/6000 43P 7043 Model 260 is a new addition to the RS/6000 workstation and workgroup server family. The Model 260 is an affordable 64-bit symmetric multiprocessing (SMP) system with true multi-threaded application support and extended floating-point capabilities. This is IBM's first announced 64-bit workstation and workgroup server. With the new GXT3000P graphics adapter, it provides significant performance enhancements over the Model 150 and the Model 240.



Figure 22. RS/6000 43P 7043 Model 260 (with Peripherals, and Spaceball)

The Model 260 is a 1- or 2-way SMP system utilizing the 64-bit 200 MHz POWER3 processor with 4 MB of Level 2 (L2) cache per processor. The system memory can be expanded up to 4 GB. Ethernet, Ultra SCSI, and a service processor are integrated within the Model 260. The Model 260 also contains five PCI slots and five bays. Two of the slots are 64-bit slots and three are 32-bit slots. A 1.44 MB 3.5-inch diskette drive and a 32X maximum speed CD-ROM occupy two of the bays. A variety of 2D and 3D graphics adapters, including a new advanced 3D graphics adapter, are available for the Model 260. This graphics capability, along with excellent price and

performance, position the Model 260 as a single-seat MCAD design and analysis solution.

In addition, the Model 260's SMP capability and robust set of disk drive and communications features make it an excellent entry technical and commercial server. It is the lowest cost 64-bit SMP server in the RS/6000 family of products. The AIX Version 4.2.1 or 4.3.2 operating system is included with each system and can be pre-installed, if desired.

The Model 260 is designed for use as a mid- to high-range graphics workstation, or as an entry server. It supports MCAD, CAE, Geophysical, and entry technical server customers. It is also well suited for use as an entry SMP/64-bit server for ISVs and customers who want a cost-effective development platform for developing and testing applications that will run on larger RS/6000 systems. The dual processing power of the Model 260 and its small package make it an excellent solution for Internet service providers and customers that need or want a stand-alone internet server.

The system is designed for customer setup of the machine and for subsequent addition of most features (adapters/devices).

3.2.1 Model 260 System Architecture

The RS/6000 43P 7043 Model 260 is the first product of the RS/6000 family based on the POWER3 architecture. The design incorporates high-speed processor and memory buses, a more efficient 4 MB Level 2 (L2) cache design, a high-performance memory-I/O controller, and fast Synchronous Dynamic Random Access Memory (SDRAM).

For more information about the POWER3 Architecture refer to 2.4.1, "The POWER3 Microprocessor" on page 35.

3.2.2 Standard Features

In this section, the standard features of the Model 260 and their most important characteristics are described.

3.2.2.1 Processor Subsystem

The PCI-based RS/6000 43P 7043 Model 260 features the POWER3 64-bit microprocessor at 200 MHz. The processor includes an L1 cache split into 32 KB instruction and a 64 KB data cache. Each processor has 4 MB of Level 2 (L2) cache. Systems with only one processor must have installed the CPU card in slot 1.

The processor subsystem is implemented on pluggable cards. Two processor card slots are available. If two processor cards are installed in the system, they must both be the same type or the system will not power on.

3.2.2.2 System Planar Board

One planar is installed in a Model 260. It provides the system bus and card connections for the riser card, a voltage regulator, and clock module and four memory DIMM slots. Power supply connections, I/O connections, and a pin-grid array socket for the PowerPC processor.

One or two 200 MHz POWER3 processors can be installed.

Figure 23 shows the location of the I/O board components, Table 10 lists their attributes.




Figure 23. RS/6000 43P 7043 Model 260 - System Planar & I/O Board

Table 10. RS/6000 43P 7043 Model 260 - System Planar & I/O Board Components

Location	Description	Location	Description
J1	Operator panel power connector	J2	Operator panel audio connector

Location	Description	Location	Description
J3	PCI fan connector	J4	I/O board connector to system board
J5	VPD module connector	J6	ESP external connector
J7	Internal SCSI connector	J8, J9	64-bit PCI connectors
J10	CD ROM Audio	J11	Battery connector
J12, J13, J14	32-bit PCI connectors	J15	Service processor external connector
J16	Power connector	J17	Diskette drive connector
J18	Tablet connector	J19	Power connector
J20	AUI Ethernet connector	J21	Cec fan connector
J23	Serial port connector	J24	Keyboard/Mouse connector
J25	RJ45 Ethernet connector	J26	Audio in/out
J27	Parallel port connector	J28	External SCSI connector

3.2.2.3 Memory Subsystem

The base configuration comes with 256 MB Error Checking and Correcting (ECC) SDRAM (2 x 128 MB DIMMs).

The Model 260 has an enhanced memory controller running at 100 MHz memory bus speed. See 3.1.3.1, "Memory" on page 58 for more information how the memory can be expanded.

Synchronous DRAM

The SDRAM operates faster than conventional DRAM. The SDRAM can supply data every 10 nanosecond clock cycle once its initial access time of 60 nanoseconds has been satisfied. As with conventional memory, the SDRAM banks have a period of latency after an access. During this period, the memory bank would be designated as busy. To achieve better performance, another memory bank (already precharged) needs to be accessed during that period of latency. The memory I/O controller uses heavy queueing to be able to make an intelligent decision on which memory bank to access next. The controller will queue up to eight read commands and eight write commands and attempts to schedule all reads to non-busy banks. It fills in with writes when no reads are pending.

3.2.2.4 Operator Panel

The operator panel on the front of the Model 260 contains a power-on switch (on/off), a reset button, a power-on light, a hard disk drive in use light, and an LCD display for function and diagnostic feedback.

Table 11 lists the operator panel components shown in Figure 24.

Table 11. RS/6000 43P 7043 Model 260 - Operator Panel Components

Number #	Description
1	Power-On Switch: Turns computer power on and off.
2	Reset Switch: Resets your system unit.
3	Power-On LED: Glows when computer is on.
4	Hard Disk Drive Status LED: Glows when computer is reading from or writing to the disk drive.
5	Operator Panel Display: Displays current status of system unit startup, or diagnostic information in the event of a hardware problem.
6	Headset Receptacle: For connecting audio headphones.
7	Microphone Receptacle: For connecting audio microphone.




Figure 24. RS/6000 43P 7043 Model 260 - Operator Panel

3.2.2.5 Bays

Five bays are available for the system, as shown in Figure 25. One is reserved for the diskette drive (bay 1) and one for the internal CD-ROM drive (bay 2).

The Model 260 comes with a standard Ultra-SCSI 16-bit Disk Drive installed in bay 4. Up to three internal disks can be installed in the system.

Bay 3, a 5.25 inch bay, can accommodate a second CD-ROM, tape drive, or other disk device. If you want to install a 4.5 GB Ultra SCSI Disk Drive (# 2900) or a 9.1 GB Ultra SCSI 16-Bit 1-Inch High Disk Drive (# 2908), you have to use the 5.25 inch Media Bay Mounting Kit (# 6508).

Bays 4 and 5, which are both 3.5 inch bays, can contain a 4.5 GB Ultra SCSI Disk Drive (# 3028) or 9.1 GB Ultra SCSI Disk Drive (# 3029).

Table 12 lists a summary of the different storage devices and in which bay they can be installed:

Table 12. Storage Device Features

Feature Code	Device	Max. Quantity	Bay #
2624	600 MB 32X CD-ROM	2	2, 3
2900	4.5 GB Ultra SCSI Disk	1	3
2908	9.1 GB Ultra SCSI Disk	1	3
3028	4.5 GB Ultra SCSI Disk	2	4, 5
3029	9.1 GB Ultra SCSI Disk	2	4, 5
6156	8 mm 20/40 GB Tape Drive	1	3
6159	4 mm 12/24 GB Tape Drive	1	3
6508	5 1/4" Media Bay Mounting Kit	1	3

For more detailed information about the disk drives, see Table 29 on page 121.

Note

The Model 260 cannot boot from drives connected to SCSI external connector.




Figure 25. RS/6000 43P 7043 Model 260 - Front View

3.2.2.6 Internal CD-ROM Drive

The 32X SCSI-2 Internal CD-ROM Drive is a tray-loading CD-ROM drive providing up to 4800 KB/s maximum media data transfer rate. It is a 5.25 inch, half-high form factor, single ended, 8 bit, multi-session capable, CD-ROM drive that provides state-of-the-art performance and supports existing 600 MB CD-ROM discs. It supports either the standard 12 cm or the smaller 8 cm CD-ROM disk.

For more detailed technical information about the CD-ROM drive, refer to Table 30 on page 122.

3.2.2.7 Rear Bulkhead Connections and Components

Table 13 lists the connections and components found on the rear bulkhead as shown in Figure 26.

Table 13. RS/6000 43P 7043 Model 260 - Rear View Components

Number #	Description	Number #	Description
1	Power Supply LED	2	Power Supply Test Switch
3	Power Connector	4	Serial Connector S1
5	Serial Connector S2	6	AUI Ethernet Connector
7	Tablet Connector	8	Keyboard Connector
9	Mouse Connector	10	Audio Line Out
11	Audio Line In	12	RJ45 Ethernet Connector
13	Parallel Connector	14	External SCSI Connector
15	PCI Slot 5 (32-bit)	16	PCI Slot 4(32-bit)
17	PCI Slot 3 (32-bit)	18	PCI Slot 2 (64-bit)
19	PCI Slot 2 (64-bit)		

3.2.2.8 Standard I/O Ports

The system provides the following standard I/O ports:

- Keyboard port** The quiet touch keyboard, an optional feature, can be plugged into this standard 6-pin mini-DIN port.
- Mouse port** A port is provided to support an optional three-button mouse.
- Serial port** Two serial ports are provided for a TTY terminal, Modem and other serial devices.
- Parallel port** One parallel port is provided.

3.2.2.9 PCI Adapter Slots

There are five PCI slots in the Model 260:

- Slots 1 and 2 are 64-bit and run at 50 MHz.
- Slots 3, 4 and 5 are 32-bit and run at 33 MHz.
- All slots are long slots.




Figure 26. RS/6000 43P 7043 Model 260 - Rear View

3.2.2.10 Integrated Ethernet Adapter

An integrated 10/100 Mbps Ethernet AUI/RJ45 port (10Base5/100BaseT) is part of the base system.

3.2.2.11 Integrated Ultra SCSI Controller

The Model 260 has two integrated Ultra SCSI controllers. One for internal and one for external connections. Both are operated at a speed of 40 MB/s.

3.2.2.12 AIX Operating System

The RS/6000 43P 7043 Model 260 includes a one to two user, full-function AIX Version 4.2.1 or 4.3.2 package at no additional charge.

3.2.2.13 Year 2000 and EuroReady

The RS/6000 43P 7043 Model 260 is Year 2000 ready. Please refer also to 1.8, "Year 2000" on page 18.

When used with AIX 4.3.2, this product is EuroReady.

3.2.2.14 Publications

Table 14 provides the publications shipped with the Model 260 (additional copies are available):

Table 14. Publications Shipped with the Model 260

Order Number	Title
SA38-0555	7043 Model 260 Hardware Setup
SA38-0553	7043 Model 260 User's Guide
SA23-2690	Customer Support Information (US customers only)
SA23-2652	System Unit Safety Information
SA38-0538	PCI Adapter Placement Reference Guide

Table 15 provides the additional publications that are available. To order, contact your IBM representative.

Table 15. Available Publications for Model 260

Order Number	Title
SA38-0554	7043 Model 260 Service Guide
SA38-0509	Diagnostics Information for Multiple Bus Systems
SA38-0516	IBM RS/6000 Adapters, Devices and Cables Information for Multiple Bus Systems

3.2.3 Reliability, Availability, and Serviceability (RAS)

The RS/6000 43P 7043 Model 260 integrates key reliability, availability, and serviceability attributes that are described in the following sections.

3.2.3.1 Reliability, Fault Tolerance, and Data Integrity

The reliability of the Model 260 system starts with reliable components, devices, and subsystems. During the design and development process, subsystems go through rigorous verification and integration testing processes. During system manufacturing, systems go through a thorough testing process to ensure the highest product quality level.

The Model 260 system L2 Cache and system memory offers Error Checking and Correcting (ECC) fault-tolerant features. ECC corrects environmentally induced single-bit intermittent memory failures, as well as single-bit hard failures. With ECC, the likelihood of memory failures will be substantially reduced. ECC also provides double-bit memory error detection that protects data integrity in the event of a double-bit memory failure. System memory also provides 4-bit packet error detection that helps protect data integrity in the event of a DRAM chip failure. The system bus, I/O bus, and PCI buses are designed with parity error detection.

Disk mirroring and disk controller duplexing capability are also provided by the AIX operating system.

The Journaled File System of the AIX operating system maintains file system consistency and reduces the likelihood of data loss when the system is abnormally halted due to a power failure.

An available RAID hardware feature for the system provides data integrity and fault tolerance in the event of disk drive failure.

3.2.3.2 Service Processor

A service processor is included in the system, providing an immediate means to diagnose, check status, and sense operational conditions of a remote system, even when the main processor is inoperable. The service processor enables firmware and operating system surveillance, several remote power controls, environmental monitoring, reset and boot features, along with remote maintenance and diagnostic activities, including console mirroring. The service processor can place calls to report surveillance failures, critical environmental faults and critical processing faults.

For more information about the Service Processor please refer to Chapter 7, "Service Processor" on page 185.

3.2.3.3 Fault Monitoring Functions

Built-in Self-Test (BIST) and Power-on Self-Test (POST) checks processor, L2 cache, memory and associated hardware, that are required for proper booting of the operating system every time the system is powered on. If a non-critical error is detected, or if the error(s) occur in the resources which can be removed from the system configuration, the booting process will proceed to completion. The error(s) are logged in the system non-volatile RAM.

Disk drive fault tracking that can alert the system administrator of potential disk failure before it impacts customer operation.

The AIX log facility where hardware and software failures are recorded and analyzed (by Error Log Analysis routine) to provide warning to the system administrator on the causes of system problems. This also enables IBM service representatives to bring along needed replacement hardware components when a service call is placed, thus minimizing system repair time.

3.2.3.4 Mutual Surveillance

The service processor can monitor the operation of the firmware during the boot process, and it can monitor the operating system for loss of control. It also allows the operating system to monitor for service processor activity. The service processor can take appropriate action, including calling for service, when it detects that the firmware or the operating system has lost control. Likewise, the operating system can request a service processor repair action if necessary.

3.2.3.5 Environmental Monitoring Functions

The following is a list of the environmental monitoring functions.

- Temperature monitoring that increases the fan speed rotation when ambient temperature is above the normal operating range
- Temperature monitoring to warn the system administrator of potential environmental related problems (for example, air conditioning and air circulation around the system) so that appropriate corrective actions can be taken before a critical failure threshold is reached, and to provide orderly system shutdown when operating temperature exceeds the critical level
- Fan speed monitoring to provide warning and an orderly system shutdown when the speed is out of operational specification
- DC voltages monitoring to provide warning and an orderly system shutdown when the voltage(s) are out of operational specification

- AC power loss sensing to provide an orderly system shutdown

3.2.3.6 Error Handling and Reporting

In the unlikely event of system hardware or environmentally induced failure, the system Run-Time Error Capture capability systematically analyzes the hardware error signature to determine the cause of failure. The analysis result will be stored in system Non-Volatile RAM (NVRAM). When the system can be successfully rebooted either manually or automatically, the error will be reported to the AIX operating system. The Error Log Analysis (ELA) can be used to display the failure cause and the physical location of failing hardware.

With the integrated service processor, the system has the ability to automatically send out an alert through the phone line to a pager, or call for service in the event of critical system failure.

Environmental and Power Warnings (EPOW) is an option that is implemented in the Model 260. This option provides environmental conditions of the system to the user in order to try and minimize damage to the system if limits are close to being exceeded. This option also indicates failure conditions to AIX for appropriate action.

3.2.3.7 Availability Enhancement Functions

The following enhancements provide additional system availability function.

- On-line (Concurrent) Diagnostics with Error Log Analysis and Service Aids that allow administrators or IBM service representatives to diagnose potential system malfunction without interrupting end-user operation
- Auto-restart (reboot) option, when enabled, to automatically reboot the system following an unrecoverable software error, software hang, hardware failure, or environmental-induced (AC power) failure

3.2.3.8 Serviceability

The Model 260 is designed for Customer setup of the machine and for subsequent addition of most features (adapters/devices). For a fee, IBM Service can perform the installation.

The Model 260 allows support personnel to remotely log into a system to review error logs and perform remote maintenance. The Model 260 Service Processor enables the analysis of a system that will not boot.

The diagnostics consist of Stand-alone Diagnostics (system dedicated to service) and On-line Diagnostics. On-line Diagnostics, when installed, are resident with AIX on the disk or system. They can be booted in single user

mode (service mode), run in maintenance mode, or run concurrently (concurrent mode) with other applications. They have access to the AIX Error Log and the AIX Configuration Data.

Service Mode allows checking of system devices and features. Concurrent Mode allows the normal system functions to continue while selected resources are being checked. Maintenance Mode allows checking of most system resources.

The System Management Services (SMS) error log is accessible from the SMS menu for tests performed through SMS programs. The Service Processor error log is accessible from the Service Processor menu for results of Service Processor tests.

3.2.3.9 Service Director for RS/6000

The Service Director for RS/6000 is available at no additional charge if the RS/6000 processor is covered by an IBM Warranty or IBM Maintenance Service Agreement. When installed on the RS/6000, the Service Director for RS/6000 can enhance IBM's ability to provide the system with maintenance service.

The Service Director for RS/6000:

- Monitors and analyzes system errors, and if needed, can automatically place a service call to IBM without customer intervention
- Can reduce the effect of business disruptions due to unplanned system outages and failures
- Performs problem analysis on a subset of hardware-related problems and, with customer authorization, automatically reports the results to IBM Service

3.2.3.10 On-line Customer Support (OCS)

On-line Customer Support for hardware problem reporting may be performed by remote login by RS/6000 specialists using normal AIX facilities. The Service Director for RS/6000 software can also be used for this capability.

AIX Support offerings will be under AIXSERV and Service Director for RS/6000.

3.2.4 Supported Optional Features

In this section, *optional features* are defined as internal optional devices. Adapters are described in more detail in Chapter 5, "Adapters, Overview and

Configuration” on page 127 and storage subsystems are described in Chapter 4, “Storage Architectures and Devices” on page 97.

3.2.4.1 Additional Processor

Two processor card slots are available on the system board. Since the system comes with one standard POWER3 processor, a second processor can be added (# 4342). As already mentioned, if two processor cards are installed in the system, they must both be the same type or the system will not power on.

3.2.4.2 Memory

The RS/6000 43P 7043 Model 260 comes standard with 256 MB of ECC SDRAM memory and can be expanded up to 4 GB.

Note

As a 64 bit operating system, AIX 4.3.2 fully supports the maximum memory configuration. AIX 4.2.1, however, limits support to a maximum of 3 GB of real memory to allow mapping of I/O in the address space.

The DIMMs can be installed in one or two memory cards. Each card can contain 16 DIMMs, which are available as 32 MB DIMMs (# 4107) or 128 MB DIMMs (# 4110). If you need more than 16 DIMMs, you have to order the Memory Expansion Feature (# 4098).

The first memory card must be installed in memory card slot 2. Memory DIMMs must be ordered and installed in pairs on the memory cards. They must be installed starting at the bottom of each card (card slot J1 and J2). See Figure 27 for the memory card configuration.




Figure 27. Slot Configuration on Memory Card

3.2.4.3 Internal Disk Drives

The system comes with a standard 4.5 GB Enhanced Ultra SCSI Disk. There is also a 9.1 GB Enhanced Ultra SCSI Disk available. You can add one additional 4.5 GB (# 3028) or 9.1GB (# 3029) disk into the second media bay.

A third 4.5 GB (# 2900) or 9.1 GB (# 2908) disk can be installed into bay 3 with the Media Bay Mounting Kit (# 6508).

For more detailed information about the disk drives, see Table 29 on page 121.

3.2.4.4 Internal CD-ROM Drive

One additional 32X CD-ROM drive can be added to the system (# 2624) into media bay 3, with the Media Bay Disk Drive Mounting Kit (# 6508).

More detailed information about the CD-ROM drive can be found on Table 30 on page 122.

3.2.4.5 Internal Tape Drives

Two internal tape drives are available:

- 20/40 GB 16-Bit 8mm Internal Tape Drive (Black Bezel) (# 6156)

This feature consists of a 5.25-inch half-high, single-ended, 16-bit tape drive. This drive provides a high capacity tape drive for save/restore and archiving functions. This tape drive uses IBM 8 mm data cartridges and is compression capable, providing a capacity of up to 40 GB, a 400 percent increase over the previous 5/10 GB 8mm internal tape drive.

- 12/24 GB 4 mm Internal Tape Drive (# 6159)

The 12/24 GB 4 mm Internal Tape Drive is a 5.25 inch, half-high, single-ended 8-bit tape drive, which provides a high capacity for save/restore and archive functions. This tape drive uses IBM 4 mm data cartridges and is compression capable, providing a capacity of up to 24 GB, a significant increase in capacity over the previous 4/8 4-mm internal tape drives (when using DDS-3 media).

For more information about the tape drives, see Table 32 on page 123.

3.2.4.6 Keyboard

Stealth Black Quiet Touch Keyboards (# 87xx¹) are available in 37 languages. They have the same quiet and soft touch as the Pearl White Quiet Touch Keyboards (# 66xx¹). Like their predecessors, these keyboards have removable wrist rests to add comfort, and a 3-meter cable to add flexibility.

(¹ xx depends on the country)

The Pearl White Quiet Touch Keyboards will continue to be available with black systems giving users maximum flexibility in choosing their desktop color scheme.

There is Euro currency symbol support on appropriate national language versions.

3.2.4.7 Mouse

A Stealth Black 3-Button (# 8741) mouse has a 3-meter cable with a standard 6-pin mini-DIN connector providing the flexibility users want. It is still available in pearl white (# 6041).

3.2.4.8 Graphics Adapters

For 2D graphics applications, the Model 260 offers the POWER GXT120P, GXT250P, and GXT255P accelerator cards. The GXT120P meets the graphics need for many commercial applications and provides an entry point for workstation users. It is also a good price/performance choice for server environments.

For advanced graphics functions, the GXT3000P graphics accelerator for 3D visualization marks a breakthrough in performance and functionality for design and visualization solutions. This graphics accelerator, matched with IBM POWER3 workstation, utilizes IBM's multi-threaded PHIGS and OpenGL libraries, which exploit the 260's SMP architecture to provide excellent speed and performance for demanding applications in:

1. Mechanical Computer-Aided Design (MCAD) and engineering for aerospace and automotive
2. Petroleum exploration and production
3. Scientific visualization
4. Other technical industries

Refer to 3.2.5, "RS/6000 43P 7043 Model 260 Adapter Placement Guidelines" on page 92, or to 5.2, "Graphics Adapters" on page 135 for more information.

Note

The GXT3000P requires AIX 4.3.2. Support for AIX 4.2.1 is planned for a later date.

3.2.4.9 Graphics Peripheral Devices

The following graphic peripheral devices are officially supported:

- 6093-011 CursorPad Tablet
- 6093-012 Large Tablet
- 6093-021 IBM Tablet Model 021
- 6094-010 Dials
- 6094-020 Lightened Programming Function Keyboard (LPFK)
- 6094-030 Spaceball
- 6094-031 Spaceball
- 6094-040 Spaceball
- 6094-600 Magellan Enhanced Input Device

3.2.4.10 Monitors

Three monitor types are provided. Each of them can be obtained in pearl white and black color:

Table 16. List of Orderable Monitors for Model 260

Description	Feature Code	Screen Size
IBM P72 Color Monitor	# 3620	17"
IBM P72 Color Monitor, Stealth Black, Captured Cable	# 3623	17"
IBM P92 Color Monitor	# 3621	19"

Description	Feature Code	Screen Size
IBM P92 Color Monitor, Stealth Black, Captured Cable	# 3625	19"
IBM P202 Color Monitor	# 3622	21"
IBM P202 Color Monitor, Stealth Black, Captured Cable	# 3626	21"

All monitors include the following features:

- Video Electronics Standards Association (VESA) Display Data Channel (DDC) capability for Plug and Play and asset tracking when the attaching system unit's hardware and software is DDC-enabled
- Energy Star guidelines meet energy efficiency
- Power management capability to meet the power consumption requirements of Swedish National Board for Industrial and Technical Development (NUTEK)
- TCO-95 compliant - incorporates a universal power supply with power factor correction (EN 61000-3-2:1995) providing worldwide operating capability in many instances
- Microprocessor-controlled, multi-frequency operation so that you can change, replace, or upgrade most system units without replacing these monitors
- Timings for the VESA 85 Hz non-interlaced refresh rate display modes for 640 x 480, 800 x 600, and 1024 x 768 addressabilities supported

Also earlier announced displays, like the 5081, 6091-19i, PowerDisplay17, PowerDisplay20, P50, P70, P200, or P201, are supported.

To attach a 5081 or 6091 display, the 6091 attachment cable (# 4217) is required, which is 1.8 meters (6 feet) in length.

3.2.4.11 Personal Microphone

The IBM Personal Microphone (# 6308) is a lightweight microphone and headset that conveniently mounts over and behind the user's ear. The microphone element is attached to the earpiece with a thin tubular boom. The IBM Personal Microphone allows adjustments to be made to position the microphone and headset for either the left or right ear.

The IBM Personal Microphone is a perfect addition to the microphone supplied with your system unit. It will allow each user to have their own microphone.

The microphone is suitable for a wide range of applications, but is especially suited for speech recognition applications including dictation and navigation in an office environment. The microphone allows consistent performance by providing for by an optimum distance between the voice source and the microphone element and providing concurrent feedback to the user through the personal headset.

3.2.4.12 Software Pre-Install

Licensed programs may be pre-installed on the media selected. To order this option, specify code (# 5005) on the initial system order. The pre-installed licensed programs will be at the same release level as programs being shipped from the program libraries at the time the system is manufactured.

3.2.4.13 PCI Storage Adapters

The following PCI SCSI adapters are supported in the Model 260:

- PCI Single-Ended Ultra SCSI Adapter
- PCI Differential Ultra SCSI Adapter
- Serial Storage Architecture (SSA)

PCI Single-Ended Ultra SCSI Adapter (# 6206)

The PCI Single-Ended Ultra SCSI has a maximum data transfer rate of 40 MB/s, which is two times the maximum data transfer rate of Fast/Wide adapters and requires one PCI slot.

PCI SCSI-2 Differential Ultra SCSI Adapter (# 6207)

The PCI Differential Ultra SCSI adapter is the next generation of SCSI-2 performance with a maximum data transfer rate of 40 MB/s. It occupies one PCI slot.

PCI SSA Multi-Initiator/RAID EL Adapter (# 6215)

The PCI SSA Multi-Initiator/RAID EL adapter allows up to eight RS/6000 to attach to an SSA loop in non-RAID mode and two RS/6000 to attach in RAID-5 mode and requires one PCI slot.

You can find more detailed information about the PCI storage adapters in 4.8, "PCI Storage Adapters" on page 124.

3.2.4.14 Communications Adapters

The following communication adapters are officially supported:

- IBM PCI Token-Ring Adapter (# 2920)
- PCI Ethernet adapters
 - IBM 10/100 Mbps Ethernet PCI Adapter (# 2968)
 - Gigabit Ethernet - SX PCI Adapter (# 2969)
- ATM adapters
 - Turboways 155 PCI UTP ATM Adapter (# 2963)
 - Turboways 155 PCI MMF ATM Adapter (# 2988)
 - Turboways 25 ATM PCI Adapter (# 2998)
- FDDI adapters
 - SysKonnnect SK-NET FDDI-LP SAS PCI (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS PCI (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS PCI (# 2743)
- Asynchronous adapters
 - 8-Port Asynchronous Adapter EIA-232/RS-422 PCI (# 2943)
 - 128-Port Asynchronous Controller (# 2944)
- 2-Port Multiprotocol PCI Adapter (# 2962)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)
- ARTIC960Hx 4-Port T1/E1 PCI Adapter (# 2948)

More detailed technical information about these PCI adapters can be found in 5.1, "Communication Adapters" on page 127.

For more information about the adapter placement guidelines, see 3.2.5, "RS/6000 43P 7043 Model 260 Adapter Placement Guidelines" on page 92.

3.2.4.15 Supported Devices

The following is a list of the external subsystems that are officially supported by the RS/6000 Model 260. This list is not all-inclusive. Please refer to the sales manual of the external machine type for further attach support information.

External Disk Subsystems

- 7131-105 Multi-Storage Tower
- 7131-405 SSA Multi-Storage Tower

- 7133-010 SSA Disk Subsystem (Rack-Mounted)
- 7133-020 SSA Disk Subsystem (Rack-Mounted)
- 7133-500 SSA Disk Subsystem
- 7133-600 SSA Disk Subsystem
- 7137-412 Disk Array Subsystem (1.97 - 6.91 GB)
- 7137-413 Disk Array Subsystem (3.94 - 13.77 GB)
- 7137-414 Disk Array Subsystem (8.39 - 29.36 GB)
- 7137-415 Disk Array Subsystem (17.7 - 70.83 GB)
- 7203-001 Portable Disk Drive
- 7204-010 1 GB Disk Drive
- 7204-112 1.1 GB Fast/Wide Disk Drive
- 7204-113 2.2 GB Fast/Wide Disk Drive
- 7204-114 4.5 GB Fast/Wide Disk Drive
- 7204-118 18 GB SE External Disk Drive
- 7204-139 9.1 GB Fast/Wide Disk Drive
- 7204-215 2 GB Differential Disk Drive
- 7204-315 2 GB Differential Fast/Wide Disk Drive
- 7204-317 2.2 GB Differential Fast/Wide Disk Drive
- 7204-325 4.5 GB Differential Fast/Wide Disk Drive
- 7204-339 9.1 GB External Disk Drive
- 7204-402 External 2.2 GB Disk Drive
- 7204-404 External 4.5 GB Disk Drive
- 7204-418 18 GB DE External Disk Drive

External SCSI Devices

- 3995-C60 Optical Library Dataserver (1 Drive)
- 3995-C62 Optical Library Dataserver (2 Drives)
- 3995-C64 Optical Library Dataserver (2 Drives)
- 3995-C66 Optical Library Dataserver (4 Drives)
- 3995-C68 Optical Library Dataserver (4 Drives)
- 7209-002 Optical Disk Drive
- 7209-003 Optical Disk Drive

- 7210-005 CD-ROM Drive
- 7210-010 CD-ROM Drive (Quad speed)
- 7210-015 CD-ROM Drive (8x speed)

Tape Subsystems

- 3490-F00 IBM Magnetic Tape Subsystem
- 3490-F01 IBM Magnetic Tape Subsystem
- 3494-L12 Tape Library Dataserver
- 3570-B00 IBM Magstar MP Tape Subsystem
- 3570-B01 IBM Magstar MP Tape Subsystem (Floor Library, 1 Drive)
- 3570-B02 Tape Subsystem (2 Drives)
- 3570-B11 IBM Magstar MP Tape Subsystem (Rack Mounted, 1 Drive)
- 3570-B12 IBM Magstar MP Tape Subsystem (Rack Mounted, 2 Drives)
- 3570-B1A IBM Magstar MP Tape Drive Model B1A
- 3570-C00 Magstar Magnetic Tape Subs * 1 Tape Lib (1), SCSI
- 3570-C01 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C02 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3570-C11 Magstar Magnetic Tape Subs * 1 Tape Lib (20), SCSI
- 3570-C12 Magstar Magnetic Tape Subs * 2 Tape Lib (20), SCSI
- 3575-L06 IBM Magstar MP Tape Library DataServer - 60 Cartridge
- 3575-L12 IBM Magstar MP Tape Library DataServer - 120 Cartridge
- 3575-L18 IBM Magstar MP Tape Library DataServer - 180 Cartridge
- 3575-L24 IBM Magstar MP Tape Library DataServer - 240 Cartridge
- 3575-L32 IBM Magstar MP Tape Library DataServer - 320 Cartridge
- 3590-B11 High Performance Tape Subsystem
- 3590-B1A High Performance Tape Subsystem
- 3590-C12 Silo Compatible Tape Subsystem
- 7205-311 DLT Tape Drive Bridge Box
- 7206-005 4 GB 4mm Tape Drive
- 7206-110 Cartridge Autoloader 4mm, DDS-3, SCSI-2, DAT
- 7207-011 525 MB 1/4" Tape Drive
- 7207-012 1.2 GB 1/4" Tape Drive

- 7207-122 4 GB External SLR5 QIC Tape Drive
- 7207-315 13 GB 1/4" Tape Drive
- 7208-011 5 GB 8 mm Tape Drive
- 7208-341 20 GB External 8 mm Tape Drive
- 7331-205 8 mm Tape Library
- 7331-305 8 mm Tape Library
- 7332-005 4 mm DDS-2 Tape Autoloader
- 7332-310 4 mm Tape Autoloader
- 7336-205 4 mm Tape Library
- 7337-305 Digital Linear Tape Subsystem
- 7337-306 Digital Linear Tape Subsystem
- 9348-012 1/2" 9-Track Magnetic Tape Unit

Communications Subsystems

- 7318 Model P10 Serial Communication Network Server
- 7318 Model S20 Serial Communication Network Server

Printers, Plotters and ASCII Terminals

These devices are all supported through standard ports. Refer to the list of supported devices in the appropriate section of the AIX sales manual pages.

Modems

- IBM 5841, 5853, 5865
- Hayes compatible
- Recal-Vadic compatible

3.2.5 RS/6000 43P 7043 Model 260 Adapter Placement Guidelines

Table 17 lists the slot plugging guidelines for the Model 260 system.

When determining a slot in which to place a new adapter, you should start from the top of the table. The list of slot numbers represent the order in which the slots should be used.

Table 17. RS/6000 43P 7043 Model 260 - Adapter Placement Guidelines

Adapter	Feature Code	Slot Usage	max. Quantity
POWER GXT120P	# 2838	Any slot	2
POWER GXT250P	# 2851	Any slot	4
POWER GXT255P	# 2852	Any slot	4
POWER GXT3000P	# 2825	Slot 2	1
PCI Single Ended Ultra SCSI	# 6206	Any slot	2
PCI Differential Ultra SCSI	# 6207	Any slot	2
SysKonnect SK-NET FDDI-LP SAS PC	# 2741	Any slot	2
SysKonnect SK-NET FDDI-LP DAS PCI	# 2742	Any slot	2
SysKonnect SK-NET FDDI-UP SAS PC	# 2743	Any slot	2
Token Ring PCI	# 2920	Any slot	4
10/100 Mbps Ethernet PCI	# 2968	Any slot	4
1 GB Ethernet PCI	# 2969	Slot 1, 2	2
Eicon ISDN DIVA PRO 2.0 PCI S/T	# 2708	Any slot	1
2-Port Multiprotocol PCI	# 2962	Any slot	2
ARTIC960Hx 4-Port Selectable PCI	# 2947	Any slot	2
ARTIC960Hx 4-Port T1/E1 PCI	# 2948	Slots 4, 3, 2	2
8-Port Asynchronous EIA-232E/RS-422A PCI	# 2943	Any slot	2
128-Port Async Controller PCI	# 2944	Any slot	2
155 TURBOWAYS ATM PCI MMF	# 2988	Any slot	2
TURBOWAYS 25 ATM PCI	# 2998	Any slot	4
155 TURBOWAYS ATM PCI UTP	# 2963	Any slot	2
PCI SSA Multi-Initiator/RAID EL	# 6215	Any slot	1
Ultimedia Video Capture	# 2639	Any slot	1

For more information on the adapter refer to Chapter 5, “Adapters, Overview and Configuration” on page 127 or see *PCI Adapter Placement Reference*, SA38-0538.

