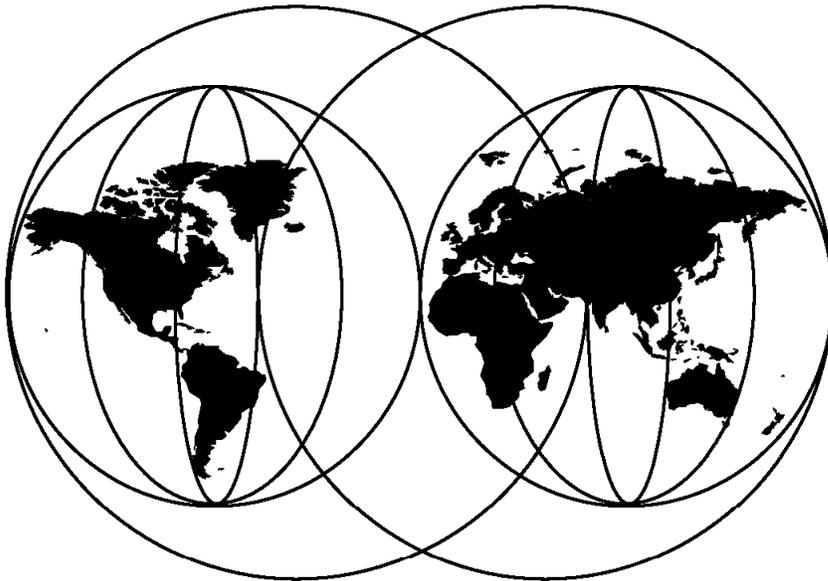




RS/6000 Models E30, F40, F50, and H50 Handbook

*Verena Gschell
Scott Vetter*



International Technical Support Organization

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



SG24-5143-00

International Technical Support Organization

RS/6000 Models E30, F40, F50, and H50 Handbook

April 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, "Special Notices" on page 249.

First Edition (April 1998)

This edition applies to the RS/6000 7024-E30, 7025 Models F40 and F50, and 7026-H50 servers for use with AIX Version 4.3 and subsequent releases.

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

Figures	ix
Tables	xi
Preface	xiii
The Team That Wrote This Redbook	xiii
Comments Welcome	xiv
Chapter 1. RS/6000 Introduction	1
1.1 RS/6000 Server Offerings	1
1.2 RS/6000 Model E30 Overview	2
1.2.1 RS/6000 Model E30 Key Attributes	2
1.2.2 RS/6000 Model E30 Feature Summary	3
1.3 RS/6000 Model F40 Overview	4
1.3.1 RS/6000 Model F40 Key Attributes	5
1.3.2 RS/6000 Model F40 Feature Summary	5
1.4 RS/6000 Model F50 Overview	7
1.4.1 RS/6000 Model F50 Key Attributes	7
1.4.2 RS/6000 Model F50 Feature Summary	8
1.5 RS/6000 Model H50 Overview	10
1.5.1 RS/6000 Model H50 Key Attributes	10
1.5.2 RS/6000 Model H50 Feature Summary	11
1.6 Performance Results	12
1.6.1 SPEC95 Software Benchmark	15
1.6.2 Performance Summary	16
1.7 Industry Trends and Directions	18
1.7.1 PowerPC Futures	18
1.7.2 Universal Serial Bus	18
1.7.3 IEEE 1394	19
1.7.4 Storage Devices	20
Chapter 2. Hardware Overview	23
2.1 PowerPC Microprocessor Common Hardware Reference Platform	23
2.1.1 The PowerPC Platform Introduction	23
2.1.2 PowerPC Platform Achievements	24
2.2 The Hardware Design	25
2.2.1 The PCI Bus Architecture	26
2.2.2 The ISA Bus Architecture	28
2.2.3 The 60X and 6XX Bus	29
2.2.4 The 6XX-MX I/O Bus	30
2.2.5 The X-Bus	30