3.2.6 Service Package

This kit (# 3752) contains service publications, wrap plugs, and a diagnostic CD-ROM. It also includes Maintenance Analysis Procedures (MAPs) that are not common to other RS/6000 systems.

The *Diagnostics Information for Multiple Bus Systems*, SA38-0509, contains information common to all multiple bus system units and is used primarily for problem isolation by a customer's service specialist trained on the unit being serviced. These publications are intended to be used in the customer's enterprise by the customer's own service specialists.

Two wrap plugs are included for use in diagnostic testing of the following connections to the system unit:

- RS-232 serial I/O port
- Parallel I/O port

This kit also contains a system diagnostic CD-ROM, along with a test diskette, for use by the customer's service specialist.

3.2.6.1 Options Library (Hard Copy)

A hard copy version of the options publications (# 3753) can be obtained. This hard copy is also available as *CIO - Option Pubs Library English*, SK2T-7452.

3.3 IBM Multimedia Kit for RS/6000

The IBM Multimedia Kit for RS/6000 (# 8679) consists of a video camera, speakers, and connecting cables, as shown in Figure 28.



Figure 28. Model 150 with Multimedia Kit for RS/6000

The following are part of the package:

- Philips Desktop Video Camera
 - Flicker-free picture mode (shutter speed is fixed to avoid noticeable flickering in environments with fluorescent lighting)
 - Microphone for live audio (connection with a 3.5 mm miniature audio plug)
 - Capability to focus manually close to 0 cm
 - Back-light compensation (improved image with back-light situations)
 - White-balance fixed mode (for optimal color stability)
 - An LED that shows the camera is on
 - Camera output through an RCA composite video connection connects to the Ultimedia Video Capture Adapter

- NTSC format
- High-power output speaker system includes two premium Infinity(R) sound quality speakers specially designed to fit and match IBM equipment.

With the Multimedia Kit for RS/6000 and the Ultimedia Video Capture adapter (# 2639), IBM provides the hardware you need to capture and display video contents within IBM Workgroup Conferencing for AIX and Ultimedia Services. Both software products are included in the AIX Bonus Pack CD, which is shipped with the base operating system.

For more information about Workgroup Conferencing see Chapter 1.7, "Workgroup Conferencing" on page 15.

Note

The Multimedia Kit (# 8679) was announced in October 1998 and will be available in the US and Latin America. IBM plans to announce the Multimedia Kit for other countries at a later date.

Chapter 4. Storage Architectures and Devices

In this chapter, the internal storage features of the Models 150 and 260 are described. Storage standards, SCSI, and SSA are covered. The different technologies are discussed, and additional information is provided in the cabling sections. As well, a description of the storage devices and adapters follows as well.

4.1 Storage Boot Devices

The system boot image can be executed from the following base storage devices:

- SCSI CD-ROM
- SCSI DASD devices connected to a base SCSI controller
- SCSI tape drives

4.2 SCSI Overview

A Small Computer System Interface (SCSI) is a bus-level interface. Computers may communicate with a large number of devices of different types connected to the system unit through a SCSI controller and daisy-chained cable. The attached devices include such peripherals as fixed disks, CD-ROMs, printers, plotters, and scanners. The SCSI controller may be in the form of an adapter, or may be integrated on the motherboard. There are several terms and concepts used in discussing SCSI technology that require definition.

Note

The American National Standards Institute (ANSI) refers to the different SCSI specifications using the SCSI-I and SCSI-II type nomenclature. IBM uses SCSI-1 and SCSI-2 nomenclature in official product names.

4.2.1 SCSI-I

SCSI is a standard defined by the American National Standards Institute (ANSI). The original SCSI standard is defined in ANSI standard X3.131-1986. It defines an 8-bit interface with a burst-transfer rate of 5 MB/s and a 5 MHz clock (1 byte transferred per clock cycle).

It is sometimes referred to as SCSI-I to differentiate it from the generic term SCSI. SCSI-I was the first of all SCSI technologies to come about and was the fastest controller interface at the time.

4.2.2 SCSI-II

The SCSI-II specification gained final approval from ANSI in 1994 as standard X3T9.2/375R Revision 10K. SCSI-II allowed far better performance than SCSI-I. It defines extensions to SCSI that allow for 16- and 32-bit devices, a 10 MB/s synchronous transfer rate for 8-bit transfers and 20 MB/s for 16-bit transfers. Other enhancements are discussed in the text that follows. SCSI-II comes in many varieties, such as SCSI-II, SCSI-II Fast and SCSI-II Fast/Wide.

The interface for SCSI-II also defined additional control signals, as well as additional data signals. This meant that the maximum number of devices supported by one SCSI channel was increased from 8 to 16.

4.2.2.1 Common Command Set

The SCSI-II standard defines a set of commands that must be interpreted by all devices that are attached to a SCSI bus. This is called the common command set. Unique devices may implement their own commands, which can be sent by a device driver and interpreted by the device. The advantage of this architecture is that the SCSI adapter does not have to change when new devices with new capabilities are introduced.

4.2.2.2 Tagged Command Queuing

Tagged Command Queuing (TCQ) is a SCSI-II enhancement. It increases performance in disk-intensive server environments. With SCSI-I systems, only two commands could be sent to a fixed disk. The disk would store one command while operating on the other. With TCQ, it is possible to send multiple commands to the hard disk because the disk stores the commands and executes each command in the sequence that gives optimal performance.

Also with TCQ, the adapter has more control over the sequence of disk operations. For example, the adapter can tell the device to execute the next command immediately, or it can instruct it to finish everything it already has been given before completing the new command.

4.2.2.3 Disconnect/Reconnect

Some commands take a relatively long time to complete (for example, a seek command could take roughly 10 ms). With this feature, the target can disconnect from the bus while the device is positioning the heads (seeking).

Then, when the seek is complete and data is ready to be transferred, the device can arbitrate for the bus and then reconnect with the controller to transfer the data.

If the device is very efficient, it will even begin reading the data and place it into a buffer before it reconnects. This allows it to burst the data across the SCSI bus, thereby minimizing the time it needs to use or *own* the bus.

These techniques result in a more efficient use of the available SCSI bus bandwidth. If the controller held onto the bus while waiting for the device to seek, then the other devices would be locked out.

Since, in effect, multiple operations can occur simultaneously, this is also known as overlapped operations or multi-threaded I/O on the SCSI bus. This feature is very important in multi-tasking environments.

If you only have one SCSI device, then disable this feature. You will gain a small amount of performance since there is a slight overhead associated with the disconnect/reconnect sequence.

4.2.3 SCSI-III

ANSI continues to develop the SCSI-II specification to address issues of cable definition, termination, confusing SCSI-II commands, and electrical and signal timing definitions. The SCSI-III architecture encompasses the following commands, interconnects, and protocols:

- The SCSI-III command set consists of five command sets that are derived from SCSI-II command sets and a new SCSI-III command set for RAID controllers.
- The interconnect technologies for SCSI-III specifications are:
 - Fibre Channel Physical and Signaling Interface (FC)
 - IEEE 1394 High Performance Serial Bus (FireWire)
 - SCSI-III Parallel Interface (SPI)
 - Serial Storage Architecture Bus (SSA)
- The SCSI-III protocol standards are:
 - SCSI-III Interlock Protocol (SIP)
 - SCSI-III Serial Storage Protocol (SSP)
 - SCSI-III Serial Bus Protocol (SBP)
 - Fiber Channel Protocol for SCSI (FCP)

Currently, SCSI-III has a transfer rate of 40 MB/s and includes Fast-20 Narrow and Fast-20 Wide.

SCSI-III will be a major step forward in the development of disk subsystems. It further enhances the SCSI-II interface in the following ways:

- Provides three new physical interface layers, SSA, FC and FireWire. These new layers provide better performance, higher availability and more expendability to SCSI.
- Divides SCSI into more than 15 standards, each dealing with a separate part. Because SCSI had become a very large standard, the separation makes the SCSI standard easier to maintain and better to work with. It also allows parts of SCSI-III to be formalized much sooner.

4.2.3.1 Overview of SCSI-III Standards

The breakdown of the SCSI-III standards is shown in Figure 29 and is described in Table 18.




Figure 29. SCSI III Standards Overview

Table 18. Overview of SCSI-III Standards

Standards	Abbrev	Function
SCSI-III Interlocked Protocol	SIP	Describes the protocols used on the SPI bus
Fibre Channel Protocol	FCP	Describes the protocols used on the FC bus
Serial Bus Protocol	SBP	Describes the protocols used on the IEEE 1394 bus
Serial Storage Protocol	SSP	Describes the protocols used on the SSA bus
SCSI-III Architecture Model	SAM	Describes the architecture of the SCSI-III model (includes the SCSI-III device models)
SCSI-III Primary Commands	SPC	Describes the commands that all SCSI devices must implement
SCSI-III Block Commands	SBC	Describes the commands used to transmit <i>blocks</i> of data
SCSI-III Stream Commands	SSC	Describes the commands used to transmit <i>streams</i> of data
SCSI-III Medium Changer Commands	SMC	Describes the commands used to change a medium in a device
SCSI-III Graphic Commands	SGC	Describes the commands that involve graphics
SCSI-III Controller Commands	SCC	Describes the commands used to configure and test the controller of a device
SCSI-III Multimedia Commands	MMC	Describes commands that involve multimedia data, such as audio and video

For clarity, the SCSI-III interconnects are discussed. There are four SCSI-III interconnect technologies:

- SCSI-III Parallel Interface (SPI)
- Fibre Channel Physical and Signaling Interface (FC-PH)
- IEEE 1394 High Performance Serial Bus
- Serial Storage Architecture Bus (SSA-PH)

SPI is as close to SCSI-II technology as we can get because both are parallel technologies. FC-PH, SSA-PH and IEEE 1394 introduce serial data transfer into the SCSI mix.

These three serial SCSI interfaces have two major differences over their parallel sibling:

- The cable has only six wires
- The connectors are simplified into six pins

These new physical properties translate into lower costs for users and manufacturers.

You should take note that SCSI-III Parallel Interface (SPI) is backwardly compatible with SCSI-I and SCSI-II. It is ideal for those users that have sizable investments in SCSI-II equipment because SPI integrates SCSI-II and SCSI-III devices on the same chain. The use of SCSI-III provides much more function due to the improvements in the command set and data transfer rates. The migration from SCSI-II to SPI is relatively simple because all the connectors, cables and terminators essentially remain the same.

The SCSI-III Parallel Interface specification uses the terms *Fast-20 Narrow* and *Fast-20 Wide*. Fast-20 Narrow is an 8-bit bus with a maximum data transfer rate of 20 MB/s. Fast-20 Wide is a 16-bit bus with a maximum data transfer rate of 40 MB/s. The *20* in Fast-20 refers to the clock speed of bus: 20 MHz, which is double the speed of SCSI-II Fast.

The term *Ultra SCSI* is another name for Fast-20 Wide. Ultra SCSI is not equivalent to SCSI-III, but is a subset of the SCSI-III Parallel Interface (SPI).

4.2.4 SCSI Terminology

The terminology of SCSI products can sometimes be confusing. Here are some definitions to help your understanding.

4.2.4.1 General Terminology

The following is a list of general terminology for SCSI.

Fast	Fast refers to the doubling of the data transfer rate from the SCSI 5 MB/s to 10 MB/s by doubling the clock rate. SCSI (that is, the original SCSI specification, or SCSI-I) is 5 MB/s, which is produced by a clock speed of 5 MHz sending data down eight wires. SCSI-II Fast achieves 10 MB/s by doubling the clock speed to 10 MHz.
Wide	Wide is used in reference to the width of the SCSI parallel bus between the adapter and the device. Wide means wider than the original 8-bit path defined in SCSI-I, usually 16-bit. 32-bit transmission is possible within the specification but there are no "Wide-32" devices on the

market. With a 16-bit path, the data rate is double that of an 8-bit device for the same clock speed.

Fast/Wide	Fast/Wide refers to a 16-bit data path running at 10 MHz producing a maximum data transfer rate (or <i>burst</i> rate) of 20 MB/s.
Fast-20	Fast-20 is a bus running at double the clock speed of Fast, or 20 MHz. Fast-20 typically refers to an 8-bit bus and can also be called <i>Fast-20 Narrow</i> . <i>Fast-20 Wide</i> is the 16-bit version, also known as Ultra SCSI.
Ultra SCSI	Ultra SCSI, as described above, is a subset of the SCSI-III specification. It is effectively a Fast SCSI bus running at 20 MHz. Ultra SCSI can produce a maximum transfer of 20 MB/s over an 8-bit data path. Wide Ultra SCSI, the 16-bit version of Ultra SCSI, can transmit a maximum of 40 MB/s.
Ultra2 SCSI	Ultra-2 is the latest extension to the SCSI-III specification that allows transfers to take place at 80 MB/s on a wide bus, or 40 MB/s for narrow. At twice the speed of the current Ultra SCSI, and comparable with SSA, it has the advantage of complete backwards compatibility with existing peripherals. The cable length maximum, as defined by the standard, is 12 meters. This standard requires Low Voltage Differential (LVD).

Note

Wide refers to the width of the bus between the SCSI adapter and its attached devices. Do not confuse this with the width of the host bus interface (for example, a 32-bit PCI or Micro Channel bus).

4.2.4.2 Asynchronous versus Synchronous

An asynchronous device must acknowledge each byte as it comes from the controller. Synchronous devices may transfer data in bursts, and the acknowledgments happen after the fact. Synchronous is much faster than asynchronous, and most newer devices support the synchronous mode of operation. The adapters negotiate with devices on the SCSI bus to ensure that the mode and data transfer rates are acceptable to both the host adapter and to the devices. This process prevents data from being lost and ensures that data transmission is error free.

4.2.4.3 SCSI Differential

Normally, there is one wire in a SCSI cable for each signal. However, over long distances and with high clocking rates, the signals can degrade and errors can occur. To solve this, SCSI Differential was developed, which uses two wires for each signal.

Consider Figure 30 on page 104 as an example. Signal A is transmitted along two wires, A+ and A-. A- is the mirror image of A+. When some line noise appears such as a signal spike, the spike will be of the same polarity on both A+ and A-. On the receiving end, the signals are subtracted, and the result is A+/A- as shown in the diagram. The spike is canceled out by the subtraction.




Figure 30. SCSI Differential Example

With SCSI Differential, longer cable lengths, up to 25 meters, can be maintained. The trade-off is the additional cost of the cables, connectors and the circuitry needed to perform the calculations. SCSI Differential is not commonly used for long distances. However, SCSI repeaters, as described in the below section, are more common.

4.2.5 SCSI Repeaters

A SCSI repeater is a device that enhances the signal quality of a SCSI bus and allows the bus to be physically longer. The repeater can be in the form of either an adapter or an external "black box."

As noted in Table 19 on page 105, the maximum length of the SCSI bus for a SCSI-II F/W configuration is 3.0 meters. With a repeater, it is possible to extend this distance.

4.2.6 Summary of SCSI Specifications

SCSI defines many different modes of operation, including several different data transfer modes. The most common are shown in Table 19.

Table 19. Differential SCSI Modes

Mode	Bandwidth	Frequency	Max. data Transfer Rate	Max. Cable Length Single Ended
SCSI-I	8-bit	5 MHz	5 MB/s	6 meters
Fast SCSI	8-bit	10 MHz	10 MB/s	3 meters
Fast/Wide SCSI	16-bit	10 MHz	20 MB/s	3 meters
Ultra SCSI (or Fast-20 SCSI)	8-bit	20 MHz	20 MB/s	1.5 meters
Wide Ultra (or FAST-20 Wide SCSI)	16-bit	20 MHz	40 MB/s	1.5 meters
Ultra2 SCSI	16-bit	40 MHz	80 MB/s	12 meters

Note

The single-ended length refers to the length of the SCSI cable without using devices and their main characteristics.

4.3 SCSI Cabling

It is very likely that you will face problems regarding configuration and installation of SCSI adapters and devices.

SCSI cabling can be one of the most confusing topics to encounter. Some of the questions that may arise are:

- How many SCSI devices can be attached to a single SCSI adapter?
- What are the right cable features to attach more SCSI devices on the same SCSI chain?
- What is the right SCSI cabling for an HACMP configuration?

In this section, the main SCSI cabling features used with PCI-based RS/6000 systems are described, answering some of the questions mentioned above.

The following describes the SCSI cabling of the two PCI SCSI adapters supported by the Model 150 and 260 entry workstation and workgroup servers:

- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI Differential Ultra SCSI Adapter (# 6207)

Note

Refer to the publication *IBM RS/6000 Adapters, Devices and Cables information for Multiple Bus Systems*, SA38-0516, for specific SCSI cabling information.

4.3.1 General SCSI Cabling Considerations

Before describing SCSI cabling features, it is necessary to understand the following primary specifications of SCSI technology:

- Default SCSI cable routing
- Maximum SCSI bus length
- SCSI terminators
- SCSI device addresses
- SCSI bus width

4.3.1.1 Default SCSI Cable Routing

SCSI IDs for the storage devices are defined by the location in which they are installed. The standard SCSI cable is used for all combinations. SCSI termination is done on the system planar. Termination setting of all internal SCSI devices must be set to OFF.

4.3.1.2 SCSI Bus Length

SCSI bus length is defined as the distance between terminators at either end of a SCSI bus. The SCSI specifications allow the following maximum bus length for each SCSI bus:

- SCSI-II Single Ended: 6 meters (20 feet)
- SCSI-II Fast/Wide Single Ended: 3 meters (10 feet)
- SCSI-II Fast/Wide Differential: 25 meters (80 feet)

For configurations using both internal and external cabling, length restrictions apply to the length from the end of the internal cable to the terminator on the last device on the external bus.

Devices that have two SCSI connectors have internal cabling that must be included when calculating the total cable length. When attaching these devices to a SCSI chain, connect one cable to one connector and the other cable to the other connector. Do not piggy back the second cable/terminator onto the first as you would with a device that had only one SCSI connector as shown in Figure 31.

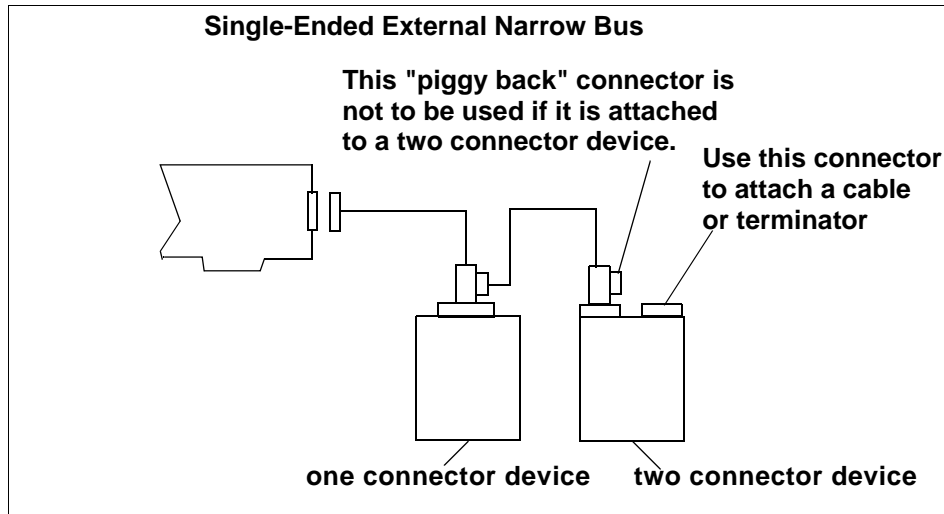


Figure 31. Piggy Back Connectors

A dedicated adapter should be used for attachment of any external enclosure containing multiple SCSI devices.

4.3.1.3 SCSI Terminators

There must be exactly two terminators on the SCSI bus, and they must be located at each end of the bus.

- For SCSI adapters with external devices only, make sure the appropriate SCSI terminator is connected to the last device on the bus.
- For SCSI adapters with internal devices only, make sure that the appropriate SCSI terminator is connected to the end of the internal SCSI cable.
- For SCSI adapter with both internal and external devices, make sure the appropriate SCSI terminator is connected to the last device on the external bus and that the end of the internal SCSI cable has been properly terminated.

4.3.1.4 SCSI Device Addresses

The Ultra SCSI adapters support a maximum of 16 SCSI addresses, including devices and adapters. The default address for an adapter is 7, but it can easily be changed if necessary.

- All devices on the SCSI bus must have a unique SCSI ID.
- The SCSI bus address determines priority on the bus. Address priority from highest to lowest is as follows:
 - The addresses 7, 6, 5, 4, 3, 2, 1, 0 and 15, 14, 13, 12, 11, 10, 9, 8, with addresses 15 through 8 being used in 16-bit configurations only.
 - Generally, the highest priority is assigned to the adapter.
- For any single adapter in an 8-bit configuration, a maximum of seven devices are permitted, provided that the supported configuration-specific bus lengths are not exceeded.
- For any single adapter in a 16-bit configuration, a maximum of 15 devices are permitted, provided that the supported configuration-specific bus lengths are not exceeded.
- The default ID of the SCSI adapter in a single adapter configuration is 7. All devices on that bus must have a unique ID from 0 to 6 (8 to 15 are also valid if SCSI-wide); two different devices may not have the same SCSI ID. In the high-availability configurations, the second adapter must have its address changed to avoid conflicts.

Note

The SCSI address switch for each device must be set while power to the system unit is off. If the SCSI device is an external type, the external device must be powered off only while switching the address. The operating system determines the system configuration during IPL (or during execution of `cfgmgr`).

- If a SCSI address is changed after the operating system is loaded, the operating system must be stopped and loaded again to have the correct configuration, or `cfgmgr` run.
- Stand-alone diagnostics always default to a SCSI ID of 7 when testing SCSI adapters and devices. Choosing SCSI IDs other than 7 for both adapters prevents problems when using stand-alone diagnostics on systems in HA clusters or in multi-initiator configurations.
- Check the documentation for your specific SCSI subsystem to insure that there are no SCSI ID conflicts if the adapters are addressed at ID(s) other than 6 or 7.

4.3.1.5 SCSI Bus Width

Operation of both 8-bit and 16-bit devices on the same external SCSI bus is not supported due to the termination and cabling restrictions.

Operation of both 8-bit and 16-bit devices on the same internal bus is supported concurrently as long as 16-bit internal cable and 68-pin to 50-pin interposers are used.

Mixed-width internal attachment is supported concurrently with single-width external attachment as long as maximum cable length restrictions are not exceeded.

4.3.2 Cabling the PCI Single-Ended Ultra SCSI Adapter

The maximum supported cable length for the PCI single-ended Ultra SCSI adapter (# 6206) depends on what type of devices are attached and where they are attached (to the internal or external connector). Device types are classified as:

- Ultra SCSI SE – maximum transfer rate of 40 MB/s (one byte transfers)
- Ultra SCSI Wide – maximum transfer rate of 40 MB/s (two byte transfers)

For this adapter:

- The maximum supported cable length for configurations without any SCSI-II Fast or Ultra SCSI devices is 6 meters (approximately 20 feet).
- The maximum supported cable length for configurations that include SCSI-II Fast (but not Ultra) is 3 meters (approximately 10 feet) with the exception of the 7027-HSC High Capacity Drawer, which can be attached with up to 6 meters of cable.
- To ensure optimum signal quality for Ultra SCSI transfers, attachment of multiple Ultra SCSI devices is only recommended for devices mounted inside the system unit. To ensure optimum signal quality for Ultra SCSI transfers, it is recommended that only Ultra SCSI devices be attached to a backplane that is driven by a PCI Single-Ended Ultra SCSI adapter.

This default setting can be changed (using SMIT or the `chdev` command) to allow attachment of external Ultra SCSI devices, with the restriction that there are no SCSI devices attached to the internal connector.

The Ultra SCSI Adapter (# 6206) has the following supported configurations:

- Internal Ultra devices running at Ultra speeds:
 - Up to six Ultra devices attached to the internal port (dependent on internal configuration and cabling)

- No external attachments are allowed
- External Ultra devices running at Ultra speeds:
 - No internal attachments allowed.
 - Up to four external Ultra wide (16-bit) devices can be attached to the external port. Maximum cable length must not exceed 3 meters.
 - Up to three external Ultra (8-bit) devices can be attached to the external port. Maximum cable length must not exceed 3 meters.
 - Ultra or SCSI-II Fast devices running at SCSI-II Fast speeds,
 - Up to six devices attached to the internal port (dependent on internal system configuration and cabling).
 - External attachment of up to four independent physical enclosures is allowed, provided each physical enclosure presents only one load to the SCSI bus. The total bus length must not exceed 3 meters. Total bus length includes internal and external cable length.
- Multiple SCSI-II Fast devices in external enclosures:
 - No internal attachments allowed.
 - Maximum combined internal (to enclosure) and external cable length is 3 meters.
 - Loads on the cable (cable length between devices). They must be 0.1 meters apart at a minimum.
 - No mixing of bus widths (8-bit and 16-bit) unless the 68-pin-to-50-pin interposer (PN 92F2565 or equivalent) is used.

The following tables describe the cables and terminator features, part numbers, and lengths for the PCI Single-Ended Ultra SCSI Adapter (# 6206).

Cables and Terminators for Single-Ended SCSI Adapters

In the following tables, the cables and terminators of the following single-ended SCSI adapters are listed:

- PCI Single-Ended Ultra SCSI Adapter (# 6206)

Adapter-to-First Device Cables

Table 20 on page 111 describes the cables for the Single-Ended Ultra SCSI Adapter for connection to the first device.

Table 20. Single-Ended Ultra SCSI Adapter-to-First Device Cables

M/T	FC	Part Number	Length (meters)	Cable Description
Host System	2111	06H6037	1.0	Adapter-to-first device (where first device has two connectors), 8-bit narrow bus
Host System	2112	06H6037	1.0	Adapter-to-first device (where first device has two connectors), 8-bit narrow bus
Host System	2113	52G0174	1.5	Adapter-to-first device (where first device has one connector), 8-bit narrow bus
Host System	2114	52G0173	.94	16-bit Y-cable
Host System	2115	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus
Host System	2116	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus

Note: When cables are ordered by feature code (FC), the appropriate terminator is included with the order. When cables are ordered by part number, only the cable is included.

The external connector on these adapters are the SCSI-III standard, 68-pin P cable connector. Many of the 16-bit SCSI devices also use this connector type, and as a result, some cables can be used as either adapter-to-first device or device-to-device cables, depending upon what type of SCSI connectors are present on the devices.

Device-to-Device Cables

Table 21 describes the device-to-device cables for single-ended applications.

Table 21. Device-to-Device Cables for Single-Ended Installations

Machine Type	FC	Part Number	Length (meters)	Cable Description
SE External Device	2840	33F4607	0.7	Device-to-Device (where second device has two connectors), 8-bit narrow bus
SE External Device	3130	31F4222	0.66	Device-to-Device (where second device has one connector), 8-bit narrow bus
SE External Device	2860 /9139	52G9921	0.3	Device-to-Device (where second device has two connectors), 16-bit wide bus
SE External Device	2884 /9160	52G4291	0.6	Device-to-Device (where second device has two connectors), 16-bit wide bus
SE External Device	2883 /9150	52G4233	2.5	Device-to-Device (where second device has two connectors), 16-bit wide bus
7027 HST	2425			
7027 HST	3132	40H7351	6.0	Device-to-Device (where second device has two connectors), 16-bit wide bus

Note: Most feature codes for cables are only orderable against the attachment device (7204, 7206, 7208, and so on.). For some cables, the feature codes have been made available on the system units. In these cases, the system feature code will be listed. Otherwise, the attachment device feature codes will be used.

System-To-System Cables

Table 22 on page 113 provides a 16-bit connection between any two differential or single ended SCSI devices having 68-pin connectors. It can be

used to attach an external SCSI device to a SCSI adapter card in an RS/6000 system.

Table 22. 16-Bit SCSI-II System-To-System Cable

Machine Type	FC	Length (meters)	Cable Description
SE External Devices	2424	0.6	16-bit SCSI-II system-to-system cable
SE External Devices	2425	2.5	16-bit SCSI-II system-to-system cable

Terminators for Use with These Adapters

Table 23 describes the terminators for single-ended installations.

Table 23. Terminators for Single-Ended Installations

M/T	FC	Part Number	Connector	Terminator Description
SE External Devices	Part of cable FC	52G4260	50-pin low density	8-bit external FPT18C terminator
SE External Devices	Part of cable FC	92F0432 (52G9907)	68-pin high density	16-bit external Boulay terminator
SE External Devices	Part of cable FC	92F0322 (92G2566)	68-pin high density	16-bit external bus terminator

4.3.3 Cabling the PCI Differential-Ended Ultra SCSI Adapter

To understand the cabling for the PCI differential-ended Ultra SCSI adapter (# 6207), first read 4.3.1, "General SCSI Cabling Considerations" on page 106 and then read the following for specific information.

SCSI Differential Cable Lengths Using This Adapter

The maximum supported cable length for configurations is 25 meters (approximately 80 feet).

Adapter-to-First Device Cables

Table 24 provides a list of adapter to first device cables.

Table 24. Differential-Ended Ultra SCSI Adapter-to-First Device Cables

M/T	FC	Part Number	Length (Meters)	Cable Description
Host System	2112	06H6037	1.0	Adapter-to-first device (where first device has two connectors), 8-bit narrow bus
Host System	2114	52G0173	0.94	16-bit Y-cable
Host System	2116	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus

Note: When cables are ordered by feature code, the appropriate terminator is included with the order. When cables are ordered by part number, only the cable is included. For terminator part numbers, refer to the Table 26 on page 116.

For this adapter, the same cable can be used for either single-ended or differential attachments. The difference in feature code orders is the terminator type.

The external connector on this adapter is the SCSI-III standard, 68-pin "P" cable connector. Many of the 16-bit SCSI devices also use this connector type, and, as a result, some cables can be used as either adapter-to-first devices or device-to-device cables, depending upon what type of SCSI connectors are present on the devices.

Device-to-Device Cables

Table 25 provides a list of device-to-device cables.

Table 25. Device-to-Device Cables for Differential-Ended Installations

M/T	FC	Length (Meters)	Part Number	Cable Description
DE External Device	2848 /9134	0.6	74G8511	Device-to-device (where second device has two connectors), 8-bit narrow bus
DE External Device	2860 /9139	0.3	52G9921	Device-to-device (where second device has two connectors), 16-bit wide bus

M/T	FC	Length (Meters)	Part Number	Cable Description
DE External Device	2884 /9160	0.6	52G4291	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2846 /9132	2.5	52G4233	Device-to-device (where second device has two connectors), 16-bit wide bus
7027	2425			
DE External Device	2885 /9161	4.5	88G5749	Device-to-device (where second device has two connectors), 16-bit wide bus
7027-HSD	3132	6.0	40H7351	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2870 /9146	12.0	88G5747	Device-to-device (where second device has two connectors), 16-bit wide bus
7027 HSD	3125			
DE External Device	2869 /9145	14.0	88G5748	Device-to-device (where second device has two connectors), 16-bit wide bus
DE External Device	2868 /9144	18.0	88G5746	Device-to-device (where second device has two connectors), 16-bit wide bus
7027 HSD	3136			

Note: Most feature codes for cables are only orderable against the attachment device (7204, 7206, 7208). For some cables, the feature codes have been made available on the system units. In these cases the system feature code is listed otherwise the attachment device feature code is used.

9xxx feature codes are used for new build orders, whereas 2xxx feature codes are used for MES orders.

Terminators for Use With This Adapter

This adapter has on-card SCSI terminators that must be removed before the adapter can be used in a high-availability configuration. The high-availability configuration is implemented by removing the three on-card differential terminating resistors (labeled RN1, RN2, and RN3) on the adapter, then attaching the middle leg connector of the high-availability configuration

Y-cable to the adapters 68-pin external connector. The remaining two legs of the Y-cable are used to attach other systems and devices to the SCSI bus.

If the PCI Differential Ultra SCSI adapter is at the end of the SCSI bus, the shorter leg of the Y-cable must be terminated with the appropriate terminator.

Note: The high-availability configuration (Y-cable with a terminator on the shorter leg) allows disconnection of the adapter from a *live* SCSI bus by removal of the external bus connection (the middle leg of the Y-cable). Although termination and SCSI bus continuity is maintained during removal of the adapter, the noise generated may create undetected data errors if the bus is in use during the time of removal. To maintain data integrity, the SCSI bus should be inactive during the removal of adapters, cables, or terminators.

Table 26 provides a list of required terminators.

Table 26. Terminator For Description-Ended Installations

M/T	FC	Part Number	Connector	Terminator Description
DE External Devices	Part of cable FC	87G1356	50-pin low density	8-bit external bus terminator
DE External Devices	Part of cable FC	61G8324	68-pin high density	16-bit external bus terminator

Cabling Examples for the PCI Differential Ultra SCSI Adapter

Figure 32 and Figure 33 shows how to cable the PCI Differential Ultra SCSI Adapter.




Figure 32. Differential External Narrow Bus




Figure 33. Differential External Wide Bus

Table 27 provides a list of available cables for differential adapters.

Table 27. Cabling for the PCI Differential Ultra SCSI Adapter

Item Number	Part Number	Length (Meters)	Description
1	06H6037	1.0	Adapter-to-dual-connector device (narrow 8-bit)
2	06H6036	1.0	Adapter-to-dual-connector device (wide 16-bit)
3	74G8511	0.6	Device-to-dual-connector device (narrow 8-bit)
4	52G4291	0.6	Device-to-dual-connector device (wide 16-bit)
	52G9921	0.3	
5	87G1356		Terminator (8-bit)
6	61G8324		Terminator (16-bit)

4.4 Serial Storage Architecture Overview

Serial Storage Architecture (SSA) is a high-speed serial interface that IBM originally developed as a proprietary interface. The interface was standardized in 1994 under the control of the ANSI X3T10.1 committee. SSA is designed to overcome the bottlenecks of existing SCSI bus architectures.

SSA is a single I/O interface that can effectively address the storage requirements of a variety of computers, from PCs to supercomputers. Combining the highest performance of a direct disk interface with low cost connectivity and flexibility of network-wide attachment, SSA offers a reliable and robust storage interface. Compared to newer interfaces such as Fibre Channel-Arbitrated Loop (FC-AL) connection, which IBM announced in October 1998 for the RS/6000 Enterprise Server model S7A, SSA is more cost-effective while delivering equal or better performance. Compared to today's parallel interfaces such as SCSI, SSA serial link connectivity provides many benefits:

- Enables hot-swapping and automatic configuration of new devices (without interrupting file service to users)
- Supports up to 128 total devices (including multiple initiator) in a single SSA configuration, compared to a maximum of 16 devices for SCSI
- Provides configuration flexibility with up to 25 meters of separation between devices
- Maximizes data transfer, currently enabling 80 MB/s of available bandwidth

SSA uses shielded twisted pair (STP) cabling as a transport medium. There are five conductors in the cable. Four are used for signal transmission, and one is used as a shield. The twists in the cable eliminate crosstalk, and the shielding eliminates electromagnetic interference (EMI). Thus, SSA can function reliably with its nodes placed 25 m apart. IBM offers a fiber optic extender that can extend this limit to 2400 m. The fiber optic extenders are supplied in pairs to complete a string.

4.5 Comparison between SCSI and SSA Architecture

Table 28 on page 120 lists the common features found in SCSI and SSA devices for comparison and Figure 34 shows how the subsystems differ.

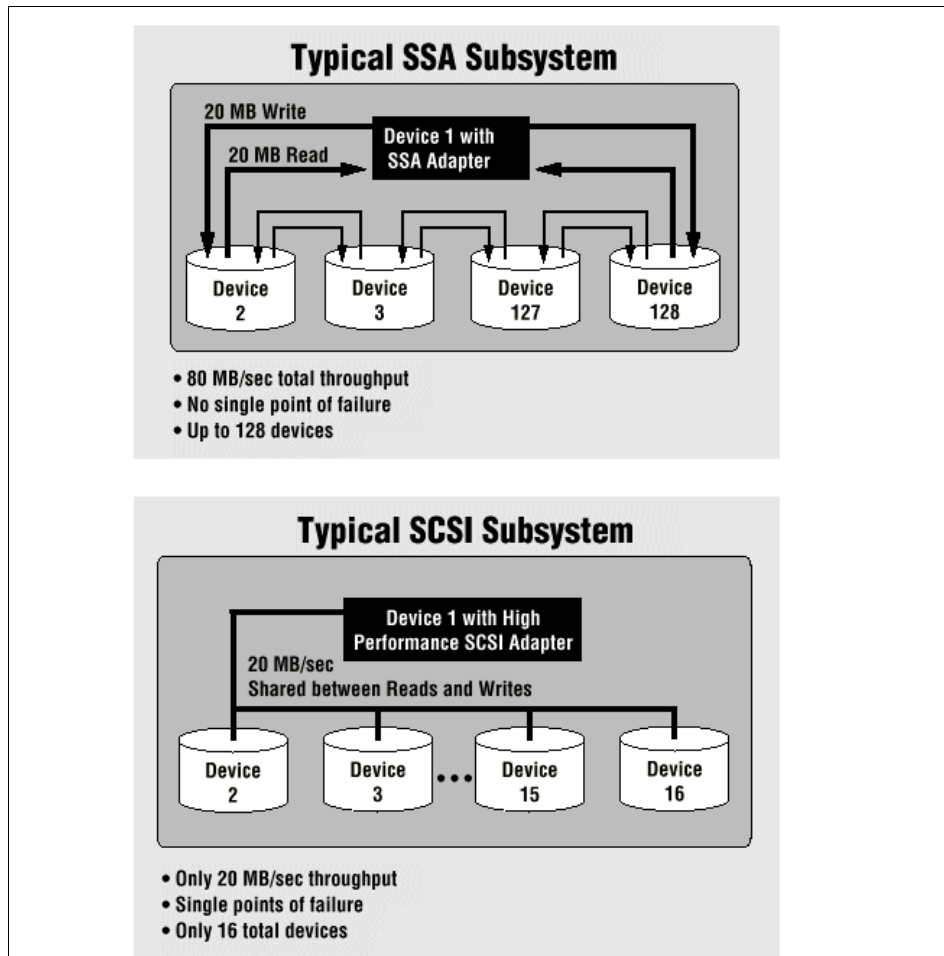


Figure 34. Comparison between SCSI and SSA Subsystem

Table 28. Comparing SSA and SCSI

Feature	SCSI	SSA
Cable Type	68 conductor shielded cable (SCSI-II Differential Fast/Wide)	4 conductor twisted-pair
Maximum Cable Lengths	25 Meters Total	25 meters between devices, Up to 2500 Meters between devices with the use of a Fiber Optic Extender
Cable Costs	Significant	Low
Individual drive addressing	Must be set manually	Dynamic
Cable Termination	Electrical terminator must be at the end of the SCSI-bus	No termination required, closed serial loops
Fault Tolerance	If single data path (bus) is compromised, access is lost	No single point of failure. Multiple data paths prevents a single loop from causing loss of access to information
Maximum Throughput	Up to 40 MB/s	Up to 160 MB/s

4.6 Internal SCSI Devices

The following section gives an overview of the internal SCSI devices for the Models 150 and 260, and their main characteristics:

- Disk drives
- CD-ROM drives
- Tape drives

4.6.1 Disk Drives for the Model 150 and Model 260

Table 29 on page 121 shows the main characteristics of the internal disk drives that are available for the Models 150 and 260. When you want to install a 4.5 GB (# 2900) or a 9.1 GB (# 2908) hard disk in bay 3 you have to use the Media Bay Disk Drive Mounting Kit (# 6509).

If you want to install the disk drives in bay 4 or 5, you have to use either the 4.5 GB (# 3028) or the 9.1 GB (# 3029) hard disk.

4.6.1.1 4.5 GB Ultra SCSI 16-Bit Disk Drive (# 2900)

The 4.5 GB Ultra SCSI 16-bit disk drive is a high performance disk drive that provides 4.5 GB capacity that supports the Ultra SCSI interface speed of 40 MB/s. To operate at 40 MB/s, Ultra-SCSI adapter (# 6206) is required.

4.6.1.2 4.5 GB Ultra SCSI 16-Bit Enhanced Disk Drive (# 3028)

The 4.5 GB Ultra SCSI 16-bit Enhanced Disk Drive is a high-performance disk drive that provides 4.5 GB capacity that supports the Ultra SCSI interface speed of 40 MB/s. It requires Ultra SCSI adapter (# 6206) to operate at 40 MB/s (only available for Model 260).

4.6.1.3 9.1 GB Ultra-SCSI 16-Bit 1-Inch High Disk Drive (# 2908)

The 9.1 GB Ultra SCSI 16-Bit 1-inch (25 mm) high disk drive offers high performance that provides 9.1 GB of storage capacity and supports the industry Ultra SCSI interface speed of 40 MB/s. To operate at 40 MB/s, Ultra SCSI adapter (# 6206) is required.

4.6.1.4 9.1GB Ultra SCSI 16-Bit Enhanced Disk Drive (# 3029)

The 9.1 GB Ultra SCSI 16-bit 1-inch (25 mm) High Disk Drive offers high-performance that provides 9.1 GB of storage capacity and supports the industry Ultra -SCSI interface speed of 40 MB/s. It requires an Ultra SCSI adapter (# 6206) to operate at 40 MB/s (only available for Model 260).

Table 29. Internal Disk Drives for the Models 150 and 260

Capacity	4.5 GB	9.1 GB
Feature Code	# 2900 or # 3028	# 2908 or # 3029
Maximum Media Data Transfer Rate	13.7 MB/s	22.4 MB/s
Average Seek Time	7.9 ms	6.7 ms
Rotational Speed	7200 RPM	7200 RPM
Average Latency	4.17 ms	4.17 ms
SCSI Interface	Ultra SCSI	Ultra SCSI
Data Width	16-bit, single-ended	16-bit, single-ended

4.6.2 Internal CD-ROM Drive

Table 30 on page 122 provides an overview of the supported CD-ROM drives for the Models 150 and 260. The 32X max SCSI-II Internal CD-ROM drive is a tray loading CD-ROM drive providing up to 4800 KB/sec maximum media data transfer rate. It is a 5.25 inch half-high form factor, single-ended, 8-bit, multi-session capable, CD-ROM drive that provides state-of-the-art

performance and supports existing 600 MB CD-ROM discs. It will operate in either the vertical or horizontal positions.

Table 30. Internal CD-ROM Drives for Models 150 and 260

CD-ROM Speed	32X
Feature Code	# 2624
Interface	SCSI-II, 8-bit single-ended
Interface Speed (average)	4800 KB/s
Average Random Access Time	90 ms
Buffer memory	512 KB
Bays	1 half-height

4.6.3 Internal Tape Drives

Table 31 and Table 32 show detailed technical information about the available internal tape drives for the Models 150 and 260.

Table 31. Internal Tape Drive for Model 150

Product Name	12/24 GB 4 mm Internal Tape Drive
Feature Code	# 6159
Tape Capacity (Native Mode)	12 GB
Tape Capacity (Compression Mode)	24 GB
Type of Cartridge	IBM 4 mm DDS-3 Data Cartridge
Data Transfer Rate	1.1 MB/s Native Mode, 2.2 MB/s Compression Mode
Interface	SCSI-II (Single-Ended)
Form Factor	5.25-Inch Half-High
Compatibility	2 GB Mode (R/W) 4 GB Compression (R/W) 4 GB Mode (R/W) 8 GB Compression (R/W) 12 GB Mode (R/W) 24 GB Compression (R/W)

Table 32. Internal Tape Drives for Model 260

Product Name	12/24 GB 4 mm Internal Tape Drive	20/40 GB 8 mm Internal Tape Drive
Feature Code	# 6159	# 6156
Tape Capacity (Native Mode)	12 GB	20 GB
Tape Capacity (Compression Mode)	24 GB	40 GB
Type of Cartridge	IBM 4 mm DDS-3 data cartridge	IBM 8 mm Data Cartridge
Data Transfer Rate	1.1 MB/s Native Mode, 2.2 MB/s Compression Mode	3 MB/s Native Mode, 6 MB/s Compression Mode
Interface	SCSI-II (Single-Ended)	SCSI-II 16-bit (Single-Ended)
Form Factor	5.25 Inch Half-High	5.25 Inch Half-High
Compatibility	2 GB Mode (R/W) 4 GB Compression (R/W) 4 GB Mode (R/W) 8 GB Compression (R/W) 12 GB Mode (R/W) 24 GB Compression (R/W)	2.3 GB Mode (R/O) 2.3 GB Mode Compression (R/O) 5 GB Mode (R/O) 5 GB Compression (R/O) 7 GB Mode (R/O) 7 GB Compression (R/O) 20 GB Mode (R/W) 20 GB Compression (R/W)

4.7 Internal SSA devices

There are no internal SSA storage devices available for the Model 150 and 260. However, you can attach external SSA devices through the PCI SSA Multi-Initiator/RAID EL adapter (# 6215).

For more information about this adapter see 4.8.2, "SSA Adapter" on page 124.

4.8 PCI Storage Adapters

IBM RS/6000 43P Models 150 and 260 support the following PCI-based storage SCSI and SSA adapters:

- PCI Single-Ended Ultra SCSI adapter (# 6206)
- PCI Differential Ultra SCSI adapter (# 6207)
- PCI SSA Multi-Initiator/RAID EL Adapter (# 6215) with SSA Fast-Write Cache option card (# 6222)

4.8.1 SCSI Adapters

This section describes the available SCSI adapters for the Models 150 and 260.

4.8.1.1 PCI Single-Ended Ultra SCSI Adapter (# 6206)

The PCI Single-Ended Ultra SCSI adapter is an ideal solution for applications requiring large block data transfers (> 64 KB block size) in multi-disk drive environments utilizing SCSI-II protocol. This adapter has a maximum data transfer rate of 40 MB/s which is two times the maximum data transfer rate of Fast/Wide adapters. This adapter requires one PCI slot.

This adapter (# 6206) will automatically sense the presence of an external device. The default mode of operation is for the adapter to limit negotiations to fast (10 MB/s for 8-bit, 20 MB/s for 16-bit) operation whenever there is an external attached.

4.8.1.2 PCI Differential Ultra SCSI Adapter (# 6207)

The PCI Differential Ultra SCSI adapter is the next generation of SCSI-II performance with a maximum data transfer rate of 40 MB/s. This adapter allows connection to external differential SCSI-II F/W of Ultra SCSI type devices up to 25 meters away. This adapter will negotiate with each external device and transfer data at the fastest SCSI data transfer rate capable by the external device. Some external Differential Ultra SCSI devices include the 7137 Disk Array subsystem model and certain external disk drives models of the 7204 family. The adapter conforms to SCSI-II standard and the Fast-20 (Ultra) documentation.

4.8.2 SSA Adapter

The PCI SSA Multi-Initiator/RAID EL Adapter (# 6215) is supported in the Model 150 and 260. It is also supported by other PCI models in the IBM RS/6000 family, including Models E30, F40, F50, H50, and S70.

4.8.2.1 PCI SSA Multi-Initiator/RAID EL Adapter

The PCI SSA Multi-Initiator/RAID EL Adapter (# 6215) can be configured as either a two-initiator non-RAID, or one-initiator RAID adapter. It also has been enabled to support an optional 4 MB Fast-Write Cache card that improves write performance in both RAID 5 or non-RAID (one-initiator configurations) operations. The PCI SSA Multi-Initiator/RAID EL Adapter has four ports, two SSA loops, and a supported configuration of 48 SSA disk drives per loop, or 96 SSA disk drives. The adapter supports target mode operations and HACMP.

The PCI SSA Multi-Initiator/RAID EL Adapter, when operated in a RAID 5 configuration, will support (2+P) to (32+P) arrays and up to six (15+P) arrays. The adapter also supports Hot Spares in RAID 5 mode. All models of the 7133 are supported, as well as the 7131-405 and internal RS/6000 SSA disk drive configurations. The SSA Fiber optic Extender is also supported. It requires one PCI adapter slot. Only AIX 4.2.1 and later version of AIX support this adapter.

The SSA Fast-Write Cache Option Card (# 6222) is a 4 MB fast-write optional feature that will plug into either SSA Multi-Initiator/RAID EL Adapter. It utilizes non-volatile RAM and has a data memory retention of over seven years. During the unlikely event of an SSA Multi-Initiator/RAID EL Adapter failure, a new SSA Multi-Initiator/RAID EL Adapter can be installed and the SSA Fast-Write Cache Option Card can be removed from the failing adapter and installed in the new SSA Multi-Initiator/RAID EL Adapter insuring data integrity and continuous operation.

The SSA Fast-Write Cache Option Card (# 6222) can provide as much as 10-times improvement of data throughput and response time during certain sequence write operations compared to SSA RAID adapters without the SSA Fast-Write Cache. The response time and data transfer improvement using the option card (# 6222) will vary depending upon the data block sizes, the percentage of sequential writes, machine type/model, and application parameters.

The SSA Fast-Write Cache Option Card (# 6222) plugged into SSA Multi-Initiator/RAID EL Adapter (# 6215) will operate in either non-RAID or RAID 5 mode, in a one-initiator configuration.

4.8.2.2 SSA Device Drivers

Two types of device drivers provide support for all SSA subsystems:

- The SSA adapter device driver, which deals with the SSA adapter.

- The SSA head device drivers, which deal with devices that are attached to the SSA adapter. The SSA disk device driver is an example of an SSA head device driver.

For subsystems that use PCI SSA Multi-Initiator/RAID EL Adapters or Micro Channel Enhanced SSA Multi-Initiator/RAID EL Adapters, the Target-Mode SSA (TMSSA) device driver is also available. This device driver provides support for communications from using system to using system.

4.8.3 Storage Adapter Placement Guidelines for Model 150 and 260

Table 33 provides information about the placement of the PCI storage adapters in the Models 150 and 260.

Table 33. Placement of PCI Storage Adapters into the Models 150 and 260

Adapter	Slot Usage Model 150	Slot Usage Model 260	System Maximum
PCI Single-Ended Ultra SCSI (# 6206)	Any slot	Any slot	2
PCI Differential Ultra SCSI (# 6207)	Any slot	Any slot	2
PCI SSA Multi-Initiator/RAID EL (# 6215)	2P, 3P	Any slot	1

4.9 External SCSI and SSA Storage Devices

There many external SCSI and SSA storage devices available which can be attached to the Model 150 and 260.

See also 3.1.3.14, "Supported Devices" on page 63 for details about the Model 150 and 3.2.4.15, "Supported Devices" on page 89 for details about the Model 260 supported devices.

Further information can be found on the Web at:

<http://www.storage.ibm.com>

Chapter 5. Adapters, Overview and Configuration

This chapter introduces the major communication and graphics adapters that are used with the RS/6000 43P 7043 Model 150 and Model 260. Following the introduction, a discussion of PCI adapter configuration is given.

For a complete summary of the adapters supported on the Model 150 and the Model 260, refer to Table 37 on page 152.

5.1 Communication Adapters

In this section, an introduction to various communication adapters is provided including:

- Token-Ring
- Ethernet
- ISDN
- FDDI
- ATM
- Multiple Port
- WAN
- Miscellaneous

5.1.1 Token-Ring Adapter

In this section, information on the available token-ring adapter is given.

5.1.1.1 PCI IBM Token-Ring Adapter (# 2920)

The IBM PCI Token-Ring Adapter is a high performance 32-bit bus master card designed to operate on selected RS/6000 models that incorporate the PCI bus interface. The following is a list of attributes for this card:

- Enhanced system interface for higher throughput and lower system utilization, coupled with RJ-45 and DB9 connectors
- Handles LAN environment requirements of today's servers and high-end workstations running I/O intensive applications on the network
- Operates at either 16 Mbps or 4 Mbps over unshielded twisted pair (UTP) or shielded twisted pair (STP)
- Automatic ring-speed detection
- External status LEDs providing visual indication of adapter and ring status

- Full duplex LAN operation
- Supports system boot over token-ring

5.1.2 Ethernet Adapters

In this section, information on the available Ethernet adapters is provided.

5.1.2.1 PCI IBM 10/100 Mbps Ethernet Adapter (# 2968)

The IBM 10/100 Mbps Ethernet PCI adapter provides a 10 Mbps (10BaseT) or 100 Mbps (100BaseTx) full-duplex Ethernet LAN connection. The adapter supports 32-bit bus master PCI connection to the host system for enhanced data throughput and lower bus utilization. The adapter supports Category-5 unshielded twisted pair cabling.

5.1.2.2 PCI Gigabit Ethernet - SX Adapter (# 2969)

The Gigabit Ethernet - SX PCI Adapter provides a 1 Gbps (1000BaseT) full-duplex Ethernet LAN connection. With a ten-time increase in performance over Fast Ethernet, Gigabit Ethernet enables customers to move data to and from their servers with ease, at only a fraction of the cost. This adapter provides a full-duplex 1000Base-SX fiber connection to a Gigabit Ethernet switch.

Note

AIX 4.3.2 or later is required for hardware checksum and jumbo frame support.

5.1.3 ISDN Adapters

In this section, information on the available ISDN adapter is given.

5.1.3.1 PCI Eicon ISDN DIVA PRO 2.0 S/T Adapter (# 2708)

The Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter connects PCI RS/6000 systems to worldwide ISDN networks. The adapter comes from Eicon Technology, the world leader in ISDN technology, as a fully supported IBM feature. The Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter for the PowerPC Systems provides basic rate 2B + D with 64 Kbps on each B channel. It requires connection to NT-1 (Network Terminator -1) from the phone line provider. TCP/IP applications are supported.

Note

AIX drivers are shipped on diskette with the hardware feature. Systems ordered with AIX preinstalled will have the device drivers installed on the hard disk drive.

5.1.4 FDDI Adapters

In this section, information on the available FDDI adapters is given.

Note

The sum of the following adapters may be used in any combination, but must not exceed four (4) total per system: # 2741, # 2742, # 2743, and # 2988 (SysKonnnect FDDI and TURBOWAYS MMF ATM).

AIX drivers are shipped on diskette with the hardware feature. Systems ordered with AIX preinstalled will have the device drivers installed on the hard disk drive.

NIM is supported on using FDDI with AIX 4.2.1 and later.

5.1.4.1 PCI SysKonnnect SK-NET FDDI-LP SAS (# 2741)

This adapter provides single attachment to an FDDI concentrator (or point-to-point) using fiber optic cabling (not supplied with the adapter). It also offers high speed, low latency, low CPU and bus utilization, and occupies a single PCI slot.

Features include:

- Occupies single PCI bus slot
- Compatible with the FDDI-ANSI X3T12 specification and FDDI Standard Series
- SMT 7.3 compliance
- AMD SUPERNET3 single chip FDDI design
- SysKonnnect designed custom ASIC for the buffer management with minimal CPU and bus load
- Four independent 128-byte FIFOs for buffering between PCI bus and buffer memory
- 128 KB static RAM on board

- Support under AIX SMIT and diagnostic routines
- Supports remote IPL and NIM with the appropriate system firmware
- Connects to fiber FDDI concentrator ports using fiber optic cable with an FDDI SC (low cost fiber) connection

5.1.4.2 PCI SysKonnnect SK-NET FDDI-LP DAS (# 2742)

This fiber optical FDDI dual-attach station adapter provides either dual attachment to the main ring path or dual homing to one or two FDDI concentrators using fiber optic cabling (not supplied with the adapter). It also offers high speed, low latency, low CPU and bus utilization, and occupies a single PCI bus slot.

Features include:

- Occupies single PCI bus slot
- Compatible with the FDDI-ANSI X3T12 specification and FDDI Standard Series
- SMT 7.3 compliance
- AMD SUPERNET3 single chip FDDI design
- SysKonnnect-designed custom ASIC for buffer management with minimal CPU and bus load
- Four independent 128-byte FIFOs for buffering between PCI bus and buffer memory
- 128 KB static RAM on board
- Support under AIX SMIT and diagnostic routines
- Supports remote IPL and NIM with the appropriate system firmware
- Connects to fiber FDDI concentrator ports using fiber optic cable with an FDDI SC (low cost fiber) connection

5.1.4.3 PCI SysKonnnect SK-NET FDDI-UP SAS (# 2743)

This copper FDDI single-attach station adapter provides single attachment to a FDDI concentrator (or point-to-point) using Category-5 unshielded twisted pair (UTP) cabling (not supplied with the adapter).

Features include:

- Occupies single PCI bus slot
- Compatible with the FDDI-ANSI X3T12 specification and FDDI Standard Series

- SMT 7.3 compliance
- AMD SUPERNET3 single chip FDDI design
- SysKconnect-designed custom ASIC for buffer management with minimal CPU and bus load
- Four independent 128-byte FIFOs for buffering between PCI bus and buffer memory
- 128 KB static RAM on board
- Support under AIX SMIT and diagnostic routines
- Supports remote IPL and NIM with the appropriate system firmware
- Connects to Category-5 unshielded twisted pair FDDI concentrator ports using a shielded RJ-45 connector

Note

The AIX Device Driver diskette is shipped with the adapter.

Limitation: A configuration using more than one SysKconnect FDDI adapter (# 2741, # 2742, or # 2743), where any one of them is feature # 2743 (SysKconnect SK-NET FDDI-UP SAS PCI), constitutes a Class A system.

NIM is supported on this system using FDDI on AIX Version 4.2.1 and later.

5.1.5 ATM Adapters

In this section, information on the available ATM adapters is provided.

5.1.5.1 PCI TURBOWAYS 155 UTP ATM Adapter (# 2963)

The TURBOWAYS 155 PCI UTP ATM Adapter provides dedicated, 155 Mbps full-duplex connection to ATM networks over either permanent virtual circuits (PVC) or switched virtual circuits (SVC). This adapter enables TCP/IP to run over an ATM network with Category-5 Unshielded Twisted Pair (UTP). It also supports communication with devices located on an ATM network or bridged to a token-ring, Ethernet, or other LAN.

5.1.5.2 PCI TURBOWAYS 155 MMF ATM Adapter (# 2988)

The TURBOWAYS 155 PCI MMF ATM Adapter provides direct access to ATM networks. The TURBOWAYS 155 PCI MMF ATM Adapter provides dedicated 155 Mbps full-duplex connection using PVCs or SVCs and enables TCP/IP to run over an ATM network. The adapter also supports communication with

devices located on an ATM network or bridged to a token ring, Ethernet, or other LAN.

Note

Systems ordered with AIX preinstalled will have the device drivers installed on the hard disk drive

The TURBOWAYS 155 PCI MMF ATM Adapter is compatible with:

- IBM 8285 ATM Workgroup Switch
- IBM 8260 ATM Subsystem
- IBM 8282 ATM Concentrator
- IBM 8281 ATM LAN Bridge

A maximum of four high-speed networking adapters are supported per system in any combination. This limitation includes the following adapters:

- TURBOWAYS 155 PCI MMF ATM Adapter (# 2988)
- SysKonnnect SK-NET FDDI-LP SAS PCI for PowerPC Systems (# 2741)
- SysKonnnect SK-NET FDDI-LP DAS PCI for PowerPC Systems (# 2742)
- SysKonnnect SK-NET FDDI-UP SAS PCI for PowerPC Systems (# 2743)

5.1.5.3 PCI TURBOWAYS 25 ATM Adapter (# 2998)

The TURBOWAYS 25 ATM PCI Adapter is for direct access to ATM networks. This adapter provides dedicated 25 Mbps full-duplex connection using PVCs or SVCs and enables TCP/IP to run over an ATM network. The TURBOWAYS 25 ATM PCI Adapter also supports communication with devices located on an ATM network or bridged to a token ring, Ethernet, or other LAN.

Note

AIX drivers are shipped on diskette with the hardware feature. Systems ordered with AIX preinstalled will have the device drivers installed on the hard disk drive.

TURBOWAYS 25 ATM PCI Adapter is compatible with:

- IBM 8285 ATM Workgroup Switch
- IBM 8260 ATM Subsystem
- IBM 8282 ATM Concentrator

- IBM 8281 ATM LAN Bridge

5.1.6 Multiple Port Adapters

In this section, information on the available multi-port asynchronous adapters is provided.

5.1.6.1 PCI 8-Port Asynchronous Adapter EIA-232/RS-422 (# 2943)

The 8-Port Asynchronous Adapter EIA-232/RS-422 is used for connection of up to eight asynchronous EIA-232 or RS-422 devices. All eight ports are software programmable to support either EIA-232E or RS-422A protocols at up to 230 Kbps.

5.1.6.2 PCI 128-Port Asynchronous Controller (# 2944)

The 128-Port Asynchronous Adapter subsystem provides attachment for a high concentration of asynchronous lines (up to 128) from a single PCI bus slot. This gives the system the ability to serve a large number of EIA-232 or RS-422 devices such as terminals, printers, and modems.

Two 2.4 Mbps synchronous channels link the adapter to a maximum of eight 16-Port Remote Asynchronous Nodes (RANs). Up to four RANs can be linked to each synchronous channel.

These RANs provide 16 EIA-232 (# 8137) or 16 RS-422 (# 8138) ports, yielding a total of 128-ports in a maximum configuration. Previous RANs feature # 8130, 8134, or 8136 can also be used with this adapter. If these RANs are connected to the synchronous channel, the data rate drops down to 1.2 Mbps. For best results, keep the previous and current RANs on separate synchronous channels.

5.1.7 WAN Adapters

In this section, information on the available WAN adapters is given.

5.1.7.1 PCI 2-Port Multiprotocol Adapter (# 2962)

The 2-Port Multiprotocol Adapter provides high-speed connections to wide area networks (WANs) through two 2 Mbps ports. To access WAN lines, the adapter connects through an external communications controller, such as a synchronous modem, Data Service Unit (DSU), or Channel Service Unit (CSU). With the appropriate cable, this adapter is compatible with the following specifications:

- X.21
- V.24

- V.35
- V.36

5.1.7.2 PCI IBM ARTIC960HX 4-Port Selectable Adapter (# 2947)

The IBM ARTIC960HX 4-Port Selectable PCI Adapter is a one slot, standard length, 32-bit PCI card. The adapter provides four ports, supporting EIA-232, EIA530, RS-449, X.21 and V.35 specifications. Each port will attach to data networks at speeds up to 2.0 Mbps.

The adapter provides support of WAN protocols such as X.25, SNA, SDLC and Bisync. Software support is provided by ARTIC960 Support for AIX, Developer's Kit, AIX 4.2.1 + PTF, AIX V4.3.2 or later.

Note

The combination of feature numbers 2947 and 2948 must not exceed seven.

5.1.7.3 PCI IBM ARTIC960HX 4-Port T1/E1 Adapter (# 2948)

The IBM ARTIC960HX 4-Port T1/E1 PCI Adapter is a one slot, standard length, 32-bit PCI card. The adapter provides four ports of channelized T1/E1. This adapter also provides SCbus external TDM interface and switch, and is intended for use with the DT/6000 DTQA adapter.

Software support is provided by ARTIC960 Support for AIX, Developer's Kit, AIX 4.2.1, AIX 4.3.2 or later.

5.1.8 Miscellaneous Adapters

In this section, information on the available miscellaneous adapters is given.

5.1.8.1 PCI Ultimedia Video Capture Adapter/S (# 2639)

The Ultimedia Video Capture Adapter/S is a PCI-based adapter that allows the capture of live video and still images. It supports three industry-standard video formats:

- NTSC
- PAL
- SECAM

This adapter, coupled with Ultimedia Services and the appropriate optional hardware, can be used for entry-level video conferencing, security monitoring, security identification, and other image applications.

5.2 Graphics Adapters

In this section, the available graphics adapters are discussed, and a table summarizing their features follows.

IBM offers a broad range of versatile graphics accelerators to meet your customer's application needs, from entry 2D design and drafting to complex 3D solid modeling. These accelerators provide a consistent implementation of open APIs that help ensure application compatibility across the entire family of graphics products.

For entry- to high-performance 2D graphics requirements, IBM has a range of industry-leading 2D accelerators designed for the RS/6000 workstation your customer chooses. These accelerators support the X Window System and provide acceleration for key 2D functions and operations. In addition, most of IBM's family of 2D graphics accelerators provide support for popular 3D application programming interfaces (APIs) when combined with optional IBM software.

From entry to high-end 3D graphics needs, IBM offers a wide range of 3D graphics accelerators that support key open APIs, like OpenGL. These graphics accelerators are designed to offload the system processor by providing varying amounts of hardware acceleration.

This section discusses the following graphics accelerators:

- GXT120P
- GXT250P and GXT255P
- GXT550P
- GXT3000P

Be sure to check the RS/6000 configurator for the application and availability of any RS/6000 adapter.

5.2.1 POWER GXT120P (# 2838)

The POWER GXT120P is a versatile, low-priced graphics accelerator for RS/6000 43P Series systems. It accelerates X Windows graphics applications, such as displaying business graphics, data base status, CASE tools, and activity for securities trading. The GXT120P can support a network control display for monitoring activity and assisting in load balancing, performance analysis and capacity planning.

The GXT120P is also ideal for process control, where, in selected configurations, up to four monitors can be operated concurrently from the same RS/6000 system for displaying graphics and text for dozens of process points. For decision support systems, the GXT120P provides opportunities for rapid, simultaneous display of large amounts of critical information at high resolution on up to four screens.

The GXT120P accelerates UMS/6000 video playback on RS/6000 systems with up to 640 Mbps memory bandwidth. UMS/6000 is an IBM Licensed Program Product that includes video capture, multiple window playbacks and video scaling. Both X and Y interpolative scaling are implemented by the accelerator for zooming and mapping to windows. It can support up to three video streams concurrently. The video scaling feature can be used for collaborative computing and playback of live and stored video.

The GXT120P supports screen resolutions of 640x480, 800x600, 1024x768, and 1280x1024 all at 8-bits per pixels. It satisfies the ISO 9241 standard of visual quality at all these resolutions, refreshing the screen at up to 85 Hz for all resolutions, including the 1280x1024 8-bit resolution.

The GXT120P graphics accelerator is a half-length PCI card with a standard 15-pin DIN output connector.

5.2.2 POWER GXT250P (# 2851) and POWER GXT255P (# 2852)

The POWER GXT250P and GXT255P graphics accelerators provide solid graphics performance at an attractive price. These accelerators supply excellent 2D 8-bit, 16-bit, or 24-bit graphics performance for the RS/6000 43P Series systems. The POWER GXT250P provides a high level of 8-bit 2D graphics performance, while the GXT255P provides 3D graphics performance at entry 2D graphics prices. When coupled with IBM's Softgraphics implementation of the OpenGL, graPHIGS, and PEX 3D APIs, these accelerators offer an exciting level of 3D performance. With high-quality 8-bit double buffer or 24-bit single buffer performance, these graphics accelerators help provide sharper graphical imagery on your screen.

The POWER GXT250P supports 8-bit color depth that allows simultaneous display of up to 256 colors from a palette of approximately 16.7 million colors. With its 3 MB of VRAM (2 MB frame buffer plus 1 MB auxiliary buffer), the POWER GXT250P can meet the demanding needs of most popular graphics applications.

The POWER GXT255P graphics accelerator extends the graphics capability to 24-bit true color required in many high-tech design applications such as

MCAD, GIS and ECAD. The GXT255P has 8-bit, 16-bit and 24-bit true color depth and allows simultaneous display of the entire palette of approximately 16.7 million colors. The GXT255P features three hardware color maps for multi-window graphics without technicolor effects. Its 8 MB of VRAM (6 MB frame buffer plus 2 MB auxiliary buffer) provides outstanding high-end graphics applications.

The GXT250P and GXT255P graphics accelerators are each packaged as a single card that attaches to the local PCI bus slot. The GXT250P and GXT255P provide hardware acceleration for points, lines, triangles, rectangles, quadrilaterals and bit-block transfer, rectangular and non-rectangular clipping, and pattern fill.

These accelerators support display resolutions of 1024x768 and 1280x1024 that are software selectable – no switches required. The POWER GXT255P also supports a display resolution of 1600x1280. Refresh rates are 60 to 85 Hz to comply with Part 3 of the ISO ergonomic standard. When used with compliant monitors, this standard provides improved viewing, reduced flicker, minimized reflections, and sharper characters. The POWER GXT250P and GXT255P automatically adjust to the highest resolution and fastest refresh rate of the attached monitor, supporting the DDC2B VESA standard.

5.2.3 POWER GXT550P (# 2845)

The POWER GXT550P graphics accelerator offers exceptional and affordable 3D graphics for the PCI-based RS/6000 43P Model 150. Coupled with IBM's implementation of many of the industry's most popular APIs, it is an excellent fit for today's most demanding 3D graphic applications, such as mechanical design and analysis, petroleum exploration and production, molecular modeling, and scientific research.

The GXT550P further establishes IBM as a leader in providing price/performance solutions for a wide range of demanding customer application requirements. The GXT550P combines 3D rasterization with the strength of the latest PowerPC microprocessors. This rasterization technology provides customers with a cost-effective solution that delivers excellent 3D graphics acceleration that scales with the PowerPC processor performance of the RS/6000 workstation.

The POWER GXT550P provides hardware acceleration for the OpenGL, PHIGS, and IBM GL 3.2 3D APIs, and additionally supports the PEX API through software. IBM is a leader in the industry by offering native support of OpenGL and PHIGS on the same accelerator. This adapter accelerates advanced 3D graphics functions, such as Gouraud shading, antialiasing,

hidden surface removal, depth-cueing, and transparency. This enables your 3D applications to run more quickly and with good interactive performance. For OpenGL, texture maps up to 1024x1024, and a 64-bit accumulation buffer are also supported through software.