2.2.6	The Processor Subsystem	31
2.2.7	SMP Hardware Characteristics	32
Chapter 3.	Detailed Product Descriptions	41
3.1	RS/6000 Model E30 Product Description	41
3.1.1	Standard Features	42
3.1.2	Supported Optional Features	44
3.1.3	External Subsystems	50
3.1.4	RS/6000 Model E30 Configuration Notes	51
3.2	RS/6000 Model F40 Product Description	52
3.2.1	Model F40 Highlights	52
3.2.2	Standard Features	54
3.2.3	Supported Optional Features	58
3.2.4	RS/6000 Model F40 Configuration Notes	66
3.3	RS/6000 Model F50 Product Description	66
3.3.1	Model F50 Highlights	67
3.3.2	Model F50 System Architecture	68
3.3.3	Standard Features	71
3.3.4	Supported Optional Features	75
3.3.5	Configuration Notes	79
3.4	RS/6000 Model H50 Product Description	79
3.4.1	Model H50 Highlights	80
3.4.2	Model H50 System Architecture	81
3.4.3	Standard Features	81
3.4.4	Supported Optional Features	84
3.4.5	Configuration Notes	89
Chapter 4.	Internal Storage Features	91
4.1	Small Computer System Interface	91
4.1.1	SCSI-I	91
4.1.2	SCSI-II	91
4.1.3	SCSI-III	93
4.1.4	SCSI Terminology	96
4.1.5	SCSI Repeaters	98
4.1.6	Summary of SCSI Specifications	98
4.1.7	Internal SCSI Devices	98
4.1.8	SCSI Adapters	100
4.1.9	SCSI Cabling	102
4.2	Serial Storage Architecture	115
4.2.1	SSA Functions	116
4.2.2	SSA Advantages	117
4.2.3	Internal SSA Devices	118
4.2.4	SSA Adapter	119
4.2.5	SSA Cabling	120

4.3 Hot Swap Subsystem	128
4.3.1 Components	128
4.3.2 Models F40 and F50 Hot Swap Disk Installation	133
4.3.3 Model H50 Hot Swap Disk Installation	136
4.3.4 SCSI Disk Carriers	138
4.3.5 SSA Disk Carriers	139
4.3.6 Using the Hot Swap Subsystem	140
Chapter 5. Adapters, Overview and Configuration	147
5.1 Communication Adapters	147
5.1.1 Token-Ring Adapters	147
5.1.2 Ethernet Adapters	148
5.1.3 X.25 Adapters	149
5.1.4 ISDN Adapters	150
5.1.5 FDDI Adapters	150
5.1.6 ATM Adapters	153
5.1.7 Multiple Port Adapters	154
5.1.8 Miscellaneous Communications Adapters	157
5.2 Graphics Adapters	158
5.2.1 MVP Power Multi-Monitor Graphics Adapter (# 2837)	158
5.2.2 POWER GXT110P (# 2839)	159
5.2.3 POWER GXT120P (# 2838)	159
5.2.4 POWER GXT250P (# 2851) and POWER GXT255P (# 2852)	160
5.2.5 POWER GXT500P (# 2854) and POWER GXT550P (# 2855)	161
5.2.6 POWER GXT800P (# 2853) and with Texture Option (# 2859)	162
5.2.7 POWER GXT1000 Internal (# 7252)	164
5.2.8 Graphics Adapter Feature Comparison	166
5.2.9 Multiple Adapter Support	169
5.2.10 Accessing the Second Display	171
5.3 Adapter and Device Configuration on PCI-Based RS/6000 Servers	172
5.3.1 Device Types	172
5.3.2 PCI Adapter Configuration	176
5.3.3 ISA Adapter Configuration	177
5.3.4 8-Port EIA-232 Asynchronous ISA Adapter (# 2931)	183
5.3.5 128-Port EIA-232 Asynchronous ISA Adapter (# 2933)	186
5.3.6 4-Port Multi-Protocol ISA Adapter (# 2701)	191
5.3.7 X.25 Interface Co-Processor ISA Adapter (# 2961)	197
Chapter 6. Diagnostics and System Configuration Management	205
6.1 Diagnostics Operating Considerations	205
6.2 Stand-Alone Diagnostics	207
6.2.1 Booting the Stand-Alone Diagnostics CD-ROM	207
6.2.2 Running Stand-Alone Diagnostics in Concurrent Mode	210
6.3 Online Diagnostics	210

6.3.1	Concurrent Mode	211
6.3.2	Service Mode	212
6.3.3	Single-User Mode	213
6.4	Understanding the Firmware	213
6.4.1	Firmware Boot Sequence	214
6.4.2	Firmware Passwords	216
6.4.3	Firmware Flash Update	217
6.4.4	Firmware Power-Up Problems	218
6.4.5	Firmware Recovery	219
6.4.6	Power-On Self Test	220
6.5	Function Keys - Key-Switches at Boot Time	221
6.6	System Management Service (SMS)	222
6.6.1	SMS Graphical Main Menu	223
6.6.2	SMS ASCII Main Menu	223
6.6.3	Start-Up Menu	224
6.6.4	Test Menu	224
6.6.5	Tools Menu	225
6.6.6	Accessing System Management Services with an ASCII Terminal	225
6.7	Service Processor	226
6.7.1	Service Aid Interface	226
6.7.2	Service Processor Menus	228
Chapter 7. AIX Common Hardware Reference Platform Support		231
7.1	Electronic Key-Switch Function	231
7.1.1	bootlist Command	232
7.2	Entering Debug Mode	232
7.3	Backup and Cloning	233
7.3.1	The bootinfo Command	233
7.3.2	Creating a System Backup	234
7.3.3	Restoring Your System Backup	234
7.3.4	Cloning	234
7.3.5	Cloning on AIX Version 4.2 and 4.3	238
7.4	Configuration Information	241
7.5	Network Boot Support	245
7.6	Using the Error Logging Facility	246
7.7	Booting Problems	246
7.7.1	Booting from CD-ROM	247
7.7.2	System Hangs during AIX Boot Process	247
7.7.3	Recovering a System with No Bootable Media	248
Appendix A. Special Notices		249
Appendix B. Related Publications		253