The GXT550P offers highly flexible frame buffers that can be dynamically configured to provide a broad set of color and feature options. When using OpenGL, PHIGS and GL 3.2, the GXT550P supports 8-bit, 12-bit, and 24-bit double-buffered color and also includes 8-bits of double-buffered alpha buffers for more realistic transparency control.

The GXT550P provides 8-bit overlay buffers, which enhance the speed of a graphical user interface (GUI), 8-bit stencil buffers, and 24-bit Z-buffer for hidden surface removal operations. This accelerator provides for the display of up to 16.7 million simultaneous colors from a palette of 16.7 million colors and four hardware color maps.

The GXT550P graphics accelerator is a 32-bit PCI cards that plugs into a single PCI slot. It supports display resolutions of 1280x1024 and 1024x768 and refresh rates from 60 Hz to 85 Hz, including refresh rates that comply with the ISO 9241 Part 3 ergonomic standard. When used with other compliant components, including monitors, this standard provides improved viewing, reduced flicker, reduced reflection and sharper characters. These features, coupled with IBM's powerful RS/6000 43P Models, make the POWER GXT550P an ideal solution for demanding 3D application needs.

5.2.4 POWER GXT3000P Graphics Accelerator (# 2825)

The POWER GXT3000P Graphics Accelerator for 3D visualization marks a breakthrough in performance and functionality for design and visualization solutions. It is the highest performing IBM graphics accelerator available for RS/6000 workstations. The graphics subsystem, combined with IBM's POWER3 and PowerPC-based workstations, delivers outstanding speed and 3D performance for demanding applications in mechanical design and engineering, for aerospace and automotive, petroleum exploration and production, scientific visualization and other technical industries.

The GXT3000P is well suited for users with applications that require high-performance 3D graphics. It delivers significant improvement in functionality and performance (up to three times the performance of its predecessor, the POWER GXT800P graphics adapter) at a lower cost.

The GXT3000P provides hardware acceleration for the new OpenGL 1.2 version of the API. This adapter accelerates advanced 3D graphics functions

such as Gouraud shading, antialiasing, depth buffering, fog and atmospheric effects, and blending. This should enable your 3D applications to run more quickly and with good interactive performance. Additional benefits to using the POWER GXT3000P graphics accelerator are:

- Increased subpixel addressing over previously supported adapters
- Face culling support
- OpenGL polygon offset support
- OpenGL polygon mode support
- OpenGL specular lighting support in hardware
- Trilinear texture mapping

Featured hardware enhancements over the POWER GXT800P Graphics Adapter include:

- OpenGL and graPHIGS lighting mode
- 3D texture mapping
- Video scaling
- Stereo in a window

Note

Stereo viewing requires additional equipment, such as emitter, cable and glasses, in addition special software support.

The POWER GXT3000P offers a highly flexible frame buffer that can be dynamically configured to provide a broad set of color and feature options. When using OpenGL and graPHIGS APIs, the GXT3000P offers support for 8-bit and 24-bit double-buffered color, and also includes 8-bits of double-buffered alpha buffers for more realistic transparency control. In addition, this accelerator provides 8-bit overlay buffers that enhance the speed of a graphical user interface (GUI), 8-bit stencil buffers, 8-bit window ID buffer, and a 24-bit Z-buffer for hidden surface removal operations. The POWER GXT3000P allows for the display of up to 16.7 million simultaneous colors.

The GXT3000P graphics accelerator is a single 64-bit PCI card that attaches to one 64-bit PCI slot in the RS/6000 Models 150 or 260, but covers the slot adjacent to it (uses two slots of space).

Note

Only one GXT3000P 3D graphics adapter may be installed per supported workstation. GXT3000P slot limitations are as follows:

- Slot 3 in the RS/6000 7043 Model 150
- Slot 2 in the RS/6000 7043 Model 260

The same GXT3000P card can be put in either the Model 150 or 260.

The GXT3000P supports display resolutions of up to 1280x1024 at 85 Hz refresh rate, including monitors that comply with the ISO 9241 Part 3 ergonomic standard.

5.2.5 Graphics Adapter Feature Comparison

Table 34 provides the specifications and features for the RS/6000 graphics accelerators currently available on the RS/6000 43P 7043 Model 150 and Model 260 systems.

Note

The Model 150 supports all the graphics adapters listed in the table. The Model 260 supports the GXT120P, GXT250P, GXT255P and GXT3000P graphics adapters.

Table 34. Graphics Adapters

Adapter	POWER GXT120P	POWER GXT250P	POWER GXT255P	POWER GXT550P	POWER GXT3000P
Feature code	# 2838	# 2851	# 2852	# 2845	# 2825
Class	2D only	2D;3D (Soft)	2D;3D (Soft)	3D	3D
PCI slots occupied	1	1	1	1	2 ²
Max/system	4 ¹	4	4	1	1
ISO compliant	Yes	Yes	Yes	Yes	Yes

Adapter	POWER GXT120P	POWER GXT250P	POWER GXT255P	POWER GXT550P	POWER GXT3000P
Resolution	640 x 480 800 x 600 1024 x 768 1280x1024	1024 x 768 1280 x 1024	1024 x 768 1280 x 1024	1024 x 768 1280 x 1024	1280 x 1024
Frame buffer	8-bit single	8-bit single (1280 x 1024); 8-bit double (1024 x 768)	8-bit double (up to 1280 x 1024); 24-bit single	8-, 12-, 24-bit double	8/24-bit double
Color palette	16.7 million	16.7 million	16.7 million	16.7 million	16.7 million
Concurrent colors	256	256	Up to 16.7 million	Up to 16.7 million	Up to 16.7 million
Z-buffer	No	No	No	24-bit	24-bit
HW color map	1	3	3	4	8
Overlay buffer	0-bit	0-bit	8-bit	8-bit	8-bit
Dithering	No	Yes	Yes	Yes	No
Depth cueing	No	No	No	Yes	Yes
Antialiasing	No	No	No	Yes	Yes
Gouraud shading	No	No	No	Yes	Yes
Local lighting	No	No	No	No ³	No ³
NURBS	No	No	No	No ³	No ³
Texture mapping	No	No	No	HA ⁴	Yes
Alpha buffer	No	No	No	8-bit double	8-bit double
Stencil buffer	No	No	No	8-bit	8-bit
Accumulation buffer	No	No	No	No ³	No ³

Adapter	POWER GXT120P	POWER GXT250P	POWER GXT255P	POWER GXT550P	POWER GXT3000P
Supported APIs	Xlib	Xlib, OpenGL, GL 3.2, PHIGS/PEX ⁵	Xlib, OpenGL, GL 3.2, PHIGS/PEX ⁵	Xlib, OpenGL, PHIGS ⁵	Xlib, OpenGL, PHIGS ⁵

¹The Model 150 can accommodate up to four POWER GXT120P adapters, while the maximum on the Model 260 is two.

²The GXT3000P adapter is attaches to one 64-bit slot in the RS/6000 Models 150 and 260, but covers the slot adjacent to it (uses two slots of space).

³Through Softgraphics, this graphics accelerator supports lighting, NURBs, alpha buffer, stenciling, and accumulation buffering.

⁴Texture Mapping is accomplished in software, but with hardware acceleration (HA) by the graphics board.

⁵The current version of AIX (V4.3.2) no longer supports PEX and GL 3.2.

5.2.6 Graphics Adapters Attachment Cables

Table 35 provides a display cable feature number matrix for the graphics adapters.

Table 35. Display Cable Feature Code Matrix

Displays	GXT120P (15D/DDC)	GXT250P/ GXT255P (15D/DDC)	GXT550P (13W3/DDC)	GXT3000P (15D/DDC)
G52 ¹	Display ²	NS ³	NS	NS
P50 ¹	Display	Display	4213	Display
P70	4238	4238	4240	4238
P200	4238	4238	4240	4238
P201	4237	4237	4241	4237
P72 ¹	Display	Display	4213	Display
P92 ⁴	Display	Display	4240	Display
P202 ⁴	Display	Display	4240	Display
6091-19I	NS	4239	4219	4217

Displays	GXT120P	GXT250P/ GXT255P	GXT550P	GXT3000P
	(15D/DDC)	(15D/DDC)	(13W3/DDC)	(15D/DDC)
POWERDP17	4217	4239	4219	4217
POWERDP20	4217	4239	4219	4217
9516	4217	4217	4214	4217

¹Captured 15-pin D-shell cable.

²Display - The appropriate cable is included with the display.

³NS = Not Supported - The display/adaptor combination is not supported.

⁴Dual input 15-pin D-shell and 13W3 soft switchable via screen menu; ships with 15-pin D-shell.

5.2.7 Multiple Adapter Support

The RS/6000 43P 7043 Model 150 and 260 machines can support multiple graphics adapters, depending on the adapter. Check the adapter documentation to determine if there are model specific limitations. This section discusses how to configure a second adapter.

5.2.7.1 Configuring a Second PCI Graphics Adapter

Installing and using a second graphics adapter is very simple:

1. Shut down and power-off the RS/6000.
2. Install the second graphics card (refer to the system's User Guide that is provided with your system).
3. Power-on the RS/6000.

During the firmware boot phase, only one monitor will display the POST icons.

After the AIX boot image takes control in the boot process, the system will present the following message on the first display:

```
***** Please define the System Console *****
Type the F1 key and press Enter to use this display as the
system console.
```

The following message will appear on the second display:

```
***** Please define the System Console *****
Type the F2 key and press Enter to use this display as the
system console.
```

This process sets the default system console for the current session and for subsequent reboots. Upon subsequent reboots, the request to define the system console will not appear.

The graphics adapter should have been configured by AIX automatically and can be used immediately. To determine if the graphics adapter has been installed, you can run the `lsdisp` command. This will show the graphics adapters available in your system:

```
lsdisp
```

DEV_NAME	SLOT	BUS	ADPT_NAME	DESCRIPTION
=====	=====	===	=====	=====
gga0	02	pci	GXT120	IBM Personal Computer Power Series GXT120 Graphics Adapter
gga1	05	pci	GXT120	IBM Personal Computer Power Series GXT120 Graphics Adapter

```
Default display = gga0
```

5.2.7.2 Changing the Default Display

In the previous example, the `gga0` device is the default graphics adapter. This can be changed to `gga1` by using the `chdisp` command. By using this command, you can temporarily change the default display or change the default display for all subsequent reboots. You can change the default display to `gga1` for the current session by typing the following:

```
chdisp -d gga1
```

To change the display permanently for all subsequent reboots:

```
chdisp -p gga1
```

If you are using the Common Desktop Environment (CDE), things are slightly different. To change the display, you have to drop down to a command line login; you cannot run the `chdisp` command from the desktop. From there, you can change the current session with the `-d` option, but when you log out and return to the CDE login, the previous display will be used. To change the current CDE login display to a different adapter, you have to drop down to the command line and run the `chdisp` command with the `-p` option. When you log

out, the CDE login window will appear on the new display. This will also change the default display for subsequent reboots.

5.2.8 Accessing the Second Display

If you have the CDE desktop installed and have another screen attached to the second adapter, you can access the adapter immediately. By addressing the adapter with the DISPLAY variable, you can open X Windows on the second screen. The last digit in the DISPLAY determines on which screen the X client will be displayed. This is shown in Table 36.

Table 36. Accessing a Second Display

DISPLAY Variable	Description
DISPLAY=hostname:0.0	Displays on the default monitor
DISPLAY=hostname:0.1	Displays on the second monitor

The mouse can be used to move between the current desktop screen on the default adapter and the screen displayed by the second adapter. While you can access the second display by moving the mouse off the first, it is not possible to drag windows with the mouse from the first screen to the second.

The following is an example of using the second display:

1. Open an xterm on the second display:

```
xterm -display hostname:0.1
```

2. Move the mouse off the right of the default screen. The mouse should appear on the other screen. You can now gain focus on the X Window. From here, you use the new window or open more windows.

Accessing a second display by using standard X Windows and Motif is also very easy. You can start X Windows and access the second display using the same method described above. In addition, you can *stack* the screens in any order you like. So, for example, you could have the two screens in a vertical environment. You can start X Windows in a vertical environment by running the command:

```
xinit -- -P21 gga0 -P11 gga1
```

This will display the first screen, using `gga0`, and the second below the first, using `gga1`. To access the second display in this environment, you would move the mouse off the bottom of the screen.

Since the CDE environment allows virtual screens, having a separate display does not provide a large benefit. There are, however, some applications

where this can be very useful. In situations where some data or application must be permanently visible, this setup can be used. This option provides more screen area and will allow you to have more windows visible without obscuring the desktop.

5.3 Adapter and Device Configuration on PCI-Based RS/6000 Systems

This section includes step-by-step procedures explaining how to configure devices on the PCI-based RS/6000 systems. It also provides cable types and cabling information that are necessary on most adapters.

5.3.1 Device Types

PCI-based RS/6000 43P systems support the attachment of devices using the following buses:

- SCSI
- PCI
- ISA (limited to built-in support)

Devices on PCI-based RS/6000 systems can be of two types: native devices, such as the integrated SCSI controller and the serial controller, or attached (non-native) devices, such as the supported PCI adapters.

The device information required to configure the integrated devices is contained in the firmware and is passed to the operating system through the residual data written to NVRAM. Thus, integrated devices are configured automatically to AIX after system reboot.

Attached PCI devices and adapters can be configured automatically because the PCI Local Bus Specification defines a configuration method for PCI components. Provided that the device support software is installed, PCI and SCSI devices are configured automatically whenever the configuration manager program, `cfmgr`, is run at system boot or run time, and when no conflict (for example, the same SCSI ID for two different SCSI devices) is found.

The ISA/EISA bus has no standard method of identifying adapters or their configuration requirements.

5.3.1.1 Device Configuration Database

Device information is contained in a predefined database or in a customized database that makes up the Device Configuration Database managed by the Object Data Manager (ODM).

The predefined database contains configuration data for all possible devices configurable for the system.

The customized database contains configuration data for all currently defined devices in the system.

The device information stored in the Device Configuration Database allows the automatic configuration PCI devices on PCI-based RS/6000 systems whenever the Configuration Manager (`cfgmgr`) program is run at system boot and run time.

The Device Configuration Database that is currently used by the system is located in the directory that is specified by the ODMDIR environment variable. This is normally the `/etc/objrepos` directory.

In AIX Version 3.2, all device software was installed when the operating system was installed. In AIX Version 4, only those software packages for devices that were detected at the installation time are automatically installed.

Refer to *Managing AIX V.4 on PCI-Based RS/6000 Workstations (40P/43P)*, SG24-2581, or to the AIX Version 4 product documentation for more information on device management and enhancements to the Device Configuration Database.

5.3.1.2 Device Location Codes

Location codes are defined slightly differently depending on the type of device they are used for, non-SCSI or SCSI.

For planars, cards and non-SCSI devices, the location code is defined as:

AB-CD-EF-GH

For SCSI devices, the location code is defined as:

AB-CD-EF-G,H

While **AB-CD-EF** have the same meaning for both non-SCSI and SCSI devices, the last two letters identify different codes.

AB identifies a bus type, and it can assume the following values:

- 00** for resources attached to the processor bus, such as system planar, processor, memory, L2 cache, and primary PCI bus.
- 01** for resources attached to ISA buses, such as diskette drive, mouse adapter, mouse, and keyboard.

04 for resources attached to PCI buses, such as ISA bus, SCSI controller, CD-ROM drive, hard disks, and graphics adapter.

- For SSA, these values are always zero.

CD identifies a slot or adapter number.

- The possible values for CD depend on the adapter or card. For pluggable adapters or cards, this will be a two-digit slot number in the range from 01 to 99.
- For integrated adapters, the first character in CD will be a letter in the range from A to Z. The letter is based on the order that the integrated adapters are defined in residual data and ensures unique location codes for the integrated adapters. The second character, D, will be set to 0.

Any adapter or card is defined only with AB-CD. SSA adapters will show the system I/O bus identifier here.

EF is the connector identifier.

- On adapters with multiple connectors, it is used to identify the adapter connector that a resource is attached to. For SSA, this will show the physical disk drive module and the logical disk drive.

GH for non-SCSI devices is a port identifier, address, DIMM, devices or a Field Replaceable Unit (FRU). It has several meanings, depending upon the resource type.

G,H for SCSI devices defines the following:

- **G** defines the control-unit address of the device
- **H** defines the logical-unit address of the device

Figure 35 shows an example of PCI-based RS/6000 systems' location codes.

name	status	location	description
sys0	Available	00-00	System Object
sysplanar0	Available	00-00	System Planar
pci0	Available	00-fef00000	PCI Bus
pcil	Available	00-fee00000	PCI Bus
isa0	Available	10-58	ISA Bus
siota0	Available	01-Q1	Tablet Adapter
ppa0	Available	01-R1	CHRP IEEE1284 (ECP) Parallel Port Adapter
sa0	Available	01-S1	Standard I/O Serial Port
sal	Available	01-S2	Standard I/O Serial Port
siokma0	Available	01-K1	Keyboard/Mouse Adapter
fda0	Available	01-D1	Standard I/O Diskette Adapter
scsi0	Available	10-60	Wide/Fast-20 SCSI I/O Controller
ent0	Available	10-80	IBM 10/100 Mbps Ethernet PCI Adapter (23100020)
scsi1	Available	10-88	Wide/Ultra-2 SCSI I/O Controller
mtn0	Available	20-60	GXT3000P Graphics Adapter
sioka0	Available	01-K1-00	Keyboard Adapter
sioma0	Available	01-K1-01	Mouse Adapter
fd0	Available	01-D1-00-00	Diskette Drive
rmt0	Available	10-60-00-0,0	SCSI 4mm Tape Drive
cd0	Available	10-60-00-1,0	SCSI Multimedia CD-ROM Drive
hdisk0	Available	10-60-00-8,0	16 Bit SCSI Disk Drive
kbd0	Available	01-K1-00-00	PS/2 keyboard
mouse0	Available	01-K1-01-00	3 button mouse
lvdd	Available		LVM Device Driver
mem0	Available	00-00	Memory
proc0	Available	00-00	Processor
proc2	Available	00-02	Processor
L2cache0	Available	00-00	L2 Cache
pmc0	Available	00-00	Power Management Controller
lft0	Available		Low Function Terminal Subsystem
rootvg	Defined		Volume group
hd5	Defined		Logical volume
hd6	Defined		Logical volume
hd8	Defined		Logical volume
hd4	Defined		Logical volume
hd2	Defined		Logical volume
hd9var	Defined		Logical volume
hd3	Defined		Logical volume
hd1	Defined		Logical volume
inet0	Available		Internet Network Extension
en0	Available		Standard Ethernet Network Interface
et0	Defined		IEEE 802.3 Ethernet Network Interface
lo0	Available		Loopback Network Interface
pty0	Available		Asynchronous Pseudo-Terminal
gxme0	Available		Graphics Data Transfer Assist Subsystem
rcm0	Available		Rendering Context Manager Subsystem
aio0	Defined		Asynchronous I/O
hd7	Defined		Logical volume
catia	Defined		Logical volume
catadm	Defined		Logical volume

Figure 35. Devices Location Codes

5.3.1.3 Configuration Information

Information about the PCI adapters can be obtained by running the `lscfg` command.

To determine how much memory you have in your machine, run the following command:

```
lsattr -El mem0
```

This will show the size of memory in your machine and how much usable memory there is.

You can also see the size of the L2 cache installed in your system by using the `lsattr` command:

```
lsattr -El L2cache0
```

5.3.2 PCI Adapter Configuration

If the required device drivers are installed, PCI adapters are configured automatically at system restart or by using the `cfgmgr` command.

When you install the AIX operating system, it will query the PCI bus to identify all AIX supported adapters. The installation process will install the required device drivers and diagnostic support (if available), and the system will automatically configure the PCI adapters.

If you have a system installed and running, you can use `smit cfgmgr`. This method will require you to put the AIX installation media in place (CD-ROM, tape, and so forth) before configuring the adapters. The media will be used only if the device drivers for your adapters are not currently installed.

You can also manually install the device drives and diagnostics before running `cfgmgr` or rebooting the system for automatic PCI adapter configuration.

When using PCI adapters, the slot number is important. If you change an adapter from slot A to slot B, the original adapter information for slot A will be kept in a `defined` state; a new definition will be created for slot B, and a new device will be configured.

5.3.2.1 Configuring Non-Graphic PCI Adapters

Configuring PCI adapters is very easy because they are actually configured automatically when their driver software is installed and the `cfgmgr` command is run.

The easiest way to configure non-graphic PCI adapters (such as LAN adapters, and so on) installed after IPL is to proceed using SMIT, enter:

```
smit devices
-> Install/Configure Devices Added After IPL
```

You may have to provide a source for the device driver software package for the new PCI device if it is not yet installed.

Note

Do not use the above procedure if you want to add or replace graphic adapters because if anything goes wrong (for example, a required device driver is not yet installed, and so forth) you will disable your graphics console.

To add or change PCI graphic adapters, proceed as described in 5.2, "Graphics Adapters" on page 135.

If, for whatever reason, you need to install software but you are unable to boot from hard disk or you have no working display, you may install the software while in maintenance mode.

5.3.2.2 Configuring SCSI Devices

The RS/6000 43P Series systems have at least one integrated Ultra SCSI Interface.

The process of configuring SCSI devices on the 7043 Series is exactly the same as it is on other RS/6000 Systems. This means you can use the SMIT support to configure devices that are added after IPL, enter the `cfgmgr` command, or reboot the system. However, remember that unlike with AIX Version 3.2, on AIX Version 4 the software support for the device you are adding may not yet be installed. Thus, check to see if the device driver for your SCSI device is installed. If not, use SMIT to configure the device and provide an input device or directory where the installation software can be accessed.

The Web-Based System Management tool, `wsm`, provides a similar function as SMIT, but using a visually rich GUI and task oriented dialogs.

5.4 Summary Chart of Adapters

Table 37 summarizes the adapters (communication, graphics, and storage) available on the RS/6000 43P 7043 Model 150 and Model 260 systems.

Table 37. Communication, Graphics, and Storage Adapters Summary

Adapter	Feature Code	On Model 150	On Model 260
Token Ring PCI	# 2920	Yes	Yes
10/100 Mbps Ethernet PCI	# 2968	Yes	Yes
Gigabit Ethernet-SX PCI	# 2969	NO	Yes
ISDN Basic Rate PCI	# 2708	Yes	Yes
SysKonnnect SK-NET FDDI-LP SAS PCI	# 2741	Yes	Yes
SysKonnnect SK-NET FDDI-LP DAS PCI	# 2742	Yes	Yes
SysKonnnect SK-NET FDDI-UP SAS PCI	# 2743	Yes	Yes
TURBOWAYS 155 UTP ATM PCI	# 2963	Yes	Yes
TURBOWAYS 155 MMF ATM PCI	# 2988	Yes	Yes
TURBOWAYS 25 ATM PCI	# 2998	Yes	Yes
8-Port Asynchronous EIA-232/RS-422 PCI	# 2943	Yes	Yes
128-Port Asynchronous Controller PCI	# 2944	Yes	Yes
2-Port Multiprotocol PCI	# 2962	Yes	Yes
ARTIC960Hx 4-Port Selectable PCI	# 2947	NO	Yes
ARTIC960Hx 4-Port T1/E1 PCI	# 2948	NO	Yes
Ultimedia Video Capture PCI	# 2639	Yes	Yes
POWER GXT120P	# 2838	Yes	Yes
POWER GXT250P	# 2851	Yes	Yes

Adapter	Feature Code	On Model 150	On Model 260
POWER GXT255P	# 2852	Yes	Yes
POWER GXT550P	# 2845	Yes	NO
POWER GXT3000P	# 2825	Yes	Yes
PCI Single-ended Ultra SCSI	# 6206	Yes	Yes
PCI Differential Ultra SCSI	# 6207	Yes	Yes
PCI SSA Multi-Initiator/RAID EL	# 6215	Yes	Yes

Chapter 6. Diagnostics and System Management Services

One of the most important features provided by the PCI-based RS/6000 workstations and servers is support of the AIX diagnostics programs. Workstations, and especially servers, should by their nature, provide enhanced reliability, availability, and serviceability (RAS) features. Access to AIX diagnostics, stand-alone and on-line, helps provide these features.

In addition to providing diagnostic support for new PCI devices available for the PCI-based RS/6000 systems, a major change had to be made to the firmware. An extra function had to be added to the firmware to allow a logical representation of the keyswitch that would allow the system to either boot in normal or service mode. The Models 150 and 260 have an electronic keyswitch function (see 8.1, "Electronic Key-Switch Function" on page 215) and the system can either be booted in Normal or Service mode by pressing a function key during the firmware POST stage. Additionally, support was added to enable the system to load diagnostics from a CD-ROM.

Two methods of running diagnostics are supported on the PCI-based RS/6000 systems:

- Stand-Alone Diagnostics
- On-line Diagnostics

Both of these methods are described in this chapter.

6.1 Diagnostics Operating Considerations

The following items identify some issues to consider before using the diagnostics programs.

- **Device Diagnostics**

If diagnostics support is not installed on your system and you install this option, device support for devices physically installed in your machine will not be automatically installed. Diagnostics support is, however, automatically installed when you install from an AIX Version 4 CD-ROM (check with the `ls1pp -L` command if `bos.diag.*` and `bos.rte.diag` filesets are already installed on your system since they are required). The On-line Diagnostics test list will display entries for all the devices in your system that have diagnostic support.

- **Selecting a Console Display**

When you run Stand-Alone Diagnostics, you will have to select the console display. The diagnostic program displays instructions on graphics displays and on terminals attached to the S1 serial port.

- **Identifying the Terminal Type to Diagnostics**

When you run diagnostics, the diagnostic program must know what type of terminal you are using. If the terminal type is not known when the Function Selection menu is displayed, the diagnostics will not allow you to continue until a terminal is selected from the Define Terminal option menu.

- **Undefined Terminal Type**

If an undefined terminal type from the Define Terminal option menu is entered, the menu will prompt you to enter a valid terminal type, and the menu will be re-displayed until a valid type is entered or you exit from the option.

- **Resetting the Terminal**

If you enter a valid terminal (according to the Define Terminal option), but it is not the correct type for the ASCII terminal being used, difficulty may be encountered in reading the screen, using the functions keys, or the Enter key. These difficulties can be bypassed by pressing **Ctl-c** to reset the terminal. The screen display that results from this resetting action varies with the mode in which the system is being run:

- On-line Concurrent or Single-User Mode - the command prompt appears
- Stand-Alone Mode or On-line Service Mode - the terminal type will be reset to dumb, the diagnostics operation instruction panel will be displayed, and you will be required to go through the Define Terminal process again

- **Running Diagnostics from a TTY Terminal**

The following should be considered when running diagnostics using a tty terminal as the console display:

- Support for some tty terminals is optionally installed. If you attach a tty terminal to a system to run diagnostics, be aware that it may not work properly since AIX support for the terminal may not be installed.
- See the operator manual provided with your terminal to find the key sequences required to respond to the diagnostics.
- When the diagnostics present display information through the S1 port, certain attributes are used. These attributes are set as if the

diagnostics were using an IBM 3161 display terminal. The attributes for the 3151, 3161 and 3164 are listed in the *Diagnostics Information for Multiple Bus Systems*, SA38-0509, Appendix D.

6.2 Stand-Alone Diagnostics

Stand-Alone Diagnostics provides a method of testing the system without running AIX or without having diagnostics installed. This is useful when On-line Diagnostics are not available, when devices which are in use during the operation of AIX are to be tested, or if the system is in a state where the On-line Diagnostics cannot be run.

The stand-alone diagnostics are packaged on removable media. They must be booted or mounted before they can be run. If booted, you have no access to the AIX error log or to the AIX configuration data, but you can test all the components in your system except for the SCSI adapter used to boot the diagnostics CD-ROM. If the diagnostics are run in concurrent mode using stand-alone diagnostics (the CD-ROM is mounted while AIX is running), then you have access to the AIX error log and the AIX configuration data, but you are limited to the devices you can test.

The following are some considerations that you should make when running stand-alone diagnostics:

- The diagnostics disc must remain in the CD-ROM drive for the entire time that diagnostics are executing. The disc can only be ejected after the system has been turned off and then turned on again (stand-alone mode), or after the diagnostics program has terminated (on-line concurrent mode).
- The CD-ROM drive from which diagnostics are loaded cannot be tested.
- The SCSI adapter (or circuitry) controlling the CD-ROM drive from which diagnostic are loaded cannot be tested.

6.2.1 Booting the Stand-Alone Diagnostics CD-ROM

To start the stand-alone diagnostics, perform the following steps:

1. Remove all diskettes and CD-ROMs.
2. Turn the system unit's power on.
3. Insert the diagnostics CD-ROM into the CD-ROM drive.
4. When the keyboard POST indicator appears, press **F5** on the directly attached keyboard, or the number **5** key on the ASCII terminal keyboard,

to indicate to the system that the default Boot Device Order List is to be used.

5. Enter any requested password, if set. (Power-On or Privileged-Access).
6. Follow the instructions to select the console.
7. Follow the displayed instructions to test the desired resources.
8. When testing is complete, use the **F3** key to return to the Diagnostic Operation Instructions screen.

6.2.2 Running Stand-Alone Diagnostics in Concurrent Mode

There might be a situation where on-line diagnostics is not available, but you have access to the stand-alone diagnostics CD-ROM. To keep the system available to the users, however, you do not want to reboot and run the diagnostics programs in service mode. In this case, you can mount the CD-ROM and run diagnostics concurrently. The procedure for this is explained as follows:

1. From the command line, type:

```
diag
```

2. If you already have a CD-ROM in the CD-ROM drive, unmount any file systems on the CD-ROM and remove the CD-ROM.
3. You will be asked whether you want to run diagnostics from CD-ROM. Answer *y*.
4. Insert the CD-ROM.
5. Press **Enter** if you want to use the default CD-ROM drive. If you are using a CD-ROM that is not the default, enter its name.

Diagnostics will then mount the CD-ROM, and you can access diagnostics as if they were loaded on your system. When you exit diagnostics from the CD-ROM, the CD-ROM will be unmounted automatically.

6.3 On-line Diagnostics

The on-line diagnostics are an AIX Version 4 installable package. They are packaged under `bos.diag` (you will also need to have the `bos.rte.diag` fileset installed). There is a separate, installable option for each device or adapter supported by diagnostics. If you install your system from an AIX Version 4 CD-ROM, diagnostics will be automatically installed on your system, including diagnostics support for any PCI devices that are physically installed in your system. In addition, if you install any PCI devices and configure them

to the system using `cfgmgr -i` (or use **SMIT -> Devices -> Install/Configure Devices Added After IPL**), specifying an installation device, then diagnostics support for the device (if available) will be automatically installed. Running on-line diagnostics provides you with access to the AIX error log and the AIX configuration data, but again, the devices you can test are limited.

If the diagnostics package, `bos.diag` is installed on your system, you can determine which devices in your system have diagnostic support loaded by performing the following steps:

1. Enter `diag`.
2. Press **Enter** on the main screen, and then select **Task Selection**.
3. Select **Display Configuration and Resource List**.

If a device installed in your system is not listed, refer to your installation media to see if the device comes with a diagnostics support fileset.

On-line Diagnostics can be run in the following three modes:

- Service Mode allows checking of most system resources.
- Concurrent Mode allows the normal system functions to continue while selected resources are being checked.
- Single-User Mode allows checking of most system resources.

These three modes are discussed in the following sections.

6.3.1 Service Mode

Running diagnostics in service mode provides the most thorough test of the system resources available with On-line Diagnostics. As in the concurrent mode, however, there are resources that cannot be tested. These resources are the SCSI adapter and disks used for paging. You can also analyze the error log in this mode by selecting the **Problem Determination** option on the Diagnostic Routine Selection menu.

To access diagnostics from Service mode, perform the following steps:

1. Stop all programs including the AIX operating system (`shutdown -F`).
2. Remove all tapes, diskettes and CD-ROM discs.
3. Turn the system unit's power off and wait for at least 30 seconds.
4. Turn the system unit's power on again.
5. After the first POST indicator appears on the system unit's console, press **F6** on the directly attached keyboard, or the number **6** key on the ASCII

terminal keyboard, to indicate to the system that the customized Boot Device Order List is to be used. Normally, the first boot device in this list is your boot disk, which should also contain the bos.diag package.

Note: The term POST indicator refers to the Icons (graphic display), or device mnemonics (ASCII display) that are displayed while the POST are executing.

6. Enter any requested password, if set (Power-On or Privileged-Access).
7. Select the terminal type (for example ibm3151 or lft).

Once diagnostics has loaded, you can follow the instructions given on the console to test your system. After you have tested your system, you should exit diagnostics by pressing the **F3** key and follow the instructions displayed to shut down AIX. Failure to shut down AIX could result in damage to data on your disks.

6.3.2 Concurrent Mode

The concurrent mode provides a way to run on-line diagnostics on the system resources while AIX is up and running.

Since the system is running in normal operation, some resources cannot be tested in concurrent mode. The following list shows which resources cannot be tested:

- SCSI adapters used by disks connected to paging devices
- The disk drives used for paging
- Memory
- Processor

There are three levels of testing in concurrent mode:

- The Share-test level tests a resource while it is being shared by programs running in normal operation. This test is mostly limited to normal commands that test for the presence of a device or adapter.
- The Sub-test level tests a portion of a resource while the remaining part of the resource is being used in normal operation. For example, this test could test one port of a multiport device while the other ports are being used in normal operation.
- The Full-test level requires the device not be assigned to, or used by, another operation. This level of test on a disk drive may require the use of the varyoff operation. The diagnostics display menus allow you to vary off the desired resource.

Error Log Analysis can be performed in Concurrent mode by selecting the **Problem Determination** option on the Diagnostic Routine Selection menu.

To run On-line Diagnostics in Concurrent mode, just enter the following command as root:

```
diag
```

After a period of time (up to three minutes), you will be presented with the diagnostics initial screen.

6.3.3 Single-User Mode

Diagnostics can run from single-user mode (maintenance mode - although this is different from booting into maintenance mode from a CD-ROM or a tape). You can shut down your system into a single-user mode, which will allow testing of many of your system resources that would normally be in use while AIX is running. All the system resources, except the SCSI adapters and disks used for paging, memory and processors, can be tested. Error log analysis is available in this mode. Other users will not be allowed to log into the system while the system is in single-user mode.

Logged on as root, you can enter diagnostics in single-user (maintenance) mode by performing the following:

1. Stop all applications and log all users off the system.
2. Enter `cd /` to change into the root directory.
3. Enter the `shutdown -m` command.
4. When you see a screen similar to this one:

```
Init: Single User Mode
Password:
```

the system prompts you for the root password (if set).

5. Enter the root password.
6. At the command prompt (`#`), enter the `diag` command.
7. When the Diagnostic Operating Instructions area is displayed, follow the instructions to check the desired resources.

After you have completed testing the system, press **F3** to return to the diagnostic operating instructions. Press **F3** again to return to the AIX operating system prompt. Press the **Ctrl-D** key sequence to log off from root.

The following messages appears:

```
Do you wish to leave single user mode?  
Enter (Y)es, (N)o or <cr> for the default runlevel:
```

Just press **Enter** to go back to the default runlevel.

6.4 Diagnostic Tasks and Service Aids

The AIX diagnostic package contains programs that are called tasks. This section describes the tasks available in AIX diagnostics Version 4.2 or later (screen output is based on your installed hardware). For more detailed information, refer to *7043-150 User's Guide*, SA38-0512 or *7043-260 User's Guide*, SA38-0553. The tasks are as follows:

- **Add Resource to Resource List:**
Use this task to add resources back to the resource list.
- **Analyze Adapter Internal Log:**
The PCI RAID adapter has an internal log that logs information about the adapter and the disk drives attached to the adapter. Whenever data is logged in the internal log, the device driver copies the entries to the AIX system error log and clears the internal log.
- **Backup and Restore Media:**
This service aid allows verification of backup media and devices.
- **Certify Media:**
This task allows the selection of diskette or hardfiles to be certified.
- **Change Hardware Vital Product Data:**
Use this service aid to display the Display/Alter VPD Selection Menu. The menu lists all resources installed on the system.
- **Configure Dials and LPFKeys:**
This service aid provides a tool for configuring and removing dials/LPFKeys to the asynchronous serial ports.
- **Configure Reboot Policy:**
This service aid controls how the system tries to recover from a system crash.

- **Configure Remote Maintenance Policy:**
This policy includes modem configurations and phone numbers to use for remote maintenance support.
- **Configure Ring Indicate Power-On Policy:**
This service aid allows you to power-on a system by telephone from a remote location (runs only on CHRP/RPA system units).
- **Configure Surveillance Policy:**
This task monitors the system for hang conditions, that is, hardware or software failures that cause operating system inactivity.
- **Create Customized Configuration Diskette:**
A Stand-alone Diagnostic Package Configuration Diskette is created in order to change the default refresh rate for an LFT from the console.
- **Delete Resource from Resource List:**
Use this task to delete resources from the resource list.
- **Disk Maintenance:**
You can make a disk to disk copy to another SCSI drive of similar size, as well as display and alter information on disk sectors.
- **Display Configuration and Resource List:**
This service aid displays the item header only for all installed resources.
- **Display Firmware Device Node Information:**
This task displays the firmware device node information that appears on CHRP/RPA platforms.
- **Display Hardware Error Report:**
A tool for viewing the hardware error log is provided.
- **Display Hardware Vital Product Data:**
This service aid displays all installed resources, along with any VPD for those resources.
- **Display Microcode Level:**
The current level of microcode on an adapter or device can be displayed.
- **Display or Change Bootlist:**
This service aid allows the bootlist to be displayed, altered, or erased.

- **Display or Change Diagnostic Run Time Options:**
The Display or Change Diagnostic Run Time Options task allows the diagnostic run time options to be set.
- **Display Previous Diagnostic Results:**
Each time the diagnostics produce a service request number (SRN) to report a problem, information about that problem is logged. The service representative can look at this log to see which SRNs are recorded.
- **Display Resource Attributes:**
This task displays the Customized Device Attributes associated with a selected resource.
- **Display Service Hints:**
This service aid reads and displays the information in the CERADME file from the diagnostics media.
- **Display Software Product Data:**
This task invokes SMIT to display information about the installed software.
- **Display System Environmental Sensors:**
This service aid displays the environmental sensors implemented on a CHRP/RPA system.
- **Display Test Patterns:**
This task provides a means of adjusting system display units by providing displayable test patterns.
- **Download Microcode:**
This service aid provides a way to copy microcode to an adapter or device.
- **Flash SK-NET FDDI Firmware:**
This task allows the Flash firmware on the SysKonnnect SK-NET FDDI adapter to be updated.
- **Format Media:**
The selection of diskettes, hardfiles, or optical media can be formatted.
- **Generic Microcode Download:**
This service aid provides a means of executing a generic script from a diskette. The intended purpose for this generic script is to load microcode to a supported resource.

- **Local Area Network Analyzer:**
This selection is used to exercise the LAN communications adapters, such as Token-Ring, Ethernet, and Fiber Distributed Data Interface (FDDI).
- **Periodic Diagnostics:**
This selection provides a tool for configuring periodic diagnostics and automatic error log analysis.
- **Run Diagnostics:**
The Run Diagnostics task invokes the Resource Selection List menu. When the commit key is pressed, diagnostics are run on all selected resources.
- **Run Error Log Analysis:**
The Run Error Log Analysis task also invokes the Resource Selection List menu. When the commit key is pressed, Error Log Analysis is run on all selected resources.
- **Save or Restore Hardware Management Policies:**
Use this service aid to save or restore the settings from Ring Indicate Power On Policy, Surveillance Policy, Remote Maintenance Policy and Reboot Policy.
- **SCSI Bus Analyzer:**
This service aid provides a means to diagnose a SCSI bus problem in a free-lance mode.
- **SCSI Device Identification and Removal:**
This service aid allows you to choose a SCSI device or location from a menu and to identify a device located in a system unit that uses a SCSI Enclosure Services (SES) backplane.
- **SCSD Tape Drive Service Aid:**
This task provides a means to obtain the status or maintenance information from a SCSD tape drive.
- **Spare Sector Availability:**
This selection checks the number of spare sectors available on the optical disk.
- **SSA Service Aid:**
This service aid provides tools for diagnosing and resolving problems on SSA attached devices.

- Update Disk Based Diagnostics:
Use this service aid to apply fixes (APARs) on your system.
- Update System or Service Processor Flash:
This selection updates the system or service processor flash for CHRP/RPA system units.

6.5 Understanding the Firmware

The firmware is the first code to execute when the machine is powered on. This code resides on the hardware, and it is loaded at the factory.

The firmware uses the console and the speaker as a means to communicate the results of certain tests. By default, the firmware uses the serial port S1 to display logos, messages, and errors. In case a supported graphics adapter is installed in the machine, the firmware will redirect the output to this adapter instead of the serial port.

6.5.1 General RS/6000 Boot Sequence

RS/6000 CHRP and RPA platforms will display a graphical or text message (hereafter referred to as the boot prompt) after the banner step that identifies the platform and communicates to the user that there are options that may be invoked to alter the boot process.

Once the boot prompt is displayed, the System Management Services (SMS) menu can be invoked (by pressing the **F1** key). SMS provides a user interface for utilities, configuration, and the Multiboot Menu for boot/install and the Open Firmware command interpreter.

For boot failures related to firmware trying to access a boot device, it is appropriate for the platform to retry the boot operation, especially in the case of booting from a network device. However, in platforms that have a service processor (Model 260), there are several other types of detected errors for which a reboot retry may be appropriate. For example, check-stops or loss of communication between firmware and the service processor. The service processor configuration provides a user setting for the number of retries that should be attempted in case of boot failures (see also 7.5.2.3, “Reboot/Restart Policy Setup Menu” on page 196).

The following sections describe the boot phase of operation.

6.5.1.1 Identify and Configure System Components

The firmware must, by various means, become aware of every component in the system and configure or reset those components into a known state. Components include all bridges and device controllers, but may exclude devices that are not involved in the boot process.

Firmware is generally written with a hardware platform in mind so that some components and their configuration data can be hardcoded. Examples of these components are: number and type of processors, cache characteristics, and the use of components on the planar. This hardcoding is not a requirement, only a practical approach to a part of this task.

Certain system information must come from walking the I/O buses. This is a technique that will yield identification of I/O device controllers and bridges that reside on modern, well-behaved buses such as PCI.

6.5.1.2 Generate a Device Tree

The firmware will build a device tree. Certain configuration information (configuration variables) may be stored in non-volatile memory.

6.5.1.3 Initialize/Reset System Components

The AIX operating system requires devices to be in a known state at the time control is transferred from the firmware. Firmware may gain control of the hardware in various states depending on what has initiated the boot process.

6.5.1.4 Locate an OS Boot Image

A device and filename can be specified directly from the AIX boot command. Once the boot image is located, the device path is set in the device tree as the bootpath property of the chosen node.

If multi-boot (multiple bootable operating systems residing on the same hardware platform) is supported, a configuration variable instructs the firmware to display a multi-boot menu from which the OS and bootpath are selected.

6.5.1.5 Load the Boot Image into Memory

After locating the image, it is loaded into memory at the location given by a configuration variable or as specified by the OS load image format.

6.5.2 Typical Boot Sequence for RS/6000 43P 7043 Model 150

After the system unit power is turned on, the system firmware tests the system memory and identifies failing memory cards and memory module locations. System firmware checks to see what devices are in the system and

performs a simple test on them. The system firmware displays the device name or device icon being tested. After the keyboard name or icon appears, you can enter the Systems Management Services menu by hitting the **1** key (on a ASCII terminal) or the **F1** key (on a graphics terminal).

You can also enter one of the following:

- **5** (on a ASCII terminal) or **F5** (on a graphics terminal) to start the Stand-Alone Diagnostics (CD).
- **6** (on a ASCII terminal) or **F6** to start the On-line Diagnostics (Hard Disk)

Depending on the boot device, a checkpoint may be displayed on the operator panel for an extended period of time while the boot image is retrieved from the device. If a bootable image is found, then the system unit loads and starts the operating system. The AIX Boot code indicates LED progress codes in the range from 05XX to 09XX. After the AIX Boot is completed, the AIX login prompt appears on the main console.

Note

For more information about the progress codes, refer to *Diagnostic Information for Multiple Bus Systems, SA38-0509*, Be sure to read the chapter on Diagnostics Numbers and Codes.

If no bootable operating system image is found on any of the storage devices listed in the system firmware boot list, then the system unit attempts to boot again. With each attempt, you have the opportunity to start the System Management Services (SMS) as described in 6.7.1, "SMS Start-Up" on page 173.

6.5.3 Typical Boot Sequence for RS/6000 43P 7043 Model 260

Because of its integrated service processor, the boot sequence of the Model 260 is different than the Model 150.

After the A/C power is turned on, the System Support Controller (SSC) start-up begins, and releases reset to the service processor. If the SSC cannot communicate with the service processor, the LED displays a 4BA00000 code. If the service processor is not present, the LED displays 4BA00001.

Service Processor Self-Test

The service processor card performs a self test and NVRAM initialization. The LED Code range is E000 - E07F and the LED shows OK code when complete.

Service processor in Standby Mode

You can enter the service processor menus whenever the LED code is OK, STBY, or has an eight digit error code on the LED display by pressing the **Enter** key on an ASCII terminal connected to serial port S1.

Built-In-Self-Test (BIST)

The service processor initiates a Built-In-Self-Test (BIST) on the Central Electronics Complex (CEC) chips when the **power** button is pressed. The VPD is read, the CRC is done, and the processor compatibility test is run. LED Code range is E0A0 - E0E1.

System Initialization

System firmware begins to execute and then initializes system registers after LED code E0E1. LED Code range is E1XX - E2XX.

Memory Test

The system firmware tests the system memory and identifies failing memory cards and memory module locations. LED Code range is E3XX.

Device Configuration and Test

System firmware checks to see what devices are in the system and performs a simple test on them.

The system firmware displays the device name or device icon being tested. After the keyboard name or icon appears, you can enter the Systems Management Services menu by hitting the **1** key (if using a ASCII terminal) or the **F1** key (if using a Graphics terminal).

You can also enter one of the following:

- **5** key (on an ASCII terminal) or **F5** (on a graphics terminal) to start the Stand-Alone Diagnostics (CD).
- **6** key (on an ASCII terminal) or **F6** (on a graphics terminal) to start the On-line Diagnostics (Hard Disk)

IPL Boot Code

The system firmware attempts to boot from the devices listed in the Boot List. The LED Code range is E1XX.

Boot Image Execution

Once a boot image is located on a device in the Boot List, the system firmware code hands off control to the AIX Boot code.

- LED Code E105 for Normal Boot or
- E15B for Service Mode Boot.

The AIX Boot code indicates LED progress codes in the range from 05XX to 09XX.

Note

For more information about the progress codes, refer to *Diagnostic Information for Multiple Bus Systems, SA38-0509*. Be sure to read the chapter on Diagnostics Numbers and Codes.

AIX Boot Complete

The AIX login prompt appears on the main console.

6.6 Power-On Self-Test

Power On Self Test (POST) is the process by which the firmware tests those areas of the hardware that are critical to its ability to carry out the boot process. It is not intended to be all-inclusive or to be sophisticated by how it communicates with the user. Diagnostics with these characteristics will generally be provided as a service aid. The following power-on self tests are performed by the firmware on PCI-based RS/6000 systems:

6.6.1 Processor POST

The firmware tests the processor very briefly. If there is a failure, the system will stop, and the initial beep of the speaker will not be heard.

6.6.2 Memory DIMM and L2 Cache POST

The firmware RAM POST tests all the memory, it is done in segments. After the initial short beep, if the system stops with the hard file indicator on solid, this is the indication that no memory was found. If some memory is found, then a working 2 KB continuous segment is sought. If the 2 KB segment cannot be found, the system stops with the hard file indicator on solid.

The L2 cache is tested. If it fails, there will be a momentary indication when the display becomes available. The error is also written to the Error Log. During the system initialization, when icons are being displayed for the

individual subsystems, one of the icons will be the memory icon. During the display of this icon, an 8-digit error code will be displayed if the L2 cache is not fully functional.

The remainder of memory is also tested, and if a failure is detected, an 8-digit error code is posted. This error indicates which Dual Inline Memory Module (DIMM) slot is at fault. The basic algorithm for the system memory configuration is as follows:

1. Query each memory slot to see if there is a memory DIMM installed.
 - If yes, determine its size. Identify each DIMM as IBM or non-IBM, ECC, or parity. This is done by querying special identification registers in the DIMMs themselves.
 - If the DIMM could not be identified, display the memory slot number, as well as a warning saying that a DIMM could not be identified. Proceed with the next DIMM.
2. Test the DIMM.
 - If this is an IBM DIMM and it passes the test, make it available and add its size to the available memory quantity. If it does not pass the test, keep the DIMM out of the system configuration and display a warning on the screen reporting the faulty DIMM slot number.
 - If this is a non-IBM DIMM and passes the test, make it available and add its size to the available memory quantity. In addition, display a message on the screen reporting that a non-supported, non-IBM memory DIMM has been found, tested, and correctly configured. If the DIMM does not pass the test, display a warning on the screen reporting that a non-IBM memory module was found and failed to pass the test.
3. Proceed with the next DIMM.

At the end of this process, the system will have available all the memory it could configure (IBM branded or not).

6.6.3 Graphics Adapter POST

The firmware tests the graphics adapter but not the display. The system will give a speaker beep sequence when the graphics adapter is found to be non-functional. The beep sequence is one long beep, followed by one short beep. The error is logged, and the system boot continues.

6.6.4 Keyboard and Mouse Controller POST

The keyboard/mouse controller has a built-in Power-On Self Test, and the firmware checks the result. If an error is detected, an 8-digit error code is displayed. The error is written to the Error Log, and the system continues.

6.6.5 SCSI Subsystem POST

The SCSI POST opens the adapter and runs several test scripts that check and exercise the various functions related to the chips on the integrated SCSI adapter, such as DMA, simulated mailbox commands, and so on. If an error is detected, an 8-digit error code is displayed, the error is written to the Error Log, and the system continues.

Note

The 8-digit POST error codes are listed and described in the Service Guide for the corresponding PCI-based RS/6000 system:

RS/6000 7043 43P Series Service Guide, SA38-0512

RS/6000 7043 Model 260 Service Guide, SA38-0554

6.7 System Management Services (SMS)

The SMS programs make it possible for you to view information about your computer, run memory tests, and set the device from which the server will boot.

By default, the firmware uses the serial port S1 to display logos, messages, errors, and also as the interface to the SMS programs. If a supported graphics adapter is installed in the machine, the firmware will redirect the output to this adapter instead of the serial port. This means, if you have installed a graphics adapter but you want to use an ASCII terminal attached to serial port S1 as the console for the SMS programs, you will have to remove the graphics adapter first.

The SMS programs provide two types of interfaces: a graphical user interface (GUI) that can be used on a graphical display and an ASCII interface that can be used either on an ASCII terminal or on a graphical display.

Note: The ASCII interface provides more functions than the graphical interface.

To start the programs:

1. Turn on or shut down and restart your computer.
2. When the first screen appears, press the **F1** key to use the GUI. On an ASCII terminal, you can press either the number **1** key.

Note: You must press keys listed above before the last indicator appears.

6.7.1 SMS Start-Up

During the boot process, function keys can be pressed to alter the electronic key-switch and determine which boot list to use during the boot procedure. There are four function keys that can be pressed that determine the type of boot that will occur. These are described as follows:

F1-Key (1 with an ASCII terminal)

The **F1** key pressed on a directly attached keyboard will invoke the SMS GUI interface. Press the **1** key at an ASCII terminal, and it will invoke the SMS ASCII interface. The electronic key-switch will be set to Normal mode.

Mksysb Booting

Booting from a tape (to restore a `mksysb` tape or boot into maintenance mode) requires SMS. From SMS, you can select an alternative boot device which includes a tape drive if one is installed.

F5-Key (5 with ASCII terminal)

The **F5** key will cause the default firmware boot list to be used. The electronic key-switch is set to Service mode. This mode attempts to boot from the first device of each found in the list. It does not search for other bootable devices of that type if the first device is not bootable. Instead, it continues to the next device type in the list. If after one pass through the boot list no boot record is found, the firmware will invoke the SMS program.

F6-Key (6 with ASCII terminal)

When **F6** is selected, the firmware looks for a boot record according to the custom boot list that was set up by SMS or by using the `bootlist` command. If the boot list was not customized, the default boot list is used. The electronic key-switch is set to Service mode. Again, if after a pass through the custom boot list no boot record is found, the firmware will invoke the SMS program.

F8-Key (8 with ASCII terminal)

Press **F8** to access the Open Firmware `OK>` prompt. This has no influence on the selection of the boot list or key-switch.

Note

When AIX is installed on a machine, it automatically sets the customized boot list to the installed hard disk. The customized boot list is thereafter used by the system when booting in Normal mode.

6.7.2 SMS Graphical Main Menu

The graphical System Management Services main menu offers the following tasks, as shown in Figure 36:

- Config** Enables you to view your system setup
- Multiboot** Enables you to set and view the default operating system, modify the boot sequence, access the Open Firmware command prompt, and other options
- Utilities** Provides several tools to set and remove passwords, to set the unattended start mode, to update the firmware, to set the IP addresses for network booting, and other utilities




Figure 36. System Management Services - Main Menu

6.7.2.1 SMS Submenu Config

Selecting this icon makes it possible for you to view information about the setup of your system unit. A list similar to the following appears when you select the Config icon. If more than one screen of information is available, a blue arrow in the top right corner of the screen appears. Use the **Page Up** and **Page Down** keys to scroll through the pages.

6.7.2.2 SMS Submenu Multiboot

The options, as shown in Figure 37, available from this screen allow you to view and set various options regarding the operating system and boot sequence:

- Select Software** The Select Software option allows you to choose which operating system to use.
- Software Default** This option enables you to select an operating system as the default operating system for your system unit.
- Install From** Enables you to select a media drive from which to install an operating system.
- Boot Sequence** Enables you to view and change the custom bootlist (the sequence in which devices are searched for operating system code).
- OK Prompt** Enables you to go directly to the Open Firmware command prompt. The Open Firmware command prompt is used for debugging and device driver development.
- Multiboot Start** Clicking on this button toggles whether the Multiboot menu will appear automatically at start-up.

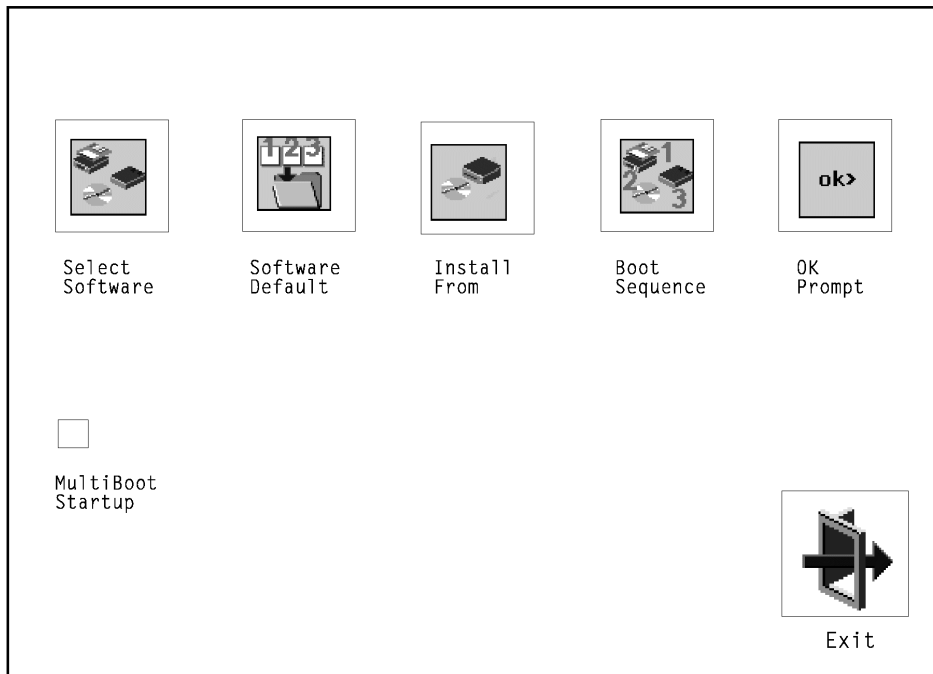


Figure 37. SMS Submenu Multiboot

6.7.2.3 SMS Submenu Utilities

Selecting this icon enables you to perform various tasks and view additional information about your system unit, shown in Figure 38:




Figure 38. SMS Submenu Utilities

1. Firmware Passwords

You can limit access to the system by using passwords. The passwords can be up to eight characters long and are stored in NVRAM. There are three categories of restricted access:

- Power-On Password

When the system is powered on, it checks whether a power-on password (POP) is present. If there is one present and the unattended start mode is not set, it means the machine's owner does not want the system to be used unless the POP password is supplied. In this case, the system will prompt for the POP password. You are given three attempts to enter the correct password. If you fail to supply the correct password, the system will go into a hung state and must be powered-off before continuing.

- Privileged Access Password

The Privileged Access Password (PAP) is used to limit access to the System Management Services (SMS) programs (also called the Supervisory Password). When you press one of the keys to access SMS, the system will check to see if a Privileged Access password exists, if it does, you are prompted to enter the Privileged Access password. You are given three attempts to supply the correct password. If you fail to do so, the system will go into a hung state and must be powered-off before continuing.

- Unattended Start Mode

To use this mode, a power-on password must be previously specified. If unattended start mode is enabled, the system will boot from the defined boot device without requiring the you to enter the power-on password. While the system can be booted without entering the POP password, the keyboard controller will be programmed to lock up until the POP password is provided. This mode is ideal for servers that run unattended. After an electrical power failure, for example, the operating system will be rebooted without waiting for you to enter the power-on password.

Figure 39 shows the SMS submenu where you can set the password.




Figure 39. SMS Submenu - Firmware Passwords

Note

For details about enabling and disabling the power-on password and privileged access password, refer to the relevant chapter in the appropriate service guide:

RS/6000 7043 43P Series Service Guide, SA38-0512

RS/6000 7043 Model 260 Service Guide, SA38-0554

2. Hard Disk Spin Up Delay

Enables you to change the spin up delay for SCSI hard disk drives attached to your system. Spin up delay values can be entered manually, or else a default setting can be used. All values are measured in seconds. The default is two seconds.

3. Error Log

Error Log enables you to view and clear the firmware error log information for your system unit.

4. RIPL (Remote Initial Program Load)

Enables you to select a remote system from which to load programs via a network adapter when your system unit is first turned on. This option also allows you to configure network adapters that require setup. You can ping a remote adapter to check if the network hardware is working correctly.

5. SCSI ID

Allows you to view and change the addresses (IDs) of the SCSI controllers attached to your system unit.

6. Update

Allows you to update the system firmware program. Refer to 6.7.4, "SMS Firmware Update" on page 180.

7. Console

If you have more than one keyboard and display attached to your system unit, or if you have an ASCII terminal attached to your system unit in addition to a keyboard and display, this tool allows you to define which one is active.

6.7.3 SMS ASCII Main Menu

The ASCII version of the System Management Services includes an additional function in the main menu not found in the graphical version. The following languages are supported:

- English
- French
- German
- Italian
- Spanish
- Swedish

All menus and options, as well as the button titles, have been translated. Your ASCII terminal must support the ISO8859-1 character set in order to properly display languages other than English.

6.7.3.1 ASCII Terminal Settings

If you have a graphics adapter in your system, then System Management Services will only display information using that adapter. If you have an ASCII terminal attached and have a graphics adapter installed, nothing will be displayed on the ASCII terminal during the firmware boot process. To be able to use the ASCII terminal to access System Management Services, you will have to physically remove the graphics adapter from the system. If there are no terminals or graphics adapters attached, the system will still be able to boot up.

If you are using an ASCII terminal, the following parameters should be used:

- 9600 baud
- 8-bit
- No parity
- 1 stop bit
- 24 lines by 80 columns
- Echo mode on

6.7.4 SMS Firmware Update

There are several situations where a firmware update may be required. For example, an update may be required after a field EC (engineering change) to provide firmware support for new adapters and peripherals or to enhance functions for existing peripherals.

6.7.4.1 How to Determine the Current SMS Firmware Version

The firmware version is displayed on the initial start-up screen on a graphical display and at the top of each System Management Services (SMS) menu on an ASCII Terminal display. It can also be found on the configuration display under the SMS menus.

To access the SMS menus, power-on the system if it is off, or shutdown and re-start the system if it is on. On a graphical display, the firmware version is displayed in the lower right corner of the initial display. Press the **F1** key while the start-up icons are displayed to stop at the SMS menus.

On an ASCII Terminal display, press the **1** key at the initial display when the ASCII icons are displayed at the bottom of the screen. Note the firmware version at the top of the menus.

Select **Configuration** from the initial SMS menu to see the system's configuration. The firmware version is near the bottom of this list.

6.7.4.2 How to Obtain the SMS Firmware Update

In order to create a firmware diskette with the latest level of SMS firmware available, download the newest firmware for your machine from:

<http://www.rs6000.ibm.com/support/micro>

There are two types of archives currently being distributed: one for AIX and one for DOS-compatible operating systems such as DOS, OS/2, and Windows. Both archives are self-extracting executables that will unpack themselves when run.

To unpack the AIX archive, filename.bin, execute the following commands:

```
chmod +x filename.bin
./filename.bin
```

To unpack the DOS-compatible archive, filename.exe, execute the archive:

```
filename1
```

(¹For example: 7043140. This depends on the model for which you want to download the firmware).

6.7.4.3 How to Create the SMS Firmware Diskette

The method of creating the firmware update diskette will vary depending on the operating system being used and the distribution archive. In all cases, 1.44 MB DOS-formatted diskettes should be used. For AIX users, DD images should be provided. To create a firmware update diskette, insert a blank

diskette into the drive of the system being used. If a makefw script is provided, execute it. Otherwise, type the following command:

```
dd if=filename.dd of=/dev/rfd0
```

For users of DOS-compatible operating systems, discrete files will be provided with a batch file named makefw to copy the appropriate files to diskette. Run the batch file to create the diskette:

```
makefw
```

6.7.4.4 How to Perform the SMS Firmware Update

Once you have a firmware update diskette available, you are ready to perform the update. Follow the steps below carefully.

The procedure below will guide you through **soft-loading** the new firmware which loads it into memory and runs it. This is a very important step which will ensure that your new firmware is not corrupted and will run on your system as it is configured *before* the new firmware is permanently flashed into the system. If the soft-load fails, *DO NOT* flash the new firmware.

1. Power-on the system, if off, or shutdown and re-start.
2. Go to the **open firmware ok** prompt:
 - From a graphics display, press **F8** when the icons appear at start-up.
 - From an ASCII terminal, press **8** when the ASCII icons appear during start-up.
3. Insert the diskette with the firmware image and soft-load it by typing the following commands at the ok prompt:
 - > setenv load-base 600000
 - > boot floppy:1,\name_of_firmware_level.img

Note: name_of_firmware_level depends on the system for which you want to update the firmware (For example: tig98153).
4. A console selection may be run at this time to allow you to choose a display. A number will appear on each display that can be used for the firmware icons and SMS menus. Press the number of the display you wish to use. A graphics display is preferable.
5. Go to the SMS menu:
 - From a graphics display, press **F1** when the icons appear at start-up.
 - From an ASCII terminal, press **1** when the ASCII icons appear during start-up.

6. Verify that the new firmware is currently running. See 6.7.4.1, “How to Determine the Current SMS Firmware Version” on page 181. If the new firmware is not running, something went wrong. Go back to step one and try again. If the new firmware IS running correctly, continue to the next steps to permanently flash the new firmware image.

Note

If power is lost during the firmware update procedure, the system could be rendered inoperable. To avoid this, DO NOT POWER OFF the system during the Firmware Update, and take any steps to ensure that power is not interrupted during the update procedure.

7. From a graphics display, use the firmware update utility. Select **Utilities** from the menu, then **Update** and follow the on-screen instructions to complete the update procedure.
8. From an ASCII terminal, use the firmware update utility. Select **Utilities** from the main menu, then **Update system firmware** and follow the on-screen instructions to complete the procedure.

6.7.4.5 SMS Firmware Recovery

If the system will not boot up fully, and the checkpoint display reads FEB, the firmware is partially corrupted, but should be recoverable. Make a copy of the firmware update diskette, rename the .img file on this diskette to PRECOVER.IMG and insert this recovery diskette in the drive when the FEB appears on the checkpoint display.

If the system drops to the ok prompt at checkpoint F4D, this may also be a partial corruption of the firmware. Use the ok prompt to attempt the soft-load procedure and re-flash the new firmware. Start at step three under 6.7.4.4, “How to Perform the SMS Firmware Update” on page 182.

If the checkpoint display reads 888 when the power button is pushed, the system firmware is completely corrupted. Report the problem to your service organization.

Note

The diskette containing the new firmware update will be your firmware recovery diskette.

6.7.5 SMS Firmware Checkpoints

Firmware Checkpoints are intended to let users and service personal know what the system is doing, in some detail, as it initializes. These checkpoints occur during system start-up and are not intended to be error indicators, but in some cases a system could hang at one of the checkpoints. The most appropriate action is included with each checkpoint.

6.7.5.1 RS/6000 43P 7043 Model 150 Firmware Checkpoints

The firmware for the Model 150 uses progress codes (checkpoints) in the range of E1xx to EFFF. The Exxx-Codes are documented in the *RS/6000 7043 43P Series Service Guide*, SA38-0512.

6.7.5.2 RS/6000 43P 7043 Model 260 Firmware Checkpoints

The firmware for the Model 260 uses progress codes (checkpoints) in the range of E1xx to EFFF. The Exxx-Codes for the Model 260 are documented in the *RS/6000 7043 Model 260 Service Guide*, SA38-0554.

Chapter 7. Service Processor

This chapter gives you an overview of the functionality of the Service Processor (SP) and how to use the SP. The SP is currently available on the Model 260, but not the Model 150.

7.1 Introduction to the Service Processor

The Service Processor is a hardware component that is imbedded onto the I/O planar of the Model 260.

The SP has the capability to control the system NVRAM and real time clock interface. The SP card has a interface to assist with the bring-up and testing of complex processor systems. In addition, the SP has the capability to take over either S1 or S2 serial ports on the I/O planar.

The SP card can use S1 or S2 to make a call home through an externally attached modem when it detects one of the following conditions:

- System firmware surveillance failure
- Operating system surveillance failure
- Critical environmental failures
- Restarts

The SP also has the capability to answer modem calls and generate a host interface to an external caller. The host functions of the SP card are contained in flash memory, resident on the SP card. The host modes are password protected and feature many levels of menus with support for several languages.

Machine diagnostics can be launched from the host session, allowing you to remotely diagnose the status of the server. Almost all of this function is available even if the main processor is completely dead. The SP provides an immediate means to diagnose, check status, and sense the operational conditions of a remote server.

The Model 260 can also be booted remotely, and a remote user may enter the Systems Management Services (SMS) screens (only in ASCII mode).

7.2 Customize and Access the Service Processor Menus

The SP menus enable you to configure SP options and to enable and disable functions. Service Processor menus are available using an ASCII terminal when the server is powered off and the Service Processor is operating with standby power. SP menus are also available when server power is on and the SP has detected a server problem (such as loss of surveillance). During the first power up (when the power cord is plugged into the outlet), Service Processor menus are not available for 45 seconds while the SP is running self-tests and initializing the server. If the server powers down, SP menus become available after 15 seconds.

The Service Processor menus are divided into two groups:

- General user menus
- Privileged user menus

To prevent loss of control in unstable power environments, the Service Processor leaves the menu mode after 5 minutes of inactivity. Menus may be resumed by pressing any key on the terminal, local or remote.

Note

You must have a ASCII terminal or a modem to access the SP menus. There is no graphical interface.

7.2.1 How to Customize the Service Processor

There are two methods through which the service processor's functions can be customized:

- SP menus can be customized while AIX is up and running (use the `diag` command) through Service Aid in AIX Diagnostics. Although accessible using concurrent AIX diagnostics, the SP functions you can customize there are relatively few.
- SP menus are accessible when AIX is not running, either locally or remotely.

The following chapters cover how to customize the SP when AIX is not running.

7.2.2 How to Access Service Processor Menus Locally

Service Processor menus may be accessed locally by connecting an ASCII terminal to either serial port. You must press a key on the ASCII terminal to confirm the presence of the Service Processor. The SP then prompts you for a password (if set), and when verified, displays the SP menus.

7.2.3 How to Access Service Processor Menus Remotely

Service Processor menus may be accessed remotely by connecting a modem to serial port 1 or serial port 2.

1. Power-off the server, unplug the power cord, and press the **power** button to drain capacitance while power is disconnected.
2. Connect the modem to the appropriate serial port and turn the modem on.
3. Plug in the server.

Note: Your modem has to be configured. For modem configurations, see 7.5.5.1, “Modem Configuration Menu” on page 203.

With a terminal or terminal emulator, call the server's modem. The Service Processor prompts you for a password (if set), and when verified, displays the SP menus.

7.3 Security Within the Service Processor

For security purposes, the Service Processor counts the number of attempts to enter correct passwords. The results of not recognizing a correct password within this error threshold are different, depending on whether the attempts are being made locally (at the server) or remotely (using a modem). The error threshold is three.

If the error threshold is reached by someone entering passwords at the server, the Service Processor exits the menus. This action is taken based on the assumption that the server is in an adequately secure location with only authorized users having access. Such users must still successfully enter a login password to access AIX.

If the error threshold is reached by someone entering passwords remotely, the Service Processor disconnects the modem to prevent potential security attacks on the server by unauthorized remote users.

There are two different kind of passwords:

- Privileged Access Password

- General Access Password

Passwords can be any combination of up to 8 alphanumeric characters. You can enter longer passwords, but the entries are truncated to include only the first 8 characters. Passwords can be set from the Service Processor menu or from the SMS menus.

Table 38 illustrates which menu can be accessed with the Privileged Access Password and the General Access Password.