B.1 International Technical Support Organization Publications 253
B.2 Redbooks on CD-ROMs 253
B.3 Other Publications 253
B.4 Web Resources 254
B.5 Tools Disks 255

How to Get ITSO Redbooks 257
How IBM Employees Can Get ITSO Redbooks 257
How Customers Can Get ITSO Redbooks 258
IBM Redbook Order Form 259

List of Abbreviations 261

Index 265

ITSO Redbook Evaluation 267

Figures

1.	RS/6000 Hardware System Offerings	1
2.	RS/6000 Model E30	2
3.	RS/6000 Model F40	4
4.	RS/6000 Model F50	7
5.	RS/6000 Model H50	10
6.	IEEE 1394 Connections	19
7.	RS/6000 Logical Block Diagram (Model F50)	25
8.	PowerPC 604e Microprocessor Logical Block Diagram	31
9.	Memory Hierarchy	33
10.	SMP Cache Coherency Problem	34
11.	MESI Protocol — Steps 1 and 2	36
12.	MESI Protocol — Steps 3 and 4	37
13.	MESI Protocol — Steps 5 and 6	38
14.	False Sharing	39
15.	RISC System/6000 Model 7024-E30	41
16.	RS/6000 Model E30 Bays and Front View	43
17.	RS/6000 Model F40	52
18.	RS/6000 Model F40 Bays	55
19.	RS/6000 Model F40 Rear View	57
20.	RS/6000 Model F50	67
21.	RS/6000 Model F50 Architecture	69
22.	RS/6000 Model F50 Bays	73
23.	RS/6000 Model F50 Rear View	75
24.	RS/6000 Model H50 Installed in the S00 Rack	80
25.	RS/6000 Model H50 Bays	82
26.	RS/6000 Model H50 Rear View	83
27.	RS/6000 Model F50 S00 Rack, Front and Back Views	88
28.	SCSI-III Standards Overview	94
29.	SCSI Differential Example	97
30.	Piggy Back Connectors	104
31.	SSA Architecture	115
32.	SSA Loops, Links, and Data Paths	121
33.	7131 Base Configuration	124
34.	Loop Configuration with Four 7131	125
35.	Internal SSA Backplanes	126
36.	SSA Loop to Internal SSA Drives	126
37.	Loop to Internal and External SSA Drives	127
38.	Hot Swap Components	129
39.	Installing Hot Swap	131
40.	RS/6000 Model H50 Bays	132
41.	Front Bays	133

42.	Removing the Front Panel	134
43.	SSA Cabling	136
44.	Accessing the Hot Swap Bays	137
45.	Hot Swap Disk Drives	138
46.	Disk Drive Light	139
47.	SSA Hot Swap Disks	140
48.	Devices' Location Codes	175
49.	List of Configured ISA Adapters	179
50.	Attributes Listed Using the Isresource Command	179
51.	Attributes Listed Using the Isattr Command	180
52.	SMIT ISA Menu	181
53.	SMIT Configuration Menu for Adding an 8-Port Asynchronous ISA Adapter	185
54.	SMIT Configuration Menu for Adding an 128-Port Asynchronous ISA Adapter	189
55.	SMIT Configuration Menu for Adding an 4-Port Multi-Protocol ISA Adapter	196
56.	SMIT Configuration Menu for Adding an X.25 Interface Co-Processor ISA Adapter	202
57.	ISA Adapter Service AID Diagnostic Menu	209
58.	ISA Adapter Attribute Selection Diagnostic Menu	209
59.	Firmware Boot Sequence	214

Tables

1.	SPEC and LINPACK Performance	16
2.	Commercial Performance	17
3.	RS/6000 Model F50 TPC-C Published Results	17
4.	TPC-D Published Results	17
5.	DVD Capacity (GB)	21
6.	PowerPC and Bus Specification	29
7.	Optional Disk Drives on Model E30	45
8.	Optional Tape Drives on Model E30	46
9.	CD-ROM Drive on Model E30	46
10.	Supported Monitors for Model E30	47
11.	SCSI Adapter Connections on Model E30	49
12.	SCSI Adapter Connections on Model E30	60
13.	Overview of SCSI-III Standards	94
14.	Different SCSI Modes	98
15.	Optional Disk Drives on Model E30, F40, F50, and H50	99
16.	Optional CD-ROM Drives on Model E30, F40, F50, and H50	99
17.	Optional Tape Drives	100
18.	Single-Ended SCSI Adapter-to-First Device Cables	109
19.	Device-to-Device Cables for Single-Ended Installations	110
20.	Terminators for Single-Ended Installations	111
21.	Differential SCSI Adapter-to-First Device Cables	111
22.	Device-to-Device Cables for Differential Installations	112
23.	Terminators for Differential Installations	114
24.	SSA Cables	122
25.	7133 Cables	122
26.	7131 Cables	123
27.	SCSI Status LEDs	139
28.	SSA Status LEDs	140
29.	Entry-Level Graphics Adapters	166
30.	3D Graphics Adapters	168
31.	Accessing a Second Display	171
32.	8-Port Asynchronous ISA Adapter DIP Switches	183
33.	128-Port Asynchronous ISA Adapter DIP Switches	187
34.	Physical Interfaces on 4-Port Multiprotocol Interface Cable	191
35.	Maximum Cable Length	192
36.	DIP-Switches on 4-Port Multi-Protocol ISA Adapter	192
37.	4-Port Multiprotocol Adapter: DIP-Switch Suggested Settings	193
38.	DIP Switches on the X.25 Interface Co-Processor ISA Adapter	198
39.	X.25 Adapter: DIP-Switch Suggested Settings	198
40.	bootinfo -p	233

Preface

IBM continues to provide products and solutions to help you grow your business. This redbook discusses the entry level to midrange RS/6000 servers, namely the 7024-E30, 7025 Models F40 and F50, and 7026-H50.

The material provided in this publication provides a single-source reference for both technical and non-technical professionals who wish to better understand these RS/6000 server offerings.

This publication combines a generalized product description, a technical overview of each of the hardware models, and a discussion of the architectures that are responsible for the success of these servers. Topics include a review of memory architecture, hot swap disks, SSA-based disks, and bus technologies. Emerging technologies, such as Ultra-2, are also discussed.

As you will soon discover, these RS/6000 servers expand the range of possible applications while providing a balance between industry-leading performance and outstanding value.

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. The authors of this document are:

Verena Gschell	IBM Munich
Scott Vetter	IBM Austin

The project that produced this publication was coordinated by:

Vetter, Scott	IBM Austin
----------------------	------------

We would also like to acknowledge the professionals who took time to review this document and provided invaluable advice and guidance during its development:

Barker, Ron	IBM Dallas
Bex, Yves	IBM Toulouse
Bove, Henry	IBM Austin
Bunte, Bob	IBM Austin
Crisanto, Miguel	IBM Munich
Eckols, Joe	IBM Austin
Hiemke, Gary	IBM Austin

Hollaway, John	IBM Austin
Hurkadli, Basu	IBM Minneapolis
Kung, Johnathan	IBM Austin
Larsen, Ken	IBM Austin
Maule, Warren	IBM Austin
Mcauley, Dave	IBM ITSO Almaden
Osmond, Roger	IBM Markham
Prokop, Thomas	IBM Minneapolis
Robertson, Simon	IBM Basingstoke
Romero, Hector	IBM Austin
Smith, Bill	IBM Austin
Thiel, Uwe	IBM Munich
Upton, John	IBM Austin
Watts, Dave	IBM ITSO Raleigh
Wright, Hamid	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

IBM greatly simplified its RS/6000 product line, throughout 1997, 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.

This publication takes a look at four servers that represent two of the larger available workgroup servers and two of the entry-level enterprise servers. The RS/6000 Models E30, F40, F50, and H50 represent computing solutions that can be configured to fit almost any data processing budget.

1.1 RS/6000 Server Offerings

Figure 1 shows the current range of RS/6000 server offerings. The servers that are the topic of this publication are outlined (- - -).