Table 38. Settings of Privileged and General Access Password

Privileged Access Password	General Access Password	Resulting Menu
None	None	MAIN MENU displayed
None	Set	MAIN MENU displayed
Set	None	Users with passwords see the MAIN MENU, others users see the GENERAL USER MENU
Set	Set	Users see menus associated with the entered password

7.4 Service Processor Operational Phases

This section provides a high-level flow of the phases of the SP, as shown in Figure 40.




Figure 40. Service Processor - Operational Phase Status Flow

7.4.1 Pre-Standby Phase

This phase is entered when the server is connected to a power source. The server may or may not be fully powered on. This phase is exited when the Power-On Self-Tests (POST) and configurations tasks are completed. The Pre-Standby phase components are:

- SP Initialization
SP performs any necessary hardware and software initializations.
- SP POST
SP conducts Power-On Self Tests on its various work and code areas.
- SP Unattended Start Mode Checks
To assist fault recovery, if unattended start mode is set, the SP automatically reboots the server. SP does not wait for a user-input or power-on command, but moves straight through the phase and into the Bring-up Phase. The unattended start mode can be reset by accessing SMS menus, or SP menus.

7.4.2 Standby Phase

The standby phase can be reached in two ways:

1. With the server OFF and power connected (the normal path), recognized by OK in the LED display.
2. With the server ON after an operating system fault, recognized by STBY or an 8-digit code in the LED display.

In the Standby phase, the SP takes care of some automatic duties and is available for menus operation. The SP remains in the standby phase until a power-on request is detected.

The Standby phase components are:

- Modem Configuration
SP will configure the modem (if installed) so that incoming calls may be received, or outgoing calls may be placed.
- Dial In
Monitor incoming phone line to answer calls, prompt for a password, verify the password and remotely display the standby menu. The remote session can be mirrored on the local ASCII console if the server is so equipped and the user enables this function.

- Menus

The SP menus are password protected. Before you can access them you need to know either General User (Power-On Password or POP) or Privileged User (Privileged Access Password or PAP).

7.4.3 Bring-Up Phase

This phase is entered upon power-on, and exited upon loading of the operating system. The Bring-Up phase components are:

- Retry Request Check

The SP checks to see if the previous IPL attempt failed. If two consecutive fails are detected, the SP displays an error code and places an outgoing call to notify an external party if the user has enabled this option.

- Dial Out

The SP can dial a pre-programmed telephone number in the event of an IPL failure. The SP issues an error report with the last reported IPL status indicated and any other available error information.

- Update Operator Panel

The SP displays Operator Panel data on the ASCII terminal if a remote connection is active.

- Environmental Monitoring

Environmental Monitoring is now controlled by the SP instead of the base system, with expanded error recording and reporting.

- System Firmware Surveillance (Heartbeat Monitoring)

The SP monitors and time the interval between system firmware heartbeats.

- Responding to System Processor Commands

The SP responds to any command issued by the system processor.

7.4.4 Run-time Phase

This phase includes the tasks that the SP performs during steady-state execution of the operating system.

- Environmental Monitoring

The SP monitors voltages, temperatures, and fan speeds.

- Responding to System Processor Commands

The SP responds to any command issued by the system processor.

- Run-Time Surveillance (Heartbeat Monitoring)

If the device driver is installed and surveillance enabled, the SP monitors the system heartbeat. If the heartbeat times out, the SP places an outgoing call. This is different from the Bring-Up phase scenario where two reboot attempts are made before placing an outgoing call.

7.5 Main Menu

At the top of the MAIN Menu is a listing containing:

- Your Service Processor's current firmware version
- The firmware copyright notice
- The system name given to your server during setup

You need the firmware version for reference when you either update or repair the functions of your service processor.

Following is an example of the SP Main Menu:

```
Service Processor Firmware
  SP Level: px980902
  EPROM: 19980902
  FLASH: 19980902
  Copyright 1997, IBM Corporation
  SYSTEM NAME

MAIN MENU

1. Service Processor Setup Menu
2. System Power Control Menu
3. System Information Menu
4. Language Selection Menu
5. Call-In/Call-Out Setup Menu
6. Set System Name
99. Exit from Menus

1>
```

Note: The Service Processor prompt reads either 1> or 2> to indicate which serial port on the system unit is being used to communicate with the SP.

The following sections cover the different submenus.

7.5.1 Service Processor Setup Menu

The service processor setup menu appears as follows:

```
Service Processor SETUP MENU

1. Change Privileged Access Password
2. Change General Access Password
3. Enable/Disable Console Mirroring:
   Currently Disabled
4. Start Talk Mode
5. OS Surveillance Setup Menu
6. Reset Service Processor
7. Reprogram Service Processor Flash EPROM
98. Return to Previous Menu
99. Exit from Menus

1>
```

7.5.1.1 Change Privileged Access Password

Set or change the Privileged Access Password. It provides you with the capability to access all Service Processor functions. This password is usually used by the system administrator or root user.

7.5.1.2 Change General Access Password

Set or change the General Access Password. It provides limited access to SP menus, and is usually available to all users who are allowed to power-on the server.

7.5.1.3 Enable/Disable Console Mirroring

When Console Mirroring is enabled, the Service Processor sends information to both serial ports. This capability may be enabled by local or remote users. This provides local users the capability to monitor remote sessions. Console mirroring may be enabled for the current session only and ends when the SP releases control of the serial ports to the system firmware.

7.5.1.4 Start Talk Mode

In a console mirroring session, it is useful for those that are monitoring the session to be able to communicate with each other. Selecting this menu item activates the keyboards and displays for such communications while console mirroring is established. This is a full duplex link, so message interference is possible. Alternating messages between users works best.

7.5.1.5 OS Surveillance Setup Menu

This option may be used to setup operating system surveillance. If the surveillance is currently disabled you have to enable it before you can

change the surveillance time interval or the surveillance delay, as shown in the following example.

```
OS Surveillance Setup Menu

1. Surveillance:
   Currently Enabled

2. Surveillance Time Interval:
   Currently 5

3. Surveillance Delay:
   Currently 1

98. Return to Previous Menu

1>
```

Surveillance

May be set to `Enabled` or `Disabled`. You can switch between those two options by simply entering **1** and then press **Enter**.

Surveillance Time Interval:

May be set to any number from 2 to 255 minutes.

Surveillance Delay:

May be set to any number from 0 to 255 minutes.

Surveillance is a function in which the SP monitors the system, and the system monitors the Service Processor. This monitoring is accomplished by periodic samplings called *heartbeats*.

Surveillance is available during two phases:

1. System firmware bringup (automatic)
Provides the SP with a means to detect boot failures while the system firmware is running. System firmware surveillance is automatically enabled during system power-on. It cannot be disabled with a user selectable option. If the SP detects no heartbeats during system IPL (for 7 minutes), it cycles the system power to attempt a reboot. The maximum number of retries is set from the SP menus. If the fail condition repeats, the SP leaves the machine powered on, logs an error and offers menus to the user. If Call-out is enabled, the SP calls to report the failure and displays the operating system surveillance failure code on the operator panel.
2. Operating system run time (optional)

Provides the SP with a means to detect hang conditions, hardware or software failures while the operating system is running. It also provides the operating system with a means to detect the SP failure by the lack of a return heartbeat. Operating system surveillance is enabled by default. This allows you to run operating systems that do not support this Service Processor option. Operating system surveillance can be enabled and disabled through:

- Service Processor Menus
- Service Processor Service Aids

Three parameters must be set for operating system surveillance:

- Surveillance enable/disable
- Surveillance interval

This is the maximum time in minutes the Service Processor should wait for a heartbeat from the operating system before timeout.

- Surveillance delay

This is the length of time in minutes for the Service Processor to wait from when the operating system is started to when the first heartbeat is expected.

Surveillance takes effect immediately after setting the parameters from the SP menus. If operating system surveillance is enabled (and system firmware has passed control to the operating system), and the Service Processor does not detect any heartbeats from the operating system, the Service Processor assumes the system is hung. The machine is left powered on and the Service Processor enters standby phase, displaying the operating system surveillance failure code on the operator panel. If Call-out is enabled, the Service Processor calls to report the failure.

7.5.1.6 Reset Service Processor

Allows you to reinitialize the Service Processor.

7.5.1.7 Reprogram Service Processor Flash EPROM

This is an automatic process. An update diskette can be acquired from the RS/6000 Support page on the Internet or from your service team. The Internet address is:

<http://www.rs6000.ibm.com/support/micro>

For more information how to update the SP firmware refer to the README file that comes with the firmware update and 7.7, "Service Processor Firmware Update" on page 210.

7.5.2 System Power Control Menu

The following shows a system power control menu.

```
SYSTEM POWER CONTROL MENU

1. Enable/Disable Unattended Start Mode:
   Currently Disabled
2. Ring Indicate Power-On Menu
3. Reboot/Restart Policy Setup Menu
4. Power-On System
5. Power-Off System
6. Enable/Disable Fast System Boot:
   Currently Disabled
98. Return to Previous Menu
99. Exit from Menus

1>
```

7.5.2.1 Enable/Disable Unattended Start Mode

This option may be used to instruct Service Processor to immediately power-on the server after a power failure, bypassing power-on password verification. Unattended Start Mode can also be set using SMS Menus. It is intended to be used on servers that require automatic power-on after a power failure.

7.5.2.2 Ring Indicate Power-On Menu

Ring Indicate Power-On is an alternate method of dialing in, without establishing a Service Processor session. If the system is powered off, Ring Indicate Power-On is enabled, the server is powered on at the predetermined number of rings. If the server is already on, no action is taken. In either case, the telephone call is not answered. The caller receives no feedback that the server powered-on. The Ring Indicator Power-On Menu and defaults are shown in the following example:

```
Ring Indicator Power-On Menu

1. Ring Indicate Power-On:
   Currently Disabled
2. Number of Rings:
   Currently 6

98. Return to Previous Menu

1>
```

- Ring Indicate Power-On may be set to *Enabled* or *Disabled*.
- Number of Rings may be set to any number from 1 to 255.

7.5.2.3 Reboot/Restart Policy Setup Menu

Reboot describes bringing the system hardware back up from scratch, for example, from a system reset or power on. The boot process ends when control passes to the operating system process.

Restart describes activating the operating system after the system hardware reinitialized. Restart must follow a successful reboot.

The following is an example of the Reboot/Restart Policy Setup Menu.

```
Reboot/Restart Policy Setup Menu

1. Number of reboot attempts:
   Currently 3

2. Use OS-Defined restart policy?
   Currently Yes

3. Enable supplemental restart policy?
   Currently No

4. Call-Out before restart:
   Currently Disabled

98. Return to Previous Menu

1>
```

Number of Reboot Attempts

If the server fails to successfully complete the boot process, it attempts to reboot the number of times specified. Entry values equal to or greater than 0 are valid. Only successive failed reboots attempts count, not reboots that occur after a restart attempt. At restart, the counter is set to 0.

Use OS-Defined Restart Policy

This lets the SP react or not react the same as the operating system to major system faults, by reading the setting of the operating system parameter Automatically Restart/Reboot After a System Crash. This parameter may, or may not be defined depending on the operating system or its version/level. If the operating system automatic restart setting is defined, then it may be set to respond to a major fault by restarting or by not restarting. See your operating system documentation for details on setting up operating systems automatic restarts. The default value is *YES*.

Enable Supplemental Restart Policy

The default setting is *NO*. If set to *YES*, the SP restarts the system when the system loses control as detected by the SP surveillance, and either:

1. The Use OS-Defined restart policy is set to `NO`.

OR

2. The Use OS-Defined restart policy is set to `YES` and the operating system has *no* automatic restart policy.

Call-Out Before Restart (Enabled/Disabled)

If a restart is necessary due to a system fault, you can enable the SP to call out and report the event. This item is valuable if the number of these events becomes excessive, signalling a bigger problem.

7.5.2.4 Power-On System

Lets you power-on the system immediately.

7.5.2.5 Power-Off System

Allows you to power-off the server following a surveillance failure.

7.5.2.6 Enable/Disable Fast System Boot

Allows you to perform a fast boot. With disabling the fast boot the system needs a longer time in order to do a much more detailed BIST. You will see the E0B0 code with a spinning cursor for a longer time on the LED.

7.5.3 System Information Menu

The following is a sample of the system information menu.

```
SYSTEM INFORMATION MENU

1. Read VPD Image from Last System Boot
2. Read Progress Indicators from Last System Boot
3. Read Service Processor Error Logs
4. Read System POST Errors
5. Read NVRAM
6. Read Service Processor Configuration
7. View System Environmental Conditions
8. Processor Configuration/Deconfiguration Menu
9. Memory Configuration/Deconfiguration Menu
98. Return to Previous Menu
99. Exit from Menus

1>
```

7.5.3.1 Read VPD Image from Last System Boot

Displays the VPD information that was in effect after the last system boot. This information is usually identical with the results from the menu selection `Read VPD`, but in the case of configuration updates or certain faults, this

historical comparison can be useful to System Administrators and service personnel.

7.5.3.2 Read Progress Indicators from Last System Boot

Displays the boot progress indicators (check points), up to a maximum of 100, from the system boot prior to the one in progress now. This historical information may be useful to help diagnose system faults.

7.5.3.3 Read Service Processor Error Logs

Displays error conditions detected by the Service Processor.

7.5.3.4 Read System POST Errors

Selecting this item lets you review the results of the POST (Power-On Self Test). Your system unit may be able to start in the presence of POST errors if there is sufficient working system resources.

7.5.3.5 Read NVRAM

Displays non-volatile random access memory (NVRAM) contents. The SP will prompt you for an address. Enter 0 as a start address or any other.

7.5.3.6 Read Service Processor Configuration

The following menu shows you a summary of all SP settings. You are not able to change anything in these three screens. To change the settings you have to go into the different menus.

```
Current Service Processor Configuration

SYSTEM POWER CONTROL MENU
1. Enable/Disable Unattended Start Mode: Currently Disabled
6. Enable/Disable Fast System Boot: Currently Enabled

OS Surveillance Setup Menu
1. Surveillance: Currently Disabled
2. Surveillance Time Interval: Not Applicable
3. Surveillance Delay: Not Applicable

Ring Indicate Power-On Menu
1. Ring indicate power-on: Currently Disabled
2. Number of rings: Currently 6

Reboot/Restart Policy Setup Menu
1. Number of reboot attempts: Currently 3
2. Use OS-Defined restart policy? Currently Yes
3. Enable supplemental restart policy? Currently No
4. Call-Out before restart: Currently Disabled

(Press Return to Continue)
```


Press **Enter** to continue and you will see the second screen.

```
Current Service Processor Configuration

Serial Port Selection Menu
1. Serial Port 1 Call-Out: Currently Disabled
2. Serial Port 2 Call-Out: Currently Disabled
3. Serial Port 1 Call-In: Currently Disabled
4. Serial Port 2 Call-In: Currently Disabled

Call-Out Policy Setup Menu
1. Call-Out policy (First/All): Currently First
2. Remote timeout, (in seconds): Currently 120
3. Remote latency, (in seconds): Currently 2
4. Number of retires: Currently 2

Telephone Number Setup Menu
1. Service Center Telephone Number: Currently Unassigned
2. Customer Administration Center Telephone Number: Currently Unassigned
3. Digital Pager Number: Currently Unassigned
4. Customer Voice Telephone Number: Currently Unassigned
5. Customer System Telephone Number: Currently Unassigned

(Press Return to Continue)
```

Now press **Enter** to see the last screen.

```
Current Service Processor Configuration

Customer Account Setup Menu
1. Customer Account Number: Currently Unassigned
2. Customer RETAIN login userid: Currently Unassigned
3. Customer RETAIN login password: Currently Unassigned

Serial Port Speed Setup Menu
1. Serial port 1 speed: Currently 9600
2. Serial port 2 speed: Currently 9600

Modem Configuration Menu
Port 1 Modem Configuration File Name:
Port 2 Modem Configuration File Name:

SERVICE PROCESSOR SETUP MENU
3. Enable/Disable Console Mirroring: Currently Disabled

(Press Return to Continue)
```

7.5.3.7 View System Environmental Conditions

The SP reads all environmental sensors and reports the results. This option is most useful when surveillance fails, since it allows you to determine the environmental conditions that may be related to the failure. The following is

an example of the system environmental conditions screen and depends on your system:

```
System Environmental Conditions
(System Power is currently off.)

Fan 0: A stopped fan detected
Fan 1: A stopped fan detected
Fan 2: A stopped fan detected
Fan 3: A stopped fan detected
MEM Temperature is operating within normal tolerances
I/O Temperature is operating within normal tolerances
CPU Temperature is operating within normal tolerances
5.0 Volts: A low 5.0 Voltage reading detected
3.3 Volts: A low 3.3 Voltage reading detected
5.0 Standby Volts: is operating within normal tolerance
+12.0 Volts: A low +12.0 voltage reading detected
-12.0 Volts: A high -12.0 voltage reading detected
(Press Return to Continue)
```

7.5.3.8 Processor Configuration/Deconfiguration Menu

This option allows you to see and modify processor configuration. The following is an example of the Processor Configuration/Deconfiguration Menu:

```
PROCESSOR CONFIGURATION/DECONFIGURATION MENU

Processor Number

0. Configured by system (0x11)
2. Configured by system (0x11)

98. Return to Previous Menu

To change the configuration, select the processor number
1>
```

You can manually configure or deconfigure any processor, regardless of failure status, through this SP menu. The configuration process takes place during the system power-up. Therefore, the configuration displayed in STANDBY mode reflects the configuration during the last boot. To view the current configuration, access the SP menu after the system boots up. When you select a processor, its state toggles between configured and deconfigured.

Processors that are not present are not listed. A processor can be in any of the following four states:

Configured by System	The processor is present, and has not exceeded the number of failures threshold. It is configured by the system and is available.
Deconfigured by System	The processor is present, but has exceeded the number of failures threshold. It is deconfigured by the system and is currently unavailable.
Manually configured	The processor is present and available. It is configured by the user through the Service Processor menus.
Manually deconfigured	The processor is present, but unavailable. It has been deconfigured by the user through the Service Processor menus.

Note

You can also disable the second processor with the following steps:

1. Power-on the system and go into the Open Firmware mode (see 6.7.1, “SMS Start-Up” on page 173.)
2. After the `ok>` prompt appeared enter: **boot -scpu=1**

This will automatically disable the second processor and cause the system to continue the boot process. After a reboot the second processor will be automatically enabled by the system.

7.5.3.9 Memory Configuration/Deconfiguration Menu

This option allows you to see and modify memory configuration. The following is an example of the Processor Configuration/Deconfiguration Menu:

```

MEMORY CONFIGURATION/DECONFIGURATION MENU

DIMMs on Memory Card #0:

DIMMs on Memory Card #1:
  1. Configured by system (0x0)

98. Return to Previous Menu
Enter card number _AND_ DIMM number separated by a space.
1>

```

You can manually configure or deconfigure any memory DIMM, regardless of failure status, through this SP menu. The configuration process takes place during the system power-up. Therefore, the configuration displayed in

STANDBY mode reflects the configuration during the last boot. To view the current configuration, access the SP menu after the system boots up. When you select a memory DIMM, its state will toggle between configured and deconfigured. Memory DIMMs that are not present are not listed. A memory DIMM can be in any of the following four states:

Configured by System	The memory DIMM is present, and has not exceeded the number of failures threshold. It is configured by the system and is available.
Deconfigured by System	The memory DIMM is present, but has exceeded the number of failures threshold. It is deconfigured by the system and is currently unavailable.
Manually configured	The memory DIMM is present and available. It is configured by the user through the Service Processor menus.
Manually deconfigured	The memory DIMM is present, but unavailable. It has been deconfigured by the user through the Service Processor menus.

7.5.4 Language Selection Menu

The following is an example of the Language Selection Menu.

```
LANGUAGE SELECTION MENU
1. English
2. Francais
3. Deutsch
4. Italiano
5. Espanol
6. Svenska
98. Return to Previous Menu
99. Exit from Menus
1>
```

Note: Your ASCII terminal must support the ISO-8859 character set in order to properly display languages other than English. This menu allows selecting languages into which Service Processor and System Firmware menus and messages are displayed.

7.5.5 CALL-IN/CALL-OUT Setup Menu

The Following is an example of a Call-In/Call-Out Setup Menu.

CALL-IN/CALL-OUT SETUP MENU

1. Modem Configuration Menu
2. Serial Port Selection Menu
3. Serial Port Speed Setup Menu
4. Telephone Number Setup Menu
5. Call-Out Policy Setup Menu
6. Customer Account Setup Menu
7. Call-Out Test
8. Ring Indicate Power-On Menu
98. Return to Previous Menu
99. Exit from Menus

1>

7.5.5.1 Modem Configuration Menu

The first two lines of the Modem Configuration Menu (shown following) are status lines showing the current selections. Selections are made in the two sections labeled Modem Ports and Modem Configuration File Name. Select the serial port that you want to activate and then select the modem configuration file for the modem on the port. If you wish to set up both serial ports with modems, make your selections one port at a time.

Modem Configuration Menu

Port 1 Modem Configuration File Name:
Port 2 Modem Configuration File Name:

To make changes, First select the port and then the configuration file name

Modem Ports:

1. Serial port 1
2. Serial port 2

Modem Configuration File Name:

3. none
4. modem_f_sp
5. modem_f0_sp
6. modem_f1_sp
7. modem_z_sp
8. modem_z0_sp
9. modem_m0_sp
10. modem_m1_sp

30. Save configuration to NVRAM and Configure modem
98. Return to Previous Menu

1>

To configure the modem, you first have to choose the serial port on which your modem is connected to the system. Press **1** or **2** for the appropriate

serial port. The SP will mark your choice with an arrow (Port 2 was used in the example for on the following screen).

```
Modem Configuration Menu

Port 1 Modem Configuration File Name:
--> Port 2 Modem Configuration File Name:

To make changes, First select the port and then the configuration
file name

Modem Ports:
 1. Serial port 1
 2. Serial port 2

Modem Configuration File Name:
 3. none
 4. modem_f_sp          9. modem_m0_sp
 5. modem_f0_sp         10. modem_m1_sp
 6. modem_f1_sp
 7. modem_z_sp
 8. modem_z0_sp

30. Save configuration to NVRAM and Configure modem
98. Return to Previous Menu

1>
```

Then you have to select a modem type. For example, if you have a Hayes Smart Modem 300 you have to select option 7 `modem_z_sp`. For more information about modem configurations and customization, please see the *RS/6000 7043 Model 260 User's Guide, SA38-0553* or the *RS/6000 7043 Model 260 Service Guide, SA38-0554*.

You will see your choices selected in the following screen:

```
Modem Configuration Menu

Port 1 Modem Configuration File Name:
--> Port 2 Modem Configuration File Name: modem_z_sp

To make changes, First select the port and then the configuration
file name

Modem Ports:
  1. Serial port 1
  2. Serial port 2

Modem Configuration File Name:
  3. none
  4. modem_f_sp          9. modem_m0_sp
  5. modem_f0_sp         10. modem_ml_sp
  6. modem_f1_sp
  7. modem_z_sp
  8. modem_z0_sp

 30. Save configuration to NVRAM and Configure modem
 98. Return to Previous Menu

1>
```

7.5.5.2 Serial Port Selection Menu

In any combination, the call-in and call-out functions of each serial port can be enabled and/or disabled, as shown in the following screen. The call-in for serial port 1 or 2 must be disabled if you want to activate the ring indicate power-on functionality. You cannot activate both functions.

```
Serial Port Selection Menu

 1. Serial Port 1 Call-Out:
    Currently Disabled

 2. Serial Port 2 Call-Out:
    Currently Disabled

 3. Serial Port 1 Call-In:
    Currently Disabled

 4. Serial Port 2 Call-In:
    Currently Disabled

98. Return to Previous Menu

1>
```

7.5.5.3 Serial Port Speed Setup Menu

The Serial Port Speed Setup Menu appears as follows:

```
Serial Port Speed Setup Menu

1. Serial Port 1 Speed:
   Currently 9600

2. Serial Port 2 Speed:
   Currently 9600

98. Return to Previous Menu

1>
```

The serial port speed can be set for terminal performance or to accommodate modem capabilities. A speed of 9600 baud or higher is recommended. Valid serial port speeds are shown below:

Table 39. Valid Serial Port Speeds

50	75	110	134	150	300	600	1200
1800	2000	2400	3600	4800	7200	9600	19200
38000	57600	115200					

7.5.5.4 Telephone Number Setup Menu

The Telephone Number Setup Menu appears as follows:

```
Telephone Number Setup Menu

1. Service Center Telephone Number:
   Currently Unassigned

2. Customer Administration Center Telephone Number:
   Currently Unassigned

3. Digital Pager Telephone Number:
   Currently Unassigned

4. Customer Voice Telephone Number:
   Currently Unassigned

5. Customer System Telephone Number:
   Currently Unassigned

98. Return to Previous Menu

1>
```


This menu may be used to set or change the telephone numbers for reporting a system failure. The SP allows setting or changing telephone numbers for:

Service Center Telephone Number

The number of the service center computer. The service center usually includes a computer that takes calls from servers with call-out capability. This computer is referred to as the *catcher*. The catcher expects messages in a specific format to which Service Processor conforms. For more information about the format and catcher computers, refer to the README file in the AIX `/usr/samples/syscatch` directory.

Customer Administration Center Telephone Number

The number of the System Administration Center computer (catcher) that receives problem calls from servers.

Digital Pager Telephone Number

The number for a pager carried by someone who responds to problem calls from your server.

Customer System Telephone Number

The telephone number of a phone near the server or answered by someone responsible for the server. This is the telephone number left on the pager for callback.

7.5.5.5 Call-Out Policy Setup Menu

The Call-Out Policy Setup Menu appears as follows.

```
CALL-OUT POLICY SETUP MENU

1. Call-Out policy (First/All):
   Currently First

2. Remote timeout, (in seconds):
   Currently 120

3. Remote latency, (in seconds):
   Currently 2

4. Number of retries:
   Currently 2

98. Return to Previous Menu

1>
```

Call-Out Policy (First/All)

The Call Out policy may be set to first or all. If the call out policy is set to first, the SP stops at the first successful call out to one of the following numbers in the order listed:

1. Service Center Customer
2. Admin Center
3. Pager

If the call out policy is set to all, the Service Processor attempts a call out to all the following numbers in the order listed:

1. Service Center
2. Customer Admin Center
3. Pager

Remote Timeout, (in Seconds) and Remote Latency, (in seconds)

Remote timeout and Remote latency are functions of your service provider's catcher computer. You should take the defaults or contact your service provider for recommended settings.

Number of Retries

Number of retries is the number of times you want the server to retry calls that resulted in busy signals or other error messages.

7.5.5.6 Customer Account Setup Menu

The following is an example of a Customer Account Setup Menu.

```
Customer Account Setup Menu
1. Customer Account Number:
   Currently Unassigned
2. Customer RETAIN Login UserID:
   Currently Unassigned
3. Customer RETAIN Login Password:
   Currently Unassigned
98. Return to Previous Menu
1>
```

Customer Account Number

This number is assigned by your service provider for record keeping and billing. If you have an account number, enter it here. Otherwise, leave this field blank.

Customer RETAIN Login User ID & Customer RETAIN Login Password

Both apply to a service function to which your service provider may or may not have access. Leave these fields blank if your service provider does not use RETAIN.

7.5.5.7 Call-Out Test

Allows you to test the call-out functionality, which has been set in the telephone number setup menu (see 7.5.5.4, “Telephone Number Setup Menu” on page 206).

7.5.5.8 Ring Indicate Power-On Menu

Please refer to 7.5.2.2, “Ring Indicate Power-On Menu” on page 195. This is an additional possibility to activate ring indicate power-on.

7.5.6 Set System Name

The System Name, an optional field, is the name your server reports in problem messages. This name helps your support team, (for example, your system administrator, network administrator, or service representative) to more quickly identify the location, configuration, and history of your server.

7.6 General User Menu

The menu options presented to the General user are a subset of the options available to the Privileged user, as shown in the following screen. You must know the General Access Password in order to access these menus (if set).

```
GENERAL USER MENU

1. Power-On System
2. Read VPD Image from Last System Boot
3. Read Progress Indicators from Last System Boot
4. Read Service Processor Error Logs
5. Read System POST Errors
6. View System Environmental Conditions
99. Exit from Menus

1>
```

Note: The Service Processor prompt reads either 1> or 2> to indicate which serial port on the system unit is being used to communicate with the SP.

7.6.1 Power-On System

Allows you to power-on the system.

7.6.2 Read VPD Image from Last System Boot

Displays the Vital Product Data (VPD) information that was in effect after the last system boot. This information is usually identical with the results from the menu selection Read VPD, but in the case of configuration updates or certain faults, this historical comparison can be useful to System Administrators and service personnel.

7.6.3 Read Progress Indicators from Last System Boot

Displays the boot progress indicators (check points), up to a maximum of 100, from the system boot prior to the one in progress now. This historical information may be useful to help diagnose system faults.

7.6.4 Read Service Processor Error Logs

Displays the Service Processor error logs.

7.6.5 Read System POST Errors

Selecting this item lets you review the results of the POST (Power-On Self-Test). Your server may be able to start in the presence of POST errors if there is sufficient working system resources. If POST errors occur during start-up, this error log helps to isolate faults when used with the diagnostics.

7.6.6 View System Environmental Conditions

On selection of this menu, the Service Processor reads all environmental sensors and reports the results. This option maybe useful when surveillance fails, since it allows you to determine the environmental conditions that may be related to the failure.

7.7 Service Processor Firmware Update

The Service Processor EPROM may need to be updated for two different reasons:

1. The UPDATE (composite) portion of the EPROM has become corrupted.
2. The Service Processor firmware upgrades, without any corruption present.

The use of a Flash EPROM allows updates to occur without physically replacing the memory.

The firmware in your server can be updated using one of four available initiation processes:

1. Service Processor initiation
2. SMS Utilities initiation
3. Service Aids initiation
4. AIX initiation

Each initiation method is described below. In each case, the process prompts you for your authority and shows the contents of the update media. Verify the file with which to perform the update, and follow any other instructions that may appear. After initiation, the processes are identical and automatic.

There are two areas in each firmware module that may need updating:

1. The gold code or base code or EPROM area
2. The custom or main program or FLASH area

Each update file contains matching gold and custom firmware, so it is not possible to update to a conflicting set.

Before the update process begins, the versions of each of the two areas of the target firmware module are compared to the versions on the update diskette. Only the area(s) that need updating are updated. In most cases, only the custom area is updated.

An update diskette can be acquired from the RS/6000 Support page on the Internet or from your service team. The Internet address is:

<http://www.rs6000.ibm.com/support/micro>

The update files must be downloaded to a DOS-formatted diskette. You should format the diskette just prior to receiving the download, to be sure it is clean. Reformatting is worth the time and effort.

7.7.1 Updating Firmware from the Service Processor Menus

The Service Processor Menus are available while the server is powered off. From the Service Processor Main Menu, select **Service Processor Setup** to bring up a menu containing the item, Reprogram Service Processor Flash EPROM. Selecting that item starts the update process. The process requests the update diskette(s) as needed. After inserting the first diskette and informing the system, the process continues automatically. If the process requires user assistance, it is requested.

7.7.2 Updating Firmware from the SMS Utilities

From a powered-down or reset condition, bring up the SMS Utilities and select the menu item for updating firmware. The process requests the update diskette(s) as needed. After inserting the first diskette and informing the system, the process continues automatically. If the process requires user assistance, it is requested.

Refer also to 6.7.4.4, "How to Perform the SMS Firmware Update" on page 182 and the README file that comes with the firmware update.

7.7.3 Updating Firmware from the Service Aids

Note: This service aid is only supported for On-line Diagnostics.

This service aid allows you to update the system or service processor flash. Additional update and recovery instructions may be provided. You also need to know the fully qualified path and file name of the flash update image file. If the flash update image file is on a diskette, you need the AIX DOS utilities package to process the diskette. You can use the `dosdir` command to find out the name of the update image file. This service aid uses the `dosread` command to put the file into the `/var` file system.

Using this service aid to update the system or service processor flash does not save the current contents of the flash. The flash image file will be copied in the `/var` file system. If there is not enough space in the `/var` file system for the flash update image file, an error is reported. If this occurs, exit the service aid, increase the size of the `/var` file system and retry the service aid. After the file is copied, a warning screen asks for confirmation to continue the flash update. Continuing the flash update eventually causes the system to reboot and return to the AIX login prompt. After the system unit reboots, the file `/var/update_flash_image` can be removed to conserve space in the `/var` file system.

7.7.4 Updating Firmware from AIX

You must delete the file `/var/update_flash_image` before proceeding.

The flash update image file must have already been placed in the `/var` file system. This could have been done with a file transfer from another server or with the `dosread` command of the AIX DOS Utilities, for example. With the flash update image in place, issuing the following AIX command `shutdown -u /var/<filename>\.img` initiates the update process. Where `<filename>` is the name of the flash update image. During the process, the server powers down

and reboots. You know the process is complete when the login prompt reappears.

7.8 Service Processor Checkpoints

SP Checkpoints are intended to let users and service personal know what the service processor is doing, in some detail, as it initializes. These checkpoints are not intended to be error indicators, but in some cases a server could hang at one of the checkpoints without displaying an 8-character error code. It is for these hang conditions, only, that any action should be taken with respect to checkpoints. The most appropriate action is included with each checkpoint.

Before taking actions listed with a checkpoint, it is a good idea to look for better symptoms in the service processor error log.

Service processor checkpoints are in the range E010 to E0FF. The message OK indicates successful service processor testing and initialization. The E0xx-Codes for the SP are documented in *RS/6000 7043 Model 260 Service Guide*, SA38-0554.

Chapter 8. Hints and Tips

This chapter provides useful hints and tips, not only for the Models 150 and 260, but for all RS/6000 models.

8.1 Electronic Key-Switch Function

One of the main differences between Micro Channel-based RS/6000 systems and PCI-based RS/6000 systems is the omission of a physical key. On all Micro Channel-based RS/6000 systems, there is a physical key that can be in one of three positions:

- Normal
- Service
- Secure

Although the Models 150 and 260 (and also the other PCI-based RS/6000 models) do not have this physical key, they have an electronic key-switch instead. Currently, this provides the ability to boot the machine in Normal or Service mode. The electronic key-switch is a logical key-switch, separate from a physical one, that the platform firmware knows about. Its state is communicated from the firmware to AIX through a residual data environment variable. Currently, only the firmware can set the electronic key-switch.

8.1.1 `bootlist` Command

The `bootlist` command can be used to alter or display the boot list. For example, you can set the customized boot list with the following command:

```
# bootlist -m normal hdisk0
```

This would set the customized boot list in the NVRAM to `hdisk0. hdisk0` would then be used to attempt a boot in Normal mode.

The service bootlist can be determined with the command:

```
# bootlist -m service -o
```

The output is as follows and depends on your settings:

```
fd0  
cd0  
rmt0  
hdisk0  
ent0
```

There are two other ways to alter the boot lists: System Management Services (SMS) and the AIX tool `diag`.

8.2 Entering Debug Mode

On a Micro Channel-based RS/6000 system, you can enter debug mode, `lldb`, by switching the key to Service and pressing one of the following key combinations:

- `ctl-alt-numpad4` (on a graphical console)
- `CTL-\` (on an ASCII display)

The `CTL-\` key combination is valid on most OEM terminals. On some IBM 3151 terminals, using a cartridge, you may have to press `CTL-NUMPAD4`.