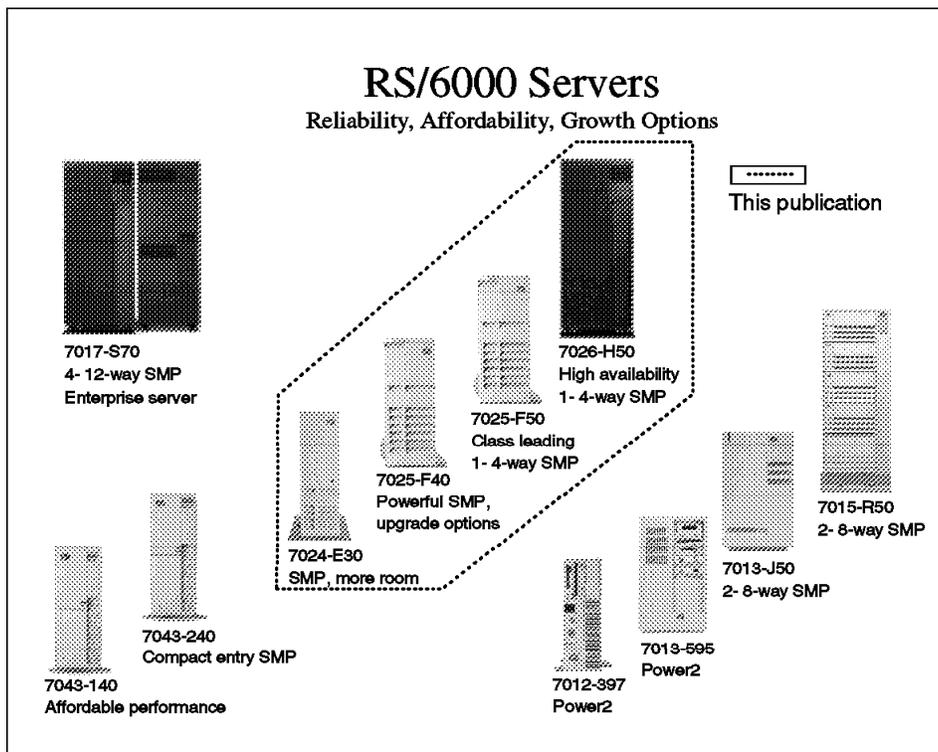


Figure 1. RS/6000 Hardware System Offerings

The selected servers share the following characteristics:

- Powered by the PowerPC 604e
- Contain PCI adapter slots
- Have ISA adapter slots
- Offer more capacity than desktop solutions
- Have available service processors (optional on E30/F40)

The RS/6000 naming convention is patterned similar 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 F40 and F50 are related, with the F50 being more powerful.

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

1.2 RS/6000 Model E30 Overview

The RS/6000 Workgroup Server Model E30 is a compact desk-side system with upgrades available for performance, storage, memory, and expansion slots. It is designed with the small business or work group in mind and is an intelligent choice for environments where a system needs to easily fit into a work area, but where future growth is a must. Figure 2 shows the RS/6000 Model E30 from the front.

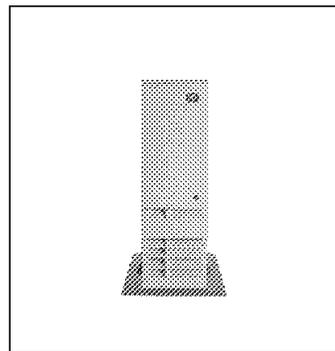


Figure 2. RS/6000 Model E30

1.2.1 RS/6000 Model E30 Key Attributes

The following is a list of the key attributes of the RS/6000 Model E30.

- Powerful server ideal for running small-business and departmental applications

- Small footprint to fit into your department or work area
- Choice of processor speeds to match your needs
- Ultra SCSI and SSA storage options for improved data transfer performance and availability
- Optional service processor for simplified remote monitoring and service
- Exceptional AIX operating system with reliability, availability, and systems-management features

The Model E30 comes in a smaller, more affordable package than the larger F-series workgroup servers, but it is still designed to help run many mission-critical business applications and networks twenty-four hours a day, seven days a week. It has reliability features typically found in systems that are much larger and more costly and can participate in clusters with other RS/6000 systems as a highly available server.

1.2.2 RS/6000 Model E30 Feature Summary

The following is a summary of the key RS/6000 Model E30 features:

- Standard configuration:
 - Microprocessor: 166 MHz PowerPC 604e (available until July 24, 1998)
 - Microprocessor: 233 MHz PowerPC 604e
 - Level 1 (L1) cache: 16 KB data/16 KB instruction
 - Level 2 (L2) cache: 1 MB on 166 MHz, 512 KB on 233 MHz PowerPC
 - RAM (memory): 64 MB ECC
 - Memory bus width: 64-bit
 - Memory slots: Eight DIMM
 - Ports: One parallel and two serial
 - Internal disk drive: 2.1 GB SCSI-2
 - Disk/media bays: Eight (five available)
 - Expansion slots: Eight (five PCI, two ISA, one PCI/ISA)
 - PCI bus width: 32-bit
- Standard features:
 - 20X (max) CD-ROM drive
 - 1.44 MB 3.5-inch diskette drive
 - Integrated SCSI-2 Fast/Wide controller

- AIX operating system:
 - Version 4.1.5, Version 4.2.1, or Version 4.3 (one- to two-user server license is standard)
- System expansion:
 - Processor upgrades (prior to July 24, 1998): 233 MHz PowerPC 604e (with 512 KB L2 cache)
 - RAM: Up to 1 GB
 - Internal Storage: Up to 27.0 GB
 - External disk storage: Up to 3.2 TB SCSI-2; up to 2.6 TB SSA
 - Service processor: Optional
- System dimensions:
 - 25.5" H x 7.7" W x 17.7" D (654.3 mm x 197.5 mm x 454.1 mm); 55 lbs (24.93 Kg)
- Warranty:
 - One year (limited)

1.3 RS/6000 Model F40 Overview

The RS/6000 Workgroup Server Model F40 is a desk side one-way or two-way symmetric multiprocessing (SMP) system with the power to run complex commercial, business, and engineering applications.

Designed with reliability, availability, and serviceability (RAS) features typically found in higher-priced systems; such as Error Checking and Correcting (ECC) memory, hot-swappable disks, internal RAID storage, and an internal service processor, the Model F40 can help keep your mission-critical business, technical applications, and networks running twenty-four hours a day, seven days a week. Figure 3 shows an RS/6000 Model F40.

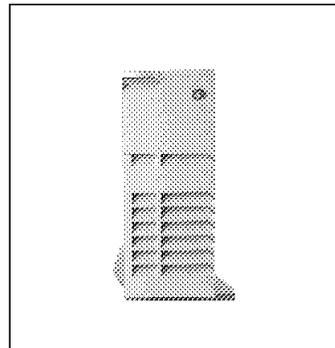


Figure 3. RS/6000 Model F40

1.3.1 RS/6000 Model F40 Key Attributes

The following is a list of the key attributes of the RS/6000 Model F40.

- Powerful SMP workgroup server ideal for collaborative computing
- Upgrade path to the next step in the SMP server family, the RS/6000 Model F50
- Easy plug-in processor design
- Outstanding expandability in a compact package
- Ultra SCSI and SSA storage options for improved data transfer performance and availability
- Built-in reliability features to help protect your valuable data
- Optional service processor for simplified remote monitoring and service
- Powerful graphics capabilities for demanding graphics workstation applications
- Exceptional AIX operating system with reliability, availability, and systems-management features

The Model F40 allows the seamless scaling of transaction processing power by adding a second processor to the 233 MHz standard processor. Multiple applications can be balanced between two processors to improve overall throughput. To convert from one-way processing to two-way processing, just plug in the second processor card.

The Model F40 has outstanding expansion capabilities within a single compact package just slightly more than two feet tall. The modular design makes adding disk storage easy. And, by using common industry storage interfaces, the Model F40 supports many available storage devices. Expandable system memory provides excellent processing capability for your business applications.

With its wide range of RS/6000 2D and 3D graphics accelerators, the Model F40 supports 2D business graphics applications through complex high-function, true-color 3D graphics applications. The entry POWER GXT120P graphics accelerator can support a graphical console as an alternative to the more traditional ASCII terminals used with many servers.

1.3.2 RS/6000 Model F40 Feature Summary

The following list defines the key RS/6000 Model F40 features:

- Standard configuration:
 - Microprocessor: 233 MHz PowerPC 604e

- Level 1 (L1) cache: 32 KB data/32 KB instruction
- Level 2 (L2) cache: 1 MB per processor
- RAM (memory): 64 MB ECC DIMM
- Memory bus width: 64-bit
- Memory slots: Eight
- Ports: Tablet, keyboard, mouse, Ethernet (thick + twisted pair), SCSI-2 F/W, serial (two), parallel, and stereo audio
- Internal disk drive: 2.2 GB SCSI-2 F/W
- Disk/media bays: Eighteen (one used) / four (two used)
- Expansion slots: Nine (32-bit: five PCI + two shared PCI/ISA; 64-bit: two PCI)
- PCI bus width: 32- and 64-bit
- Standard features:
 - LED operator panel
 - Dual Integrated SCSI-2 Fast/Wide controllers (one externalized)
 - 20X (Max) SCSI-2 CD-ROM drive
 - Integrated Ethernet Controller (10BaseT/10Base 5)
 - 1.44 MB 3.5-inch diskette drive
 - Integrated stereo audio
- AIX operating system:
 - Version 4.1.5, Version 4.2.1, or Version 4.3 (one- to two-user server license is standard)
- System expansion:
 - For 2-way SMP configuration: Additional processor of same MHz
 - RAM: Up to 1 GB
 - Internal disk storage: Up to 172.8 GB (163.8 GB hot-swappable)
 - External disk storage: Up to 4.8 TB SCSI; up to 2.6 TB SSA
 - Service processor: Optional
-

System dimensions:

- 24.3" H x 9.6" W x 27.3" D (620 mm x 245 mm x 695 mm); 75 lbs (35 Kg)
- With pedestal: 24.3" H x 13.7" W x 29.3" D (620 mm x 350 mm x 745 mm)
- Warranty: One year (limited)

1.4 RS/6000 Model F50 Overview

The RS/6000 Model F50 is a one- to four-way SMP designed to deliver the performance and reliability for your critical applications and the flexibility to add the features important to your business. This system excels as a multiuser application, database, and Internet server.