Since the Models 150 and 260 have a static electronic key that is set to Normal during the normal operation of AIX, the above key sequences will not work unless the debugger, `lldb`, is pinned. You can do this by entering the following command and rebooting your machine:

```
# bosboot -D -a -d /dev/hdiskx
```

Where `hdiskx` is the disk containing the boot logical volume. After rebooting your machine, using either of the two key sequences shown above will invoke the debugger on an ASCII terminal (although the debugger can be invoked with a graphical display keyboard, an ASCII terminal must be attached to use the debugger). Be aware that anyone using either a graphical or ASCII console will be able to make the machine unavailable to other users.

8.3 Backup and Cloning

This section describes how to create system backups that will be restored on the same or different machines (cloning).

8.3.1 The `bootinfo` Command

The type of machine (platform type) being used can be determined with the `bootinfo` command. This command uses the device configuration databases to obtain information. In addition, the `bootinfo` command uses the values that are stored in NVRAM.

To determine the platform of a running machine, use the `bootinfo -p` command if the machine is running AIX Version 4.2 or later. If the machine is running AIX Version 4.1, use the `bootinfo -T` command. However, with

`bootinfo -p`, you will get different results for PReP and CHRP/RPA machines. The result of `bootinfo -T` is always `rspc` for every PCI-based machine.

The results of `bootinfo` with the `-p` option are shown in Table 40.

Table 40. Results of `bootinfo -p`

Platform	Description
<code>chrp</code>	CHRP-compliant PCI-based RS/6000 machines
<code>rspc</code>	PReP-compliant PCI-based RS/6000 machines
<code>rs6k</code>	Micro Channel-based RS/6000 machines

8.3.2 Creating a System Backup

Creating a system backup, or `mksysb`, is very simple. If you were using a tape drive, for example `rmt0`, you could easily create a system backup from the command line by typing:

```
# mksysb -i /dev/rmt0
```

This would create a bootable image of your system that can be used to either re-create the system or to restore files.

Alternatively, you can use SMIT to create your system image. To do this, you can use the SMIT fastpath:

```
# smit mksysb
```

Note

Creating a system backup with the `mksysb` command only saves data in the `rootvg` volume group. If you want to also save data from a non-root volume group you have to use the `savevg` command or the SMIT fastpath:

```
# smit savevg
```

8.3.3 Restoring Your System Backup

You can restore your `mksysb` by booting from the tape that you have created and following the instructions that are displayed. To boot from the tape, you have to press **F1** (or **1** on an ASCII terminal) during the firmware boot, and select the tape device to boot from. Alternatively, you can place the tape drive as the first device in the customized boot list so that the tape will automatically boot when the machine is powered-on or rebooted.

When the installation of the backup has finished, the installation program reconfigures the Object Data Manager (ODM) on the target system. If the target system does not have exactly the same hardware configuration as the source system, the program may modify device attributes in the following target system files:

- All files in the `/etc/objrepos` file beginning with the letters CU
- All files in the `/dev` directory

The data restored from the `mksysb`-tape will only be the data from the `rootvg` volume group. If you want to restore data from a non-root volume group you have to run the `restorevg` command or use SMIT fastpath `smit restvg`.

Although the `mksysb` command will create a bootable tape on the PCI-based RS/6000 servers, the tapes will not boot on Micro Channel-based uni-processor or Micro Channel-based symmetric multiprocessor machines. Additionally, `mksysb` tapes created on Micro Channel-based uni-processor and Micro Channel-based symmetric multiprocessor machines will not boot on PCI-based machines.

8.3.4 Cloning

This section also covers AIX 4.1 since there are many customers still using AIX 4.1.

The term cloning describes the installation of a system backup (created with the `mksysb` command on one system) on different systems in order to make the software configuration of all systems as identical as possible. Using a backup of one machine to install several other machines can be very useful and will save the time taken to perform multiple installations.

8.3.4.1 Cloning on AIX 4.1

AIX 4.1 has been packaged so that only the devices that are needed are actually installed on the system. This has made the cloning process difficult.

Note

When cloning systems, please keep in mind that you will also clone IP addresses. If both machines are connected to the network at the same time, this will cause problems on the network. In this case, you have to comment the `tcpip` entry in the `/etc/inittab` file.

In AIX 4.1, the devices is separated from the base in separate filesets, `devices.*` (such as `devices.scsi.disk`). This, coupled with all the new hardware

that has been released in the last few years (RSPC, SMP, graphics adapters, disks, and so on) has made cloning machines almost impossible. Up to this point the only method of cloning was to install all the device drivers on a machine before you created a mksysb image. If you have a limited disk space you may run out of space. However, it is recommended to install most of the available device drivers in order to avoid cloning problems.

If you want to clone for example a PCI-based machine to a Micro Channel (MCA) based machine and you don't have the MCA device drivers installed on your mksysb-tape, there is a way to do a clone, but you still need the product media at the same level as your mksysb. This procedure will be very important not only for cloning systems, but also for developers to create a snapshot of the system in order to reproduce the environment in which a problem was detected.

In AIX Version 4.1, there are two kinds of platforms, *rs6k* and *rspc*, and two different kernels for uniprocessor (UP) and multiprocessor (MP). The *rs6k* boot images differ from the *rspc* boot images, and one will not boot the other. The *rspc* boot image has some specific boot information at the front of the image that the *rs6k* does not understand. So, to avoid this boot image mismatch problem, we advise using the product media, preferably a CD-ROM, because it contains both boot images and can boot any system.

To determine what platform type your machine is, enter:

```
# bootinfo -T          rsk6 or rspc
```

To determine what kernel your system uses, enter:

```
# bootinfo -z          0 = UP 1 = MP
```

This method allows you to clone with physical media, meaning with CD-ROM product media, a mksysb tape and a customized diskette.

If you are installing a UP mksysb to an MP system with these levels, and you are using, for example, the base 4.1.4.0 media, then the `bos.rte.mp` files set that will be installed by the cloning script will only be level 4.1.4.0. When the system reboots, you may not be able to log in. The error you will receive will be Failed setting user credentials. To fix this problem, you will need to boot into maintenance mode from the 4.1.4.0 media and update the `bos.rte.mp` files set to at least the 4.1.4.6 level. The libs and kernel will then be in sync and you will be able to log in.

So, if you can check the level of the `bos.rte.up` kernel before you create your mksysb (if you are going to clone to an MP system), you can install the MP kernel and its updates before you create your mksysb. If you are not in that

position, you may just need to boot into maintenance mode after the install and install the MP kernel updates.

8.3.4.2 On the System You Want to Clone

1. Create a mksysb tape on the system you want to clone.

Note: If your system was created from D5 install media, it can only be restored to another system that supports D5 code. To prevent problems with code levels, make sure that your system is installed from G5 media prior to creating the mksysb. Then use a G5 install CD-ROM for the steps later.

2. Create a customized diskette.

```
# mkdir /tmp/clone
# cd /tmp/clone
# echo data > signature
# cp /usr/lpp/bosinst/bosinst.template bosinst.data
# vi bosinst.data
```

Make sure that the `control_flow` stanza has this entry set. This is the most important part of the `bosinst.data` for this procedure:

```
CUSTOMIZATION_FILE = cloner
```

This will tell bos install to run the `cloner` script after it has restored the mksysb image.

Also, the `target_disk_data` stanza should be *zeroed* out as shown, so that the bos install program will install on the best fit disks.

Here is a sample `bosinst.data` file:

```
control_flow:
    CONSOLE =
    INSTALL_METHOD = overwrite
    PROMPT = yes
    EXISTING_SYSTEM_OVERWRITE = yes
    INSTALL_X_IF_ADAPTER = yes
    RUN_STARTUP = yes
    RM_INST_ROOTS = no
    ERROR_EXIT =
    CUSTOMIZATION_FILE = cloner
    TCB = no
    INSTALL_TYPE =
    BUNDLES =

target_disk_data:
    LOCATION =
```

```
SIZE_MB =
HDISKNAME =
```

```
locale:
  BOSINST_LANG = C
  CULTURAL_CONVENTION = C
  MESSAGES = en_US
  KEYBOARD = en_US
```

3. Create the customization script using an editor.

```
# vi cloner
```

Sample Script:

```
#!/usr/bin/ksh
set -x
RV=$(bootinfo -z)
if [ "$RV" -eq 1 ]
then
installp -abcgXd/../../SPOT/usr/sys/inst.images bos.rte.mp
ln -fs /usr/lib/boot/unix_mp /usr/lib/boot/unix
fi
if [ "$RV" -eq 0 ]
then
installp -abcgXd/../../SPOT/usr/sys/inst.images bos.rte.up
ln -fs /usr/lib/boot/unix_up /usr/lib/boot/unix
fi
devinstall -b -d ../../SPOT/usr/sys/inst.images -f ../../tmp/device.pkgs
cfgmgr -v -i ../../SPOT/usr/sys/inst.images
BLVDISK=`lslv -l hd5 | grep hdisk | head -1 | cut -d' ' -f1`
ln -f /dev/r$BLVDISK /dev/ipldevice
bosboot -a -d /dev/ipldevice

bootlist -m normal $BLVDISK
rm -f /etc/firstboot
sync
sync
sync
exit 0
```

4. Backup the three files to a diskette:

```
# find . -print | backup -ivqf/dev/rfd0
```

8.3.4.3 On the System You Want to Install

Insert the diskette and the AIX 4.1.4 CD and boot the target machine (the one you want to install on) from the CD product media and go into the installation main menu.

- On a PCI-based machine, just press **F5** when the icons appear on the screen. This will cause the default firmware boot list to be used.
- On a MCA-based machine, turn the key to the service mode and then power-on the machine.

Make sure the tape drive is turned on, but do not insert the tape yet since it might try to boot from the tape. Press the **F1** when the console screen comes up; then press **Enter** when the language screen comes up (to get English). When the following screen comes up, choose option **3. Start Maintenance Mode for System Recovery**.

```
                Welcome to Base Operating System
                Installation and Maintenance

Type the number of your choice and press ENTER.

1. Start Installation Now with Default Settings
2. Change/Show Installation Settings and Install
3. Start Maintenance Mode for System Recovery

88 Help ?
Choice:
```

Then choose option **4. Install from a System Backup** from the Maintenance Mode for the System Recovery submenu.

Put the mksysb tape in the tape drive and close the door. Select the tape drive from that menu and press **Enter**. You will then see the system access the tape drive and restore it from the diskette drive. Then the confirmation screen will come up. Choose **1** or **2**, press **Enter**, then press **Enter**, then press **Enter** again to confirm the install. If you are using AIX 4.1.3, you will see the language menu come up after the tape is read and the diskette is restored. Just press **Enter** for the default English, and you will get the confirmation install screen. Choose **0** and press **Enter** to start the install.

You can then turn the key (if there is a key) to normal at this time so the machine will reboot when the install is complete.

You should see the mksysb being restored, and then you should see `installp` called later when `cfgmgr` is called from the script to install any additional devices it detects.

Warning for AIX Version 4.1.4 or Prior

When installing a mksysb on a different platform type (for example, an rs6k mksysb to an rs6ksmp or an rs6k to an rspc), the first bosboot that BOS install attempts will fail, and a message will be printed out to ask if you want to do maintenance or continue. You should just continue at this point, and the cloner script will run. You will not see any output until the copyright screen at the end and the machine reboots. This step may last over half an hour.

This problem has been fixed in 4.1.5.0. The system will check to see if there is a cloner customization script in the bosinst.data file, and if there is, it will not stop when the bosboot fails. It will continue the install and run the cloner script, knowing that the cloner script will run the bosboot.

8.3.5 Cloning on AIX Version 4.2 and 4.3

The following sections describe how to clone disk images on AIX Version 4.2 and later.

8.3.5.1 What has been changed since AIX 4.1

The procedure for cloning a AIX 4.2 or 4.3 system has been greatly simplified. All you need is the AIX V4.2 or 4.3 CD-ROM and the AIX V4.2 or 4.3 mksysb tape that you wish to restore.

The mksysb images enable you to clone one system image onto multiple target systems. The target systems might not contain the same hardware devices or adapters, require the same kernel (uniprocessor or microprocessor), or be the same hardware platform (rs6k, rspc, or chrp) as the source system.

You do not need a separate diskette or allocated customization script to clone a system in 4.2 or 4.3. Cloning was incorporated into the main 4.2 BOS install program (`bi_main`). When booting from CD-ROM and installing a mksysb tape, `bi_main` will automatically determine if there are additional device filesets needed to be installed and install them from the CD-ROM. Also, the capability to boot from a product tape to install any additional device support is added.

Since AIX 4.2 there are a three new stanzas for the bosinst.data file:

- **SWITCH_TO_PRODUCT_TAPE**

This stanza allows you to boot from a product tape, then switch to a mksysb tape to install. You can then switch back to the product tape at the end of the installation if you need to install additional device filesets for support on the target machine. This procedure is usually used for cloning systems. The possible values are no (default) and yes.

Normally, you would boot from a CD-ROM (to ensure that you have the correct device support and boot image) then tell BOS install to install from the mksysb tape. At the end of the install, BOS install automatically verifies that all device support is installed on the system and installs additional device support from the CD-ROM.

If you are booting from a product tape and switching to a mksysb tape, BOS install does not automatically assume you want to do cloning and will not prompt you for the product tape again. If you want a prompt for the product tape to verify all the device support is installed, you must set this variable to yes in your bosinst.data file.

- **RECOVER_DEVICES**

Specifies whether to reconfigure the devices. For mksysb installs, the ODM configuration database is saved in the image. The device names and attributes are automatically extracted from the database, and the BOS install program attempts to recreate the devices the same way they were on the machine the mksysb was created on. This is normally what you would do for regular mksysb restores on the same system. However, for cloning, you may not want these devices configured this way, especially for network configuration. The possible values are yes (default) and no if you do not want device reconfiguration.

- **BOSINST_DEBUG**

Specifies whether to show debug output during BOS installation. The value yes will send set -x debug output to the screen during BOS installation. The possible values are no (default) and yes.

8.3.5.2 Cloning Your System

This procedure is used to install a mksysb on a target system it was not created on. Be sure to boot from the product media appropriate for your system and at the same maintenance level of BOS as the original system the mksysb was made on. For example, you can use BOS Version 4.2.1 product media with a mksysb from a BOS Version 4.2.1 system.

After booting from your chosen product media, complete the following steps when the following menu appears:

Welcome to Base Operating System
Installation and Maintenance

Type the number of your choice and press ENTER.

1. Start Installation Now with Default Settings
2. Change/Show Installation Settings and Install
3. Start Maintenance Mode for System Recovery

88 Help ?
Choice:

The Welcome to the Base Operating System Installation and Maintenance screen is displayed.

1. Select **3. Start Maintenance Mode for System Recovery** option from the installation main menu.
2. Select the **Install from a System Backup** option.
3. Select the drive containing the backup tape and insert the tape.

The system reads the tape and begins the installation.

You will be prompted again for the BOS install language, and the Welcome screen is displayed. Continue with the Prompted Installation. Cloning is not supported in a non-prompted installation.

4. If you are cloning from the product CD, do not remove the CD from the CD-ROM drive.

Notes:

- a. Booting from tape product media is not supported on rspc platform systems. You must use CD product media. To determine what your platform system is, enter the following command:

```
# bootinfo -p
```

- a. If you are cloning from the product tape, create a diskette that contains a `./bosinst.data` file with `SWITCH_TO_PRODUCT_TAPE=yes` in the `control_flow` stanza if this was not set prior to making the `mksysb`.
- a. If `SWITCH_TO_PRODUCT_TAPE` is set to `yes`, the system will prompt you to remove the `mksysb` tape and insert the product tape after the `mksysb` cloning has been restored.

After the mksysb installation completes, the installation program automatically installs additional devices and the kernel (uniprocessor or microprocessor) on your system using the original product media you booted from. Information is saved in BOS installation log files. To view BOS installation log files, enter `cd /var/adm/ras` and view the `devinst.log` file in this directory.

If the source system does not have the correct passwords and network information, you may make modifications on the target system now. Also, some products ship device-specific files such as `graPHIGS`. If your graphics adapter is different on the target system, verify that the device-specific filesets for graphics-related LPPs are installed.

If the system you have cloned is using `OpenGL` or `PEX_PHIGS`, there may be some device filesets from these LPPs that must be installed after a clone. `OpenGL` and `PEX_PHIGS` have graphics adapter-specific filesets, so if you cloned onto a system with a different graphics adapter, you will need to create a bundle as follows:

```
# echo OpenGL.OpenGL_X.dev >
  /usr/sys/inst.data/user_bundles/graphic_dev.bnd
# echo PEX_PHIGS.dev >>
  /usr/sys/inst.data/user_bundles/graphic_dev.bnd
```

Next use `smitty install_bundle` to install this bundle. If any new filesets are installed, you may want to reboot your system before using `OpenGL` and `PEX_PHIGS`. Otherwise, you may experience performance problems with these applications.

The above procedure is required because the `OpenGL` and `PEX_PHIGS` LPPs are on a separate media, and therefore the normal cloning procedure cannot automatically install any missing device filesets.

The Network Installation Management (NIM) mksysb installation also supports cloning. If you are using NIM and you have `OpenGL` and `PEX_PHIGS` in your `lpp_source`, you can create a separate bundle as described, and allocate it when you install the mksysb. The device filesets will be installed automatically.

8.3.5.3 How to Interrupt a Non-prompted Installation

If you are restoring a system using the non-prompted installation, that means you are using the `PROMPT = no` stanza in the `bosinst.data` file and the system will not show you the installation main menu to allow you to change any settings.

You can interrupt the non-prompted installation as described below:

During the boot, after a few seconds, you will see three zeroes (000) and a spinning cursor on your screen, just enter again **three zeroes (000)** immediately and press **Enter**. This causes the system to interrupt the non-prompted installation, start the installation main menu, and allows you to change the settings.

8.4 Configuration Information

Use the `lscfg` command to display the vital product data (VPD) of the system such as part numbers, serial numbers, and engineering change levels from either the Customized VPD object class or platform-specific areas.

If the command is run on a Model 260, it will output data similar to the following (extracted list):

```
lscfg -vp
```

```
INSTALLED RESOURCE LIST WITH VPD
```

```
The following resources are installed on your machine.
```

```
Model Architecture: chrp
Model Implementation: Multiple Processor, PCI bus
```

```
sysplanar0 00-00System Planar
mem0 00-00 Memory
proc0 00-00 Processor
```

```
Device Specific.(YL).....P1-C1
```

```
L2cache0 00-00 L2 Cache
pmc0 00-00 Power Management Controller
proc2 00-02 Processor
```

```
Device Specific.(YL).....P1-C2
```

```
pci1 00-fee00000 PCI Bus
```

```
Device Specific.(YL).....P2
```

```
mtn0 20-60 GXT3000P Graphics Adapter
```

```
GXT3000P 3D Graphics Adapter:
EC Level.....12605059
FRU Number.....24L0030
Manufacture ID.....IBM97N
Part Number.....08L1172
Serial Number.....198016
Version.....RS6K
Displayable Message.....GXT3000P
ROM Level.(alterable).....00
Device Driver Level.....00
Diagnostic Level.....00
```

```

        Product Specific.(Z0).....0001105A
        Device Specific.(YL).....P2-I2

pci0 00-fe00000 PCI Bus

        Device Specific.(YL).....P2

isa0 10-58 ISA Bus

        Device Specific.(YL).....P2

fda0 01-D1 Standard I/O Diskette Adapter

        Device Specific.(YL).....P2/D1

fd0 01-D1-00-00 Diskette Drive
siokma0 01-K1 Keyboard/Mouse Adapter

        Device Specific.(YL).....P2/K1

sioka0 01-K1-00 Keyboard Adapter
kbd0 01-K1-00-00 PS/2 keyboard
sioma0 01-K1-01 Mouse Adapter
mouse0 01-K1-01-00 3 button mouse
siota0 01-Q1 Tablet Adapter

        Device Specific.(YL).....P2/Q1

pau0 01-Q2 Ultimeida Integrated Audio

        Device Specific.(YL).....P2/Q2

ppa0 01-R1 CHRP IEEE1284 (ECP) Parallel Port Adapter

        Device Specific.(YL).....P2/R1

sa0 01-S1 Standard I/O Serial Port

        Device Specific.(YL).....P2/S1

tty0 01-S1-00-00 Asynchronous Terminal
sal 01-S2 Standard I/O Serial Port

        Device Specific.(YL).....P2/S2

scsi0 10-60 Wide/Fast-20 SCSI I/O Controller

        Device Specific.(YL).....P2/Z1

cd0 10-60-00-1,0 SCSI Multimedia CD-ROM Drive (650 MB)

        Manufacturer.....IBM
        Machine Type and Model.....CDRM00203
        ROS Level and ID.....1_00
        Device Specific.(Z0).....058002028F000018
        Part Number.....97H7608
        EC Level.....F15213
        FRU Number.....97H7610

hdisk0 10-60-00-11,0 16 Bit SCSI Disk Drive (9100 MB)

        Manufacturer.....IBM
        Machine Type and Model.....DCHS09U

```

```

FRU Number.....76H2698
ROS Level and ID.....36333633
Serial Number.....68048A97
EC Level.....488638
Part Number.....93G3160
Device Specific.(Z0).....000002029F00003E
Device Specific.(Z1).....RAMR6063
Device Specific.(Z2).....09RI
Device Specific.(Z3).....97350
Device Specific.(Z4).....0001
Device Specific.(Z5).....22
Device Specific.(Z6).....E29526

hdisk1 10-60-00-12,0 16 Bit SCSI Disk Drive (4500 MB)

Manufacturer.....IBM
Machine Type and Model.....DCHS04W
FRU Number.....93G2970
ROS Level and ID.....32383238
Serial Number.....68223287
EC Level.....488630
Part Number.....93G2854
Device Specific.(Z0).....000002029F00003E
Device Specific.(Z1).....RAMUR26C
Device Specific.(Z2).....09RI
Device Specific.(Z3).....97051
Device Specific.(Z4).....0001
Device Specific.(Z5).....22
Device Specific.(Z6).....484537

ent0 10-70 IBM PCI Ethernet Adapter (22100020)

Network Address.....0207011CC86D
Displayable Message.....PCI Ethernet Adapter (22100020)
Device Specific.(YL).....P2-I4

tok0 10-78 IBM PCI Tokenring Adapter (14103e00)

Loadable Microcode Level...001PX11B2
Part Number.....093H9181
EC Level.....00E76815
Serial Number.....007E9237
FRU Number.....093H6594
Manufacturer.....9BM982
Network Address.....0004AC487F7C
Displayable Message.....IBM PCI Token-Ring
Device Specific.(YL).....P2-I5

ent1 10-80 IBM 10/100 Mbps Ethernet PCI Adapter (23100020)

Network Address.....006094E90507
Displayable Message.....PCI Ethernet Adapter (23100020)
Device Specific.(YL).....P2/E1

scsil 10-88 Wide/Ultra-2 SCSI I/O Controller

Device Specific.(YL).....P2/Z2

PLATFORM SPECIFIC

Name: IBM,7043-260
Model: IBM,7043-260
Node: /

```

Device Type: chrp

Processor Card:

Part Number.....08L1200
EC Level.....22605061
Serial Number.....L608182062
FRU Number.....08L0858
Manufacture ID.....1980
Version.....RS6K
Product Specific.(ZC).....PS=000BEBC200, LB=000BEBC200,
SB=0005F5E100, NP=01, L2=04096,
PF=654, CL=4, SV=2, VR=3, ER=00
Product Specific.(ZB).....BC=30602, SG=
Physical Location: P1-C1

Processor Card:

Part Number.....08L1200
EC Level.....22605061
Serial Number.....L608182061
FRU Number.....08L0858
Manufacture ID.....1980
Version.....RS6K
Product Specific.(ZC).....PS=000BEBC200, LB=000BEBC200,
SB=0005F5E100, NP=01, L2=04096,
PF=654, CL=4, SV=2, VR=3, ER=00
Product Specific.(ZB).....BC=30602, SG=
Physical Location: P1-C2

MEMORY CARD:

Part Number.....08L0947
EC Level.....E77526
Serial Number.....15281E0046
FRU Number.....07L7065
Manufacture ID.....091
Version.....RS6K
Product Specific.(ZM).....CL=4, VR=2, DE=2, EM=0100020209
1010120304 0406111412160508060A1318141A070C080E15
1C161E

Physical Location: P1-M2

I/O Planar:

Part Number.....08L0808
EC Level.....E77739A
Serial Number.....L408199043
FRU Number.....07L8509
Manufacture ID.....1980
Version.....RS6K
Product Specific.(ZB).....BC=10601, SG=0Ic+...

Physical Location: P2

System Planar:

Part Number.....08L0824
EC Level.....12605050
Serial Number.....L208170063
FRU Number.....08L0823
Manufacture ID.....1980
Version.....RS6K Product
Specific.(ZM).....MF=4, CL=4, VR=5, ES=8, EM=01F001F001F001F
001F001F001F001F002RC02RC02RC02RC0 2RC02RC02RC
Product Specific.(ZB).....BC=10701, SG=âI"@P
Product Specific.(ZB).....BC=10701, SG=2 Izj+Q
Product Specific.(ZB).....BC=11W01, SG=2 I*E

Physical Location: P1


```

System:
  Machine/Cabinet Serial No...7043265887
  Machine Type and Model.....7043-260
  Manufacture ID.....1980
  Version.....RS6K
Physical Location: L1

Op Panel:
  Part Number.....07L7234
  EC Level.....F15535
  Serial Number.....1328000190
  FRU Number.....24L0000
  Manufacture ID.....1980
  Version.....RS6K
Physical Location: L1

NOT_VALID_VPD:
  System Info Specific.(YL)...V1

Name: memory-controller
  Model: IBM, Racer/Whip, Rev-id ?.?
  Node: memory-controller@ff001000
  Physical Location: P1

Name: nvram
  Node: nvram@ff800000
  Physical Location: P2

Name: IBM,sp
  Node: IBM,sp@ff880000
  Physical Location: P2

SP_CARD_:
  ROM Level.(non-alterable)...px980902

Name: rom
  Model: AM29F040-PLCC
  Node: rom@fff00000
  Physical Location: P2

System Firmware:
  ROM Level.(alterable).....SPX98244
  Version.....RS6K
  System Info Specific.(YL)...P2

Name: interrupt-controller
  Model: IBM, Racer/Whip, Open-PIC, 00
  Node: interrupt-controller@ffc00000
  Device Type: open-pic
  Physical Location: P1

Name: openprom
  Model: IBM,SPX98244
  Node: openprom

Name: pci
  Model: IBM, Python, Rev-id 3.0
  Node: pci@fef00000 Physical Location: P2

Name: pci
  Model: IBM, Python, Rev-id 3.0
  Node: pci@fee00000 Physical Location: P2

```



```

Model: IBM,41H8862
Node: token-ring@f
Device Type: network
Physical Location: P2-I5

Name: ethernet
Model: AMD,am79C971
Node: ethernet@10
Device Type: network
Physical Location: P2/E1

Name: scsi
Node: scsi@11
Device Type: scsi-2
Physical Location: P2/Z2

Name: display
Node: display@c
Physical Location: P2-I2

GXT3000P 3D Graphics Adapter:
EC Level.....12605059
FRU Number.....24L0030
Manufacture ID.....IBM97N
Part Number.....08L1172
Serial Number.....198016
Version.....RS6K
Displayable Message.....GXT3000P
ROM Level.(alterable).....00
Device Driver Level.....00
Diagnostic Level.....00
Product Specific.(Z0).....0001105A

Name: l2-cache
Node: l2-cache
Device Type:cache
Physical Location: P1-C1

Name: l2-cache
Node: l2-cache
Device Type: cache
Physical Location: P1-C2

Name: IBM,tablet-port
Node: IBM,tablet-port@i898
Device Type: serial
Physical Location: P2/Q1

Name: rtc
Node: rtc@i70
Physical Location: P2

Name: parallel
Node: parallel@i378
Physical Location: P2/R1

Name: serial
Node: serial@i3f8
Physical Location: P2/S1

Name: serial
Node: serial@i2f8
Physical Location: P2/S2

```

```

Name: sound
  Model: CS4236B
  Node: sound@i534
  Physical Location: P2/Q2

Name: 8042
  Model: 8042@i60
  Physical Location: P2/K1

Name: fdc
  Model: National,PC87308
  Node: fdc@i3f0
  Physical Location: P2/D1

Name: timer
  Model: WINB,W83C553
  Node: timer@i40
  Device Type: IBM,timer
  Physical Location: P2

Name: interrupt-controller
  Model: WINB,W83C553
  Node: interrupt-controller@i20
  Physical Location: P2

Name: dma-controller
  Model: WINB,W83C553
  Node: dma-controller@i0
  Physical Location: P2

Name: keyboard
  Node: keyboard@0
  Physical Location: P2-K1

Name: mouse
  Node: mouse@1
  Physical Location: P2-O1

Name: disk
  Node: disk@0
  Device Type: block
  Physical Location: P2-D1

```

The output gives more detailed information about the inventory of your system.

8.4.1 Information about the Processor

You can find detailed information about the processor speed if you search for the Processor Card entry in the `lscfg -vp` output.

```

Processor Card:
  Product Specific.(ZC).....PS=00BEBC200, LB=00BEBC200,
                               SB=0005F5E100, NP=01, L2=04096,
                               PF=654, CL=4, SV=2, VR=3, ER=00

```

- PS is the processor speed in hex-number digit.

(0x000BEB C200 = 200000000 Hz = 200 MHz)

- L2=04096 shows you the size of the Level 2 cache (= 4 MB).

8.4.2 Information about the Service Processor

You can determine the actual level of the service processor firmware if you search for the `System Firmware` entry in the `lscfg -vp` output.

```
System Firmware:  
  ROM Level.(alterable).....SPX98244
```

8.4.3 Information about the Environmental Sensors

On CHRP/RPA systems (see 8.3.1, “The bootinfo Command” on page 216) there is a utility available that displays the status of the environmental sensors. The information displayed is the sensor name, physical location code, literal value of the sensor status, and the literal value of the sensor reading.

The sensor status can be any one of the following:

- **Normal:** The sensor reading is within the normal operating range.
- **Critical High:** The sensor reading indicates a serious problem with the device. Run diagnostics on `sysplanar0` to determine what repair action is needed.
- **Critical Low:** The sensor reading indicates a serious problem with the device. Run diagnostics on `sysplanar0` to determine what repair action is needed.
- **Warning High:** The sensor reading indicates a problem with the device. This could become a critical problem if action is not taken. Run diagnostics on `sysplanar0` to determine what repair action is needed.
- **Warning Low:** The sensor reading indicates a problem with the device. This could become a critical problem if action is not taken. Run diagnostics on `sysplanar0` to determine what repair action is needed.
- **Hardware Error:** The sensor could not be read because of a hardware error. Run diagnostics on `sysplanar0` in problem determination mode to determine what repair action is needed.
- **Hardware Busy:** The system has repeatedly returned a busy indication, and a reading is not available. Try the service aid again. If the problem continues, run diagnostics, on `sysplanar0` in problem determination mode to determine what repair action is needed.

To get the information, run the following command

```
/usr/lpp/diagnostics/bin/uesensor -l
```

and you will see a output similar to the this:

```
Sensor = Fan Speed
Status = Normal
Value = 2436 RPM
Location Code = F1

Sensor = Power Supply
Status = Normal
Value = Present and operational
Location Code = V1

Sensor = Power Supply
*Status = Critical low
Value = Present and not operational
Location Code = V2
```

The status of the environmental sensors can also be shown by using the diagnostics service aid. To start diagnostics, enter: `diag`

8.5 Network Boot Support

If you are booting over the network, follow these steps:

1. Begin with your machine turned off or shut down your system properly if it is already running.
2. Power on the machine and enter the System Management Services (SMS) by pressing the **F1** key.
3. The SMS menu displays on your screen. Select the **Utilities** option.
4. From the System Management Services Utilities menu, select the **Remote Initial Program Load (RIPL)** option.
5. From the screen, select the **Set Address** option.
6. Set or change the values displayed so they are correct for your client system.

Specify the IP address of:

- The client machine you are booting in the `client address` field.

- Your SPOT server in the `bootp server address` field.
 - Your client's gateway in the `gateway address` field. If broadcasting will be used to obtain a boot image, leave these IP address fields empty.
 - Specify the `subnet mask` for your client machine if you are prompted for one in the subnet mask field. All machines in your subnet have the same subnet mask.
7. After you specify the addresses, press **Enter** to save the addresses and continue.
 8. The RIPL screen is displayed. Select the **Ping** option and try to ping the server.
 9. Go back to the SMS main menu and select the **Multiboot** function. Enter the Boot Sequence menu and select the network adapter to be used as the client's boot device.
 10. After making your selections, you can boot the machine over the network.

8.6 Using the Error Logging Facility

Problem determination is much easier when you have some information about the nature of the problem that occurs. To enable access to more detailed information about problems that occur in your system, you should install the AIX error logging feature of AIX.

The error logging features are provided with the `bos.sysmgt.serv_aid` fileset. Once this fileset is installed, you will be able to use the `errpt` command to look at the AIX error log. The command:

```
errpt -a |more
```

will display all the information in the error log starting from the most recent entry. Not all entries in the error log are errors. For example, every time the system is powered on, an entry will be logged to say that the error log has been turned on.

The error log is useful for resolving problems before they make an impact on your system. By regularly checking the error log, for example, you could discover that your hard disk was receiving temporary errors and was performing block relocations. If the errors were frequent (for example, every day), you could take actions to replace the disk without impacting the users.

To clear the error log, you can use the `errclear` command. For example, to clear all the entries from the error log, you would run the command:

8.7 Booting Problems

After the firmware has passed control of the system to AIX, problems which occur are, in most cases, the same as for any other RS/6000. In this section, several problems during booting are discussed.

8.7.1 Booting from CD-ROM

While AIX supports installation from external CD-ROM, the CD-ROM drives, 7210-001 and 7210-010, cannot be used as a boot device.

8.7.2 System Hangs during AIX Boot Process

The AIX booting process on the PCI-based RS/6000 is similar to the Micro Channel-based RS/6000 systems. This means that the boot process can be monitored by the LEDs. If the system hangs, you will have to perform system maintenance and try to determine the cause of the problem by using the standard AIX LED recovery procedures. Typically, the procedure to follow is:

1. Boot the system into Maintenance mode by using either a CD-ROM or bootable tape.(if your are booting from a tape make sure that the PROMPT stanza in the bosinst.data is set to YES. Otherwise, a mksysb restore will occur, unless interrupted as described in 8.3.5.3, "How to Interrupt a Non-prompted Installation" on page 226.
2. Select the option to access the root volume group and start a shell before the file systems are mounted.
3. Check the root volume group file systems:

```
fsck -y /dev/hd1
fsck -y /dev/hd2
fsck -y /dev/hd3
fsck -y /dev/hd4
fsck -y /dev/hd9var
```
4. Format the Journaled File System (JFS) log:

```
logform /dev/hd8
```

Answer **YES** when asked if you want to destroy the log.
5. Type `exit` to mount the root file systems.
6. Check the `inittab` file for any entries which may be causing problems.
7. Check for any full file systems (particularly `/`, `/tmp`, `/var` and `/usr`).