It has the connectivity that allows it to participate in most currently installed UNIX and PC networks. From the simplest file/print serving to the most sophisticated e-business server applications in your organization, the Model F50 is ready. Figure 4 shows an RS/6000 Model F50. If the outside appearance appears familiar to the Model F40, you are not mistaken. The Model F40 and Model F50 share nearly identical packaging. Inside, the system planar, processor, and I/O cards are very different.

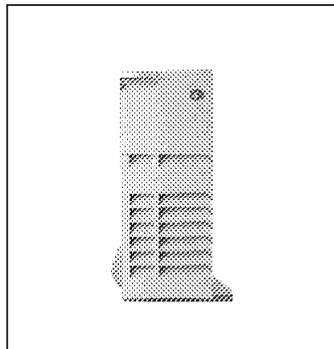


Figure 4. RS/6000 Model F50

1.4.1 RS/6000 Model F50 Key Attributes

The following is a list of the key attributes of the RS/6000 Model F50.

- Powerful one- to four-way symmetric multiprocessor (SMP) desk-side server with superior commercial application price/performance
- Two processor speeds available — 166 MHz and 332 MHz
- Industry-leading decision support performance
- Affordable upgrades, including additional processors and memory
- Outstanding reliability with ECC memory and ECC Level 2 cache, plus a standard service processor

- Hot-swappable disk bays for improved system availability, smooth growth, and data portability
- Exceptional AIX operating system with reliability, availability, and systems-management features
- E-business server with balanced price/performance

The 332 MHz processor in the F50 brings outstanding performance across a variety of benchmarks, especially with decision support performance (Transaction Processing Council TPC-D at 100 GB as of 1/26/98), where the F50/332 has achieved industry leadership for four-way systems (see 1.6, “Performance Results” on page 12 for more performance information). The F50/332 MHz has an excellent Relative OLTP (ROLTP) commercial performance value of 32.8, the highest four-processor configured server in the RS/6000 lineup. The SPECint95 performance is also outstanding, the highest of all RS/6000 servers.

The F50 system with the 332 MHz processor has an excellent SPECweb96 performance rating as well, placing the F50 among the industry leaders for four-way SMP systems. As a result of its exceptional performance, expandability, and reliability, the Model F50 is well suited as a commercial system for high-demand applications, including mission-critical OLTP, e-business, and collaborative computing, such as Lotus Domino environments.

RS/6000 servers are undergoing continuous enhancement, therefore refer to the following locations for the latest performance information.

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

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

TPC <http://www.tpc.org>

All in all, no matter what type of application performance your requirements demand, the F50 is an outstanding choice for workgroup or enterprise computing.

1.4.2 RS/6000 Model F50 Feature Summary

The following list defines the key RS/6000 Model F50 features:

- Standard configuration:
 - Microprocessors: 166 MHz or 332 MHz PowerPC 604e with X5 cache
 - Level 1 (L1) cache: 32 KB data/32 KB instruction
 - Level 2 (L2) cache: 256 KB ECC

- RAM (memory): 128 MB ECC Synchronous DRAM
- Memory bus width: 128-bit
- Memory slots: Two
- Ports: Three serial and one parallel
- Internal disk drive: 4.5 GB SCSI Fast/Wide
- Disk/media bays: Eighteen (one used)/four (two used)
- I/O expansion slots: Nine (seven PCI, two PCI/ISA)
- PCI bus widths: Two 32-bit and one 64-bit
- Standard features:
 - 20X (Max) SCSI-2 CD-ROM drive
 - 1.44 MB 3.5-inch diskette drive
 - Ethernet 10Base5/T controller
 - Dual integrated SCSI-2 Fast/Wide controllers
 - Service processor
 - GXT800P graphics adapter option
- AIX operating system:
 - Version 4.2.1 or Version 4.3 (one- to two-user license is standard)
- System expansion:
 - SMP configurations: Two, three, or four 166 MHz or 332 MHz processors (different speed processors cannot be mixed)
 - RAM: Up to 3 GB
 - Internal disk storage: Up to 172.8 GB (163.8 GB hot-swappable)
 - External disk storage: Up to 4.8 TB SCSI-2; up to 3.5 TB SSA
- System dimensions:
 - 24.3" H x 9.6" W x 27.3" D (620 mm x 245 mm x 695 mm); 75 lbs (35 Kg)
- Warranty:
 - One year (limited)

1.5 RS/6000 Model H50 Overview

The RS/6000 Enterprise Server Model H50 is a one- to four-way SMP designed to deliver the performance and reliability for your critical applications and the flexibility to add the features important to your business. This system excels as a multiuser application, database, and Internet server. It has the connectivity that allows it to participate in many currently installed UNIX and PC networks. Figure 5 shows an RS/6000 Model H50.