8. Check the following files for any modification or problems with permissions:

```
/.profile
/etc/environment
/etc/profile
/etc/inittab
```

In addition, check that the files can be executed without problems, for example:

```
/usr/bin/sh /etc/environment
```

9. Recreate the bootimage:

```
bosboot -a -d /dev/hdiskx
```

Where `hdiskx` is the hard disk containing the boot logical volume. You can determine which disk has the boot logical volume by running:

```
lslv -m hd5
```

The boot disk will be shown in the PV1 column.

If, after trying the above the problem persists, then either contact your support organization or restore from a system backup.

8.7.3 Recovering a System with No Bootable Media

If your system hangs during the AIX boot process, you can use the CD-ROM or tape drive to boot into Maintenance mode using suitable media. If your system hangs and you do not have media, then you can still gain access to a limited maintenance shell.

By using the hardware diagnostics on the hard drive, you can access an AIX shell which will allow you access to the file systems in the system. The file systems will, however, be mounted so you will be unable to run `fsck`. From this shell, you can edit `/etc/inittab` or do whatever is necessary and possible to fix the problem. To access the shell, perform the following:

1. Turn on the RS/6000.
2. Press **F5** during the firmware boot-up sequence (or **5** on an ASCII terminal) to enter Service mode.
3. If diagnostics loads, select **Service Aids**.
4. Select **AIX Shell Prompt**.
5. Enter the root password.

6. Perform maintenance work. Remember to set the `$TERM` variable so that you can use full-screen editors, for example if you are using an `ibm3151` ASCII terminal:

```
TERM=ibm3151
export TERM
```

7. Shut down and restart the system.

If problems persist, then you may require media to boot into Maintenance mode, or you may have to restore the system.

Appendix A. RS/6000 43P 7043 Model 140

The RS/6000 43P 7043 Model 140 is the volume leader of the RS/6000 product line. This appendix collects a general outline of this entry-level system to help introduce it to those not familiar with RS/6000 servers.

A.1 RS/6000 43P 7043 Model 140 Key Attributes

The RS/6000 43P 7043 Model 140 is a desktop system that provides outstanding performance for a graphics workstation or entry-level workgroup server. A wide range of options, including choice of processor speed, storage devices, and communications features enables you to configure a system suited to your business environment.

With its graphics capability, the Model 140 is an excellent workstation solution for demanding 2D or high-function 3D applications. The Model 140 allows business to be securely conducted on the Internet with an affordable, yet powerful, entry e-business server, together with the appropriate security software. It offers the opportunity to start simply, and can easily grow with the business through processor upgrades or through the creation of clusters of e-business servers.

Outstanding expansion capabilities provide confidence that an investment in a Model 140 system can keep pace with the ever-growing needs of businesses. Memory can be expanded from the base 64 MB to 768 MB. The Model 140 can handle tough 3D solid modeling, wireframe, and drafting applications. The 2.1 GB of a standard SCSI-2 Fast/Wide internal disk capacity can be increased up to 27.3 GB to satisfy the storage requirements of many entry-level server applications.

To help increase your system performance and value, the PowerPC 604e 332 MHz microprocessor is available on new orders for RS/6000 7043 Model 140. Upgrading from the base 200 MHz microprocessor to the 233 MHz or 332 MHz option is accomplished by removing the existing processor and plugging in a new one.

The Model 140 supports the MVP Power Multi-monitor adapter that allows attachment of two monitors to view multiple data sets simultaneously. This is especially beneficial for applications such as security trading, financial analysis, and software engineering. The unique design of this adapter lets a single mouse move seamlessly between the two monitors for easy operation.



Figure 41. RS/6000 43P 7043 Model 140

The Model 140 has many of the reliability, availability and serviceability (RAS) features essential to today's computing environments. The Error Checking and Correcting (ECC) memory helps detect and correct many single-bit errors. An LED operator status panel on the front of the system helps keep you informed of the system status. Concurrent diagnostics enable you to troubleshoot minor problems that may occur without requiring system shutdown.

A.2 RS/6000 43P 7043 Model 140 Feature Summary

The Model 140 optional processor features provide:

- High-end performance
- Modular processor design to help protect your investment
- Easy user installation upgrade from the current microprocessor to the PowerPC 604e 332 MHz microprocessor

Some of the standard configurations and features are described Table 41.

Table 41. RS/6000 43P 7043 Model 140 Standard Configuration

Standard Configuration	
Microprocessor	200 MHz PowerPC 604e
Level 1 (L1) cache	32 KB data/32 KB instruction
Level 2 (L2) cache	1 MB
RAM (Memory)	64 MB ECC DIMM
Memory bus width	64-bit
Internal disk drive	2.1 GB SCSI-2
Disk/Media Bay	Five (one disk and one media available)
Expansion slots	Five (three 32-bit and two shared PCI/ISA)
PCI bus width	32-bit
Memory slots	Six
Ports	Tablet, keyboard, mouse, Ethernet, SCSI-2 F/W, serial ports (2), parallel port (1), and stereo audio

A.3 Standard Features

Standard features are as follows.

- Vertical stand
- LED operator panel
- 20X SCSI-2 CD-ROM drive
- 1.44 MB 3.5-inch diskette drive
- Integrated SCSI-2 Fast/Wide controller
- Integrated Ethernet adapter
- AIX Version 4 1-2 user server license

Table 42 lists the upgrades available for the standard features.

Table 42. RS/6000 43P 7043 Model 140 System Expansion

System Expansion	
Processor upgrades	233/332 MHz PowerPC 604e

System Expansion	
RAM (Memory)	Up to 768 MB
Internal disk storage	Up to 27.3 MB
External disk storage	Up to 291 GB SCSI-2; Up to 873.6 GB SSA

A.4 Supported Optional Features

The RS/6000 43P 7043 Model 140 supports a wide range of optional features and devices, such as:

- Communications adapters for 25/155 Mbps ATM, ISDN, Token-Ring, 10/100 Mbps Ethernet, FDDI, X.25, SDLC, and asynchronous environments
- Graphics accelerators for 8-bit entry 2D, 24-bit true-color 2D, production quality 3D, and high-performance 3D graphics
- Color displays of various sizes and resolutions
- ASCII terminals and the IBM Network Stations
- Input devices, including dials, LPF keys and spaceball
- IBM and non-IBM printers, plotters and modems
- Multimedia devices, including a personal microphone and amplifies speakers

A.4.1 Storage Devices

By using common industry storage interfaces, the Model 140 enables attachment of many available storage devices. In addition to the popular and widely-used SCSI-2 interfaces, the Serial Storage Architecture (SSA) and Ultra SCSI storage interfaces are supported. SSA disks deliver outstanding performance for transaction-intensive applications and offer redundancy features to help keep your data safe. Ultra SCSI disks offer significant improvements in data transfer rates over traditional SCSI to provide the performance required by many popular e-business data-intensive applications.

A.4.2 AIX Operating System

The computing power of the Model 140 is strengthened by its operating system. AIX Version 4 adds real value over other UNIX operating systems in terms of reliability, availability, and system management. It also offers exceptional scalability with the multi-tasking, multi-user functionality for

technical and commercial distributed environments. With a vast portfolio of available commercial and technical applications, AIX can provide solutions just right for your business, regardless of its size.

The AIX operating system for 1-2 users license is now available with all models of IBM RS/6000 at no additional cost. A Bonus Pack CD is also provided as a part of operating system.

The Model 140 is supported with AIX 4.1.5, 4.2.1 and 4.3.

A.4.3 Year 2000

The Model 140 is Year 2000 ready. For more information about Year 2000 refer to 1.8, "Year 2000" on page 18.

Appendix B. Special Notices

This publication is intended to help technical and non-technical professionals market, support, configure, and manage the RS/6000 43P 7043 Models 150 and 260. The information in this publication is not intended as the specification of any programming interfaces that are provided by these systems. See the PUBLICATIONS section of the IBM Programming Announcement for RS/6000 43P 7043 Model 150, and RS/6000 43P 7043 Model 260 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

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

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM ("vendor") products in this manual has been supplied by the vendor and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate

them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX®	PowerPC Architecture
AS/400®	PowerPC 604
BookManager®	POWER2 Architecture
eNetwork	POWER3 Architecture
IBM ®	RETAIN®
IBMLink	RISC System/6000®
Magstar®	RS/6000®
Micro Channel®	Service Director®
Network Station	SP
OS/2®	System/390
POWERparallel®	TURBOWAYS®
PowerPC®	VisualAge®

The following companies are subsidiaries of IBM Corporation whose trademarks have been included in this publication:

Lotus Development Corporation
Cambridge, Massachusetts

Tivoli Systems, Inc.
9442 Capitol of Texas Highway, North
Plaza 1, Suite 500
Austin, Texas 78759

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

Java and HotJava are trademarks of Sun Microsystems, Incorporated.

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

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

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

Appendix C. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

C.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 253.

- *Managing AIX V4 on PCI-Based RISC System/6000 Workstations (40P/43P)*, SG24-2581
- *RS/6000 Models E30, F40, F50, and H50 Handbook*, SG24-5143
- *A Technical Introduction to PCI-Based RS/6000 Servers*, SG24-4690
- *A Practical Guide to Serial Storage Architecture for AIX*, SG24-4599

C.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

C.3 Other Publications

These publications are also relevant as further information sources:

- *Exploring IBM RS/6000 Computers*, GA23-2674-07

- *PowerPC Architecture: A Specification for a New Family of RISC Processors*, Second Edition, Morgan Kaufmann Publishers, Inc., San Francisco, CA, ISBN 1-55860-316-6
- *IEEE 1275, IEEE Standard for Boot (Initialization Configuration) Firmware, Core Requirements and Practices*, IEEE part number DS02683, ISBN 1-55937-426-8
- *AIX Version 4.3 Installation Guide*, SC23-4112
- *7043 Model 260 Hardware Setup*, SA38-0555
- *7043 Model 260 User's Guide*, SA38-0553
- *7043 Model 260 Service Guide*, SA38-0554
- *Customer Support Information*, SA23-2690
- *System Unit Safety Information*, SA23-2652
- *PCI Adapter Placement Reference Guide*, SA38-0538
- *CIO - Option Pubs Library English*, SK2T-7452
- *Diagnostics Information for Multiple Bus Systems*, SA38-0509
- *IBM RS/6000 Adapters, Devices, and Cable Information for Multiple Bus Systems*, SA38-0516
- *7043 43P Series Service Guide*, SA38-0512
- *7043 43P Series Users Guide*, SA38-0511
- *AIX Version 4.3 Network Installation Management Guide and Reference*, SC23-4113

C.4 Web Resources

These links will provide you with a stimulating look at UNIX server technology.

- <http://www.rs6000.ibm.com>
- <http://www.tpc.org>
- <http://www.specbench.org>
- <http://www.netlib.no/netlib/benchmark/performance.ps>
- <http://www.chips.ibm.com/products/ppc>

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at <http://www.redbooks.ibm.com/>.

How IBM Employees Can Get ITSO Redbooks

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Redbooks Web Site on the World Wide Web**

<http://w3.itso.ibm.com/>

- **PUBORDER** – to order hardcopies in the United States

- **Tools Disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLCAT REDPRINT
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
```

To get BookManager BOOKs of redbooks, type the following command:

```
TOOLCAT REDBOOKS
```

To get lists of redbooks, type the following command:

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
```

To register for information on workshops, residencies, and redbooks, type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
```

- **REDBOOKS Category on INEWS**

- **Online** – send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **Online Orders** – send orders to:

	IBMMAIL	Internet
In United States	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** – send orders to:

IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA	IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada	IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark
--	--	--

- **Fax** – send orders to:

United States (toll free)	1-800-445-9269
Canada	1-800-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1) 408 256 5422 (Outside USA)** – ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com
IBM Direct Publications Catalog	http://www.elink.ibm.link.ibm.com/pbl/pbl

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (<http://www.redbooks.ibm.com/redpieces.html>). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____
_____	_____	_____

First name _____ Last name _____

Company _____

Address _____

City _____ Postal code _____ Country _____

Telephone number _____ Telefax number _____ VAT number _____

Invoice to customer number _____

Credit card number _____

Credit card expiration date _____ Card issued to _____ Signature _____

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

List of Abbreviations

ABI	application binary interface	EC	engineering change
AC	alternating current	ECC	error checking and correction
AIX	Advanced Interactive Executive	EEPROM	electrically erasable programmable read only memory
ANSI	American National Standards Institute	EISA	extended industry standard architecture
APAR	authorized program analysis report	ELA	error log analysis
AGP	Accelerated Graphics Port	EPOW	environmental and power warning
ASCII	American National Standards Code for Information Interchange	FCP	Fibre Channel Protocol
BIST	built-in self-test	FDDI	fiber distributed data interface
BOS	base operating system	FRU	field replaceable unit
CAE	computer-aided engineering	GB	gigabytes
CEC	central electronics complex	Hz	hertz
CHRP	Common Hardware Reference Platform	HTTP	Hypertext Transfer Protocol
CMOS	complimentary metal oxide semiconductor	ID	identification
CPU	central processing unit	IDE	integrated device electronics
CRC	cyclic redundancy check	IEEE	Institute of Electrical and Electronics Engineers
DC	direct current	°C	inter integrated-circuit communications
DDC	display data channel	I/O	input/output
DIMM	dual in-line memory module	IP	internetwork protocol (OSI)
DMA	direct memory access	IPL	initial program load
DOS	disk operating system	IrDA	infrared data association (which sets standards for infrared support including protocols for data interchange)
DRAM	dynamic random access memory		

ISA	Industry Standard Architecture	RAS	reliability, availability, and serviceability
ISDN	integrated-services digital network	RAID	redundant array of independent disks
ISV	independent software vendor	RIPL	remote initial program load
ITSO	International Technical Support Organization	ROLTP	Relative On-line Transaction Processing
LAN	local area network	RPA	RS/6000 Platform Architecture
LED	light emitting diode	RTC	real-time clock
MAP	Maintenance Analysis Procedure	SCSI	small computer system interface
MCA	micro channel architecture	SDRAM	Synchronous Dynamic Random Access Memory
MCAD	mechanical computer-aided design	SDLC	synchronous data link control
MP	multiprocessor	SMS	system management services
NUMA	non-uniform memory access	SMP	Symmetric multiprocessing
NTSC	National Technical Support Center	SP	service processor
NVRAM	non-volatile random access memory	SPEC	Standard Performance Evaluation Corp.
OCS	on-line customer support	SPOT	shared product object tree
ODM	Object Data Model	SSC	system support controller
PAP	privileged access password	SSA	Serial Storage Architecture
POP	power-on password	STP	shielded twisted pair
POST	power-on self-test	TCQ	tagged command queuing
POWER	performance optimization with enhanced RISC (architecture)	UP	uniprocessor
PCI	peripheral component interconnect	USB	Universal Serial Bus
P2SC	Power2 Super Chip	VESA	Video Electronics Standards Association
RAM	random access memory	VPD	vital product data

Index

Numerics

- 1000Base-SX 128
- 1000BaseT Ethernet 128
- 100BaseTx Ethernet 128
- 10BaseT Ethernet 128
- 2639 ultimedia video capture adapter 134
- 2708 Eicon ISDN DIVA PRO adapter 128
- 2741 SysKonnnect FDDI-LP adapter (SAS) 129
- 2742 SysKonnnect FDDI-LP adapter (DAS) 130
- 2743 SysKonnnect FDDI-UP adapter (SAS) 130
- 2825 GXT3000P graphics accelerator 138
- 2838 GXT120P graphics adapter 135
- 2845 GXT550P graphics adapter 137
- 2851 GXT250P graphics adapter 136
- 2852 GXT255P graphics adapter 136
- 2900 4.5 GB Ultra SCSI 16-bit disk drive 53
- 2908 9.1 GB Ultra SCSI 16-bit disk drive 53
- 2920 PCI token-ring adapter 127
- 2943 8-port asynchronous adapter 133
- 2944 128-port asynchronous controller 133
- 2947 ARTIC960HX 4-port selectable adapter 134
- 2948 ARTIC960HX 4-port T1/E1 adapter 134
- 2962 2-port multiprotocol adapter 133
- 2963 TURBOWAYS 155 UTP ATM adapter 131
- 2968 PCI 10/100 Mbps Ethernet adapter 128
- 2969 Gigabit Ethernet adapter 89, 128
- 2988 TURBOWAYS 155 MMF ATM adapter 131
- 2998 TURBOWAYS 25 ATM adapter 132
- 3620 IBM P72 color monitor, pearl white 61
- 3621 IBM P92 color monitor, pearl white 61
- 3622 IBM P202 color monitor, pearl white 61
- 3623 IBM P72 color monitor, stealth black 61
- 3625 IBM P92 color monitor, stealth black 61
- 3626 IBM P202 color monitor, stealth black 61
- 3752 service package 67, 94
- 4098 memory expansion feature 83
- 4107 64MB (2X 32 MB) SDRAM DIMMs (Model 260) 83
- 4110 256 MB (2X 128 MB) SDRAM DIMMs (Model 260) 83
- 4149 64 MB SDRAM DIMM (Model 150) 59
- 4150 128 MB SDRAM DIMM (Model 150) 59
- 4169 256 MB SDRAM DIMM (Model 150) 59
- 5005 software pre-install 61, 88
- 6041 3-button mouse, pearl white 60
- 604e design 37
- 60X bus 34
- 6156 20/40 GB 16-bit 8mm internal tape drive (black bezel) 84
- 6159 12/24 GB 4mm internal tape drive 85
- 6206 PCI single-ended Ultra SCSI adapter 62, 88, 106, 109, 110, 124
- 6207 PCI differential Ultra SCSI adapter 62, 88, 106, 124
- 6215 PCI SSA Multi-Initiator/RAID EL adapter 62, 88, 125
- 6222 SSA fast write cache option card 125
- 6308 personal microphone 61, 87
- 6508 media bay disk drive mounting kit (Model 260) 84
- 6509 media bay disk drive mounting kit (Model 150) 53
- 6XX bus 34
- 6XX-MX bus 34
- 8130 16-port remote EIA-232 async. node 133
- 8134 16-port (non-US) remote EIA-232 async. node 133
- 8136 16-port rack mountable remote EIA-232 async. node 133
- 8137 16-port enhanced remote EIA-232 node 133
- 8138 16-port enhanced remote RS-422 node 133
- 8679 multimedia kit 95
- 8741 3-button mouse, stealth black 60

A

- accelerated graphics port 14
- adapter placement guidelines
 - Model 150 66
 - Model 260 92
- adapters
 - 128-port async 133
 - 16-port enhanced remote EIA-232 node 133
 - 8-port async EIA-232/RS-422 133
 - ATM 131
 - communication 127
 - Ethernet 128
 - FDDI 129
 - ISDN 128
 - multiple port 133
 - PCI SCSI-2 single-ended Ultra SCSI 124
 - PCI SSA Multi-Initiator/RAID EL 125
 - token-ring 127
 - ultimedia video capture 134

- WAN 133
- AGP 14
- ANSI standard
 - X3.131-1986 97
 - X3T9.2/375R Revision 10K 98
- Apple Inc., regarding CHRP 21
- architecture
 - ISA 33
 - PCI 32
 - POWER3 35
 - PowerPC 604e 37
 - SCSI-III 99
- asynchronous communication, SCSI 103
- ATM adapters 131

B

- bays
 - Model 150 53, 54
 - Model 260 74
- boot
 - from CD-ROM 238
 - problems 238
 - recover from a system hang 238
 - recovering with no bootable media 239
- boot sequence
 - in general 166
 - Model 150 167
 - Model 260 168
- bootinfo, command 216, 219
- bootlist, command 215
- bosboot, command 216
- bosinst.data
 - BOSINST_DEBUG stanza 224
 - RECOVER_DEVICES stanza 224
 - SWITCH_TO_PRODUCT_TAPE stanza 224
 - BOSINST_DEBUG stanza 224
- bus
 - 60X 34
 - 6XX 34
 - 6XX-MX 34
 - ISA 35
- bus architecture 32

C

- Cables
 - display 142
- cables, SCSI (various types) 110
- cabling the pci differential Ultra SCSI adapter 113

- cache
 - coherency 44
 - hit 40
 - L1 39
 - L2 39
 - miss 40
 - set associativity 40
- cache line 41
 - boundary 41
 - prefetch 44
- CALL-IN 202
- CALL-OUT 202
- cfgmgr, command 150, 159, 223
- chdisp, command 144
- cloner, script 223
- cloning
 - AIX 4.1 218
 - AIX 4.2 and 4.3 223
- command
 - bootinfo 216, 219
 - bootlist 215
 - bosboot 216
 - cfgmgr 150, 159, 223
 - chdisp 144
 - diag 158, 161
 - errclear 238
 - errpt 237
 - fsck 238
 - logform 238
 - lsattr 150
 - lscfg 150, 227
 - lsdisp 144
 - lslv 239
 - mksysb 217
 - savevg 217
 - shutdown 161
 - uesensor 236
 - xinit 145
- common command set, SCSI 98
- communication adapters 127, 128
 - ATM 131
 - FDDI 129
 - multiple port 133
 - token-ring 127
- communications adapters
 - ISDN 128
 - Model 260 89
 - WAN 133
- Compaq, regarding PCI involvement 32

comparison between SSA and SCSI architecture 119
configuration information 227
congruence class 42
console mirroring 192
cross invalidate 45
customer account setup 208
customization
 cloner 223

D

debug mode 216
device-to-device cables 114
diag, command 158, 161
diagnostics
 considerations 155
 on-line 158
 running from a tty 156
 stand-alone 157
 tasks and service aids 162
differential Ultra SCSI adapter, PCI 124
differential, SCSI 104
Digital (DEC), regarding PCI involvement 32
disconnect/reconnect, SCSI 98
display cable
 feature number matrix 142

E

electronic key-switch 215
errclear, command 238
errpt, command 237
Ethernet adapter
 Gigabit 89, 128
 integrated 55, 77
 Model 260 55, 77
Ethernet adapters 128
external SCSI and SSA storage devices 126

F

F1-key, boot function 173
F5-key, boot function 173
F6-key, boot function 173
F8-key, boot function 173
false sharing 46
fast SCSI 102
fast system boot 197
fast/wide, regarding SCSI-2 103

fast-20 wide
 SCSI 103
fast-20, SCSI 103
FDDI adapters 129
feature code
 2639 ultimedia video capture adapter 134
 2708 Eicon ISDN DIVA PRO 128
 2741 SysKonnnect FDDI-LP adapter 129
 2742 SysKonnnect FDDI-LP adapter (DAS) 130
 2743 SysKonnnect FDDI-UP adapter 130
 2825 GXT3000P graphics accelerator 138
 2838 GXT120P graphics adapter 135
 2845 GXT550P graphics adapter 137
 2851 GXT250P graphics adapter 136
 2852 GXT255P graphics adapter 136
 2900 4.5 GB Ultra SCSI 16-bit disk drive 53
 2908 9.1 GB Ultra SCSI 16-bit disk drive 53
 2920 PCI token-ring adapter 127
 2943 8-port asynchronous adapter 133
 2944 128-port asynchronous controller 133
 2947 ARTIC960HX 4-port selectable adapter 134
 2948 ARTIC960HX 4-port T1/E1 adapter 134
 2962 2-port multiprotocol adapter 133
 2963 TURBOWAYS 155 UTP ATM adapter 131
 2968 PCI 10/100 Mbps Ethernet adapter 128
 2969 Gigabit Ethernet adapter 89, 128
 2988 TURBOWAYS 155 MMF ATM adapter 131
 2998 TURBOWAYS 25 ATM adapter 132
 3620 IBM P72 color monitor, pearl white 61
 3621 IBM P92 color monitor, pearl white 61
 3622 IBM P202 color monitor, pearl white 61
 3623 IBM P72 color monitor, stealth black 61
 3625 IBM P92 color monitor, stealth black 61
 3626 IBM P202 color monitor, stealth black 61
 3752 service package 67, 94
 4098 memory expansion feature 83
 4107 64MB (2X32 MB) SDRAM DIMMs (Model 260) 83
 4110 256 MB (2X 128 MB) SDRAM DIMMs (Model 260) 83
 4149 64 MB SDRAM DIMM (Model 150) 59
 4150 128 MB SDRAM DIMM (Model 150) 59
 4169 256 MB SDRAM DIMM (Model 150) 59
 5005 software pre-install 61, 88
 6041 3-button mouse, pearl white 60
 6156 20/40 GB 16-bit 8mm internal tape drive

- (black bezel) 84
- 6159 12/24 GB 4mm internal tape drive 85
- 6206 PCI single-ended Ultra SCSI adapter 62, 88, 106, 109, 110, 124
- 6207 PCI differential Ultra SCSI adapter 62, 88, 106, 124
- 6215 PCI SSA Multi-Initiator/RAID EL adapter 62, 88, 125
- 6222 SSA fast write cache option card 125
- 6308 personal microphone 61, 87
- 6508 media bay disk drive mounting kit (Model 260) 84
- 6509 media bay disk drive mounting kit (Model 150) 53
- 8130 16-port remote EIA-232 async. node 133
- 8134 16-port (non-US) remote EIA-232 async. node 133
- 8136 16-port rack mountable remote EIA-232 async. node 133
- 8137 16-port enhanced remote EIA-232 node 133
- 8138 16-port enhanced remote RS-422 node 133
- 8679 multimedia kit 95
- 8741 3-button mouse, stealth black 60
- firmware
 - checkpoints 184
 - introduction 166
 - recovery 183
 - update (service processor) 210
 - update (SMS) 180
- fsck, command 238
- G**
 - general access password 188, 192
 - general SCSI consideration 106
 - graphic peripheral devices
 - Model 150 60
 - Model 260 86
 - graphics adapter
 - Model 150 140
 - Model 260 140
 - graphics adapters
 - Model 150 60
 - Model 260 85
 - graphics support matrix 142
 - GXT120P, graphics adapter 135, 140
 - GXT250P, graphics adapter 136

- GXT250P, graphics adapter 140
- GXT255P, graphics adapter 136, 140
- GXT3000P, graphics adapter 138, 140
- GXT550P, graphics adapter 137, 140

H

- H.323 standard 16, 17
- hard disk spin up delay 179
- hardware
 - design 26
 - overview 21
- heartbeat 193
- hit ratio 40

I

- I/O ports
 - Model 150 54
 - Model 260 76
- IBM, regarding PCI involvement 32
- IDE, industry-standard interface 24
- IEEE 1149.1 39
- IEEE 1275 24
- IEEE 1394 15
- industry trends and directions 13
- information about the processor 234
- information about the system 227
- integrated Ethernet adapter
 - Model 260 55, 77
- Intel, regarding PCI involvement 32
- internal CD-ROM drive
 - Model 150 59
 - Model 260 75, 84
- internal disk drives
 - Model 150 59
 - Model 260 84
- internal SCSI devices 120
- internal SSA devices 120, 123
- internal tape drive
 - Model 150 59
- internal tape drives
 - Model 260 84
- IrDA, protocol 24
- ISA
 - bus architecture 33
- ISDN adapters 128

K

keyboard 59, 85

L

L1 cache 39
 Least Recently Used 42
 Round-Robin 42
 set associativity 40
L2 cache 39
LINPACK DP description 12
LINPACK SP description 12
LINPACK TPP description 12
LocalTalk, industry-standard interface 24
logform, command 238
lsattr, command 150
lscfg, command 150, 227
lsdisp, command 144
lslv, command 239

M

Memory
 hierarchy 39
memory configuration 201
memory deconfiguration 201
memory slot configuration for Model 260 84
memory subsystem
 Model 150 52
 Model 260 72
MESI protocol 45
microphone 61, 87
Microprocessor Common Hardware Reference Platform 21
microprocessor future 13
mkysyb
 create 217
 restore 217
mkysyb, command 217
Model 140
 key attributes 241
 multi-monitor adapter 241
 optional features 244
 overview 3
 standard features 243
Model 150
 adapter placement guidelines 66
 boot sequence 167
 communications adapter 62
 firmware checkpoints 184

 graphic peripheral devices 60
 graphics adapters 60
 internal CD-ROM drive 59
 internal disk drives 59
 internal tape drive 59
 key attributes 4
 memory subsystems 52
 optional features 58
 overview 3
 PCI storage adapters 62
 performance 13
 processor subsystem 50
 reliability, availability, and serviceability 56
 standard features 50
 supported devices 63
Model 260
 adapter placement guidelines 92
 boot sequence 168
 communications adapters 89
 fast system boot 197
 firmware checkpoints 184
 graphic peripheral devices 86
 graphics adapters 85
 I/O board components 71
 integrated Ethernet adapter 55, 77
 internal CD-ROM drive 75, 84
 internal disk drives 84
 internal tape drives 84
 key attributes 7
 memory configuration/deconfiguration 201
 memory subsystem 72
 optional features 82
 overview 6
 PCI storage adapters 88
 performance 13
 processor configuration/deconfiguration 200
 processor subsystem 69
 rear view 77
 reliability, availability, and serviceability 79
 standard features 69
 supported devices 89
 system architecture 69
monitors 61, 86
Motorola Inc., regarding CHRP 21
mouse 60, 85
multiple port adapters 133
multiprocessing
 cache coherency problem 44
 cache hit versus cache miss 40

- cache line prefetch 44
 - concepts 39
 - cross invalidate 45
 - false sharing 46
 - hardware characteristics 39
 - memory hierarchy 39
 - MESI protocol 45
 - set associativity 41
 - sharing resources 39
 - snooping 45
- N**
- NCR, regarding PCI involvement 32
 - network boot 236
 - NUMA 34
- O**
- on-line diagnostics 158
 - concurrent mode 160
 - service mode 159
 - single-user mode 161
 - open firmware 173, 182
 - optional features
 - Model 150 58
 - Model 260 82
 - overview
 - adapter 127
 - hardware 21
- P**
- PC card, protocol 24
 - PCI
 - bus architecture 32
 - features and benefits 33
 - Special Interest Group 32
 - specification revision 2.0 32
 - specification revision 2.1 32
 - PCI differential Ultra SCSI adapter 124
 - PCI single-ended Ultra SCSI adapter 124
 - PCI slots
 - Model 150 54
 - Model 260 76
 - PCI storage adapters
 - Model 150 62
 - Model 260 88
 - performance summary 13
 - personal microphone 61, 87
 - POST 170
 - POWER3 design 35
 - POWER3 microprocessor 35
 - power-on password 177
 - power-on self test 170
 - PowerPC
 - introduction 21
 - PowerPC 604e microprocessor 37
 - PowerPC Platform Specification (PRep) 21
 - printers
 - Model 150 65
 - Model 260 92
 - privileged access password
 - service processor 187, 192
 - SMS 177
 - processor
 - mini-support 35
 - POWER3 35
 - PowerPC 604e 37
 - processor configuration 200
 - processor deconfiguration 200
 - processor speed 234
 - processor subsystem
 - Model 150 50
 - Model 260 69
 - POWER3 35
 - PowerPC 604e 37
 - progress indicators 198
 - publications
 - Model 150 55
 - Model 260 78
- R**
- RAS 56, 79
 - reboot policy 196
 - RECOVER_DEVICES stanza 224
 - reliability, availability, and serviceability
 - Model 150 56
 - Model 260 79
 - repeaters, SCSI 104
 - restart policy 196
 - ring indicate power-on 195
 - RIPL 179
 - RPA
 - achievements 24
 - description of 21
 - goals of 21
 - history of 21

- introduction 21
- minimum systems requirements 25
- platform topology 22
- RS/6000 43P 7043 Model 150
 - hardware design 26
- RS/6000 43P 7043 Model 260
 - hardware design 29
- RS/6000 Platform Architecture (RPA) 21
- rs6k platform 223
- rs6ksmp platform 223

S

- savevg, command 217
- SCSI
 - asynchronous 103
 - common command set 98
 - differential 104
 - disconnect/reconnect 98
 - repeaters 104
 - specifications summary 105
 - summary of specifications 105
 - synchronous 103
 - table of modes/rates 105
 - tagged command queuing 98
 - Ultra 102
- SCSI, location codes 147
- SCSI-III 99
- service center telephone number 207
- service package 67, 94
- service processor
 - checkpoints 213
 - customize and access 186
 - error log 198
 - firmware update 210
 - general access password 188, 192
 - general user menu 209
 - introduction 185
 - list configuration 198
 - local access 187
 - main menu 188, 191
 - modem configuration 203
 - operational phases 188
 - privileged access password 187, 192
 - remote access 187
 - reset 194
 - ring indicate power-on 195
 - security 187
 - self test 169
 - standby mode 169
 - system environmental conditions 199
 - unattended start mode 195
- set associative, 128-way 41
- set associative, 4-way 41
- set associativity
 - L1 cache 40
- shutdown, command 161
- single-ended Ultra SCSI adapter, PCI 124
- single-user mode 161
- SMP hardware characteristics 39
- SMS 172
 - ASCII main menu 180
 - F1-key 173
 - F5-key 173
 - F6-key 173
 - F8-key 173
 - firmware checkpoints 184
 - firmware recovery 183
 - firmware update 180
 - graphical main menu 174
 - power-on password 177
 - privileged access password 177
 - start-up 172, 173
- snooping 45
- SPECfp_base_rate95 description 12
- SPECfp_base95 description 12
- SPECfp_rate95 description 12
- SPECfp95 description 12
- SPECint_base_rate95 description 11
- SPECint_base95 description 11
- SPECint_rate95 description 11
- SPECint95 description 11
- SPECweb96 description 12
- stand-alone diagnostics 157
- standard features
 - Model 150 50
 - Model 260 69
- standards
 - ANSI X3.131-1986 97
 - X3T9.2/375R Revision 10K 98
- storage boot devices 97
- stride 42
- supported devices
 - Model 150 63
 - Model 260 89
- surveillance
 - setup menu 192
- SWITCH_TO_PRODUCT_TAPE stanza 224

- synchronous, SCSI 103
- system backup
 - create 217
 - restore 217
- system environmental conditions 199
- system management services 172

T

- tagged command queuing, SCSI 98
- talk mode 192
- token-ring adapter 127
- topology, general platform (RPA) 23

U

- uesensor, command 236
- ultimedia video capture adapter 134
- Ultra SCSI 102
- Ultra SCSI adapter, PCI differential 124
- Ultra SCSI adapter, PCI single-ended 124
- Ultra SCSI Wide vs Ultra SCSI 103
- Ultra SCSI, defined 103
- unattended start mode 195
- universal serial bus 14
- USB 14

V

- vital product data 197
- VPD 197

W

- WAN adapters 133
- Web-Based System Management 151
- wide SCSI 103
- workgroup conferencing 15
- workstation family 1
- workstation offering 1

X

- X-bus 35
- xinit, command 145

Y

- Year 2000 18

ITSO Redbook Evaluation

RS/6000 43P 7043 Models 150 and 260 Handbook
SG24-5144-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at <http://www.redbooks.ibm.com>
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?

Customer **Business Partner** **Solution Developer** **IBM employee**
 None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes___ No___

If no, please explain:

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