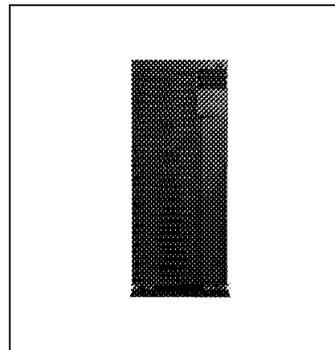


Figure 5. RS/6000 Model H50

1.5.1 RS/6000 Model H50 Key Attributes

The following is a list of key attributes:

- Powerful one- to four-way symmetric multiprocessor (SMP) rack-drawer server with superior price/performance
- Industry-leading decision support performance
- Affordable upgrades, including additional processors, memory, and storage
- Outstanding reliability with ECC memory and ECC Level 2 cache, plus a standard service processor, redundant hot-swappable cooling fans, and an available redundant hot-swappable power supply
- Hot-swappable disk bays for improved system availability, smooth growth, and data portability
- Exceptional AIX operating system with reliability, availability, and systems-management features
- E-business server with balanced price/performance

The Model H50 system is an eight-EIA high rack drawer for mounting in a 19-inch rack enclosure, such as IBM's S00 rack (separately purchased). Clustering an H50 with additional H50 drawers and shared-disk subsystems provides outstanding shared data availability and flexibility.

The Model H50 has achieved excellent performance benchmark results. Along with its sister product, the desk-side Model F50/332, the H50 has

captured industry leadership in decision support performance (TPC-D at 100 GB as of 1/26/98) for four-way servers. (see 1.6, "Performance Results" on page 12 for more performance information). Relative OLTP (ROLTP) commercial performance is excellent at a value of 32.8, the highest four-processor configured server in the RS/6000 lineup (equivalent to the Model F50/332 MHz). SPECint95 is also outstanding, the highest of all RS/6000 servers. SPECweb96 results place the Model H50 among the industry leaders. All in all, no matter what type of application performance your requirements demand, the H50 is an outstanding choice for workgroup or enterprise computing.

RS/6000 servers are undergoing continuous enhancement, therefore refer to the following locations for the latest performance information.

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

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

TPC <http://www.tpc.org>

1.5.2 RS/6000 Model H50 Feature Summary

The following list defines the key RS/6000 Model H50 features:

- Standard configuration:
 - Microprocessor: 332 MHz PowerPC 604e with X5 cache
 - Level 1 (L1) cache: 32 KB data/32 KB instruction
 - Level 2 (L2) cache: 256 KB ECC
 - RAM (memory): 128 MB ECC Synchronous DRAM
 - Memory bus width: 128-bit
 - Ports: Three serial and one parallel
 - Internal disk drive: 4.5 GB SCSI Fast/Wide
 - Disk/media bays: Thirteen (one used)/three (two used)
 - I/O expansion slots: Nine (seven PCI, two PCI/ISA)
 - PCI bus widths: Two 32-bit and one 64-bit
 - Memory slots: Two
- Standard features:
 - 20X (Max) CD-ROM drive
 - Ethernet 10Base5/T controller
 - Service processor

- 1.44 MB 3.5-inch diskette drive
- Dual integrated SCSI-2 Fast/Wide controllers
- AIX operating system:
 - Version 4.2.1 or Version 4.3 (one- to two-user server license is standard)
- System expansion:
 - SMP configurations: Two, three, or four processors
 - RAM: Up to 3 GB
 - Internal disk storage: Up to 118.2 GB
 - External disk storage: Up to 4.8 TB SCSI-2; up to 3.5 TB SSA
- System dimensions:
 - 13.8" H x 17.5" W x 33.2" D (350 mm x 443 mm x 844 mm); 120 lbs (55 Kg)
- Warranty:
 - One year (limited)

The ability to easily cluster the H50 using HACMP makes the H50 a secure, reliable, comprehensive system for business critical applications including Web serving and data mart applications. It is also an ideal departmental server for commercial customers like banks or insurance companies, where processing power, transaction processing capability and reliability are at a premium.

1.6 Performance Results

This section contains the results of several industry standard performance measurements for the RS/6000 Models E30, F40, F50, and H50.

Values shown here are the results of production-level systems. All performance data was obtained in a specific environment and is presented *as is* for illustrative purpose only. While these values should be indicative of customer systems, no warranties or guarantees are stated or implied by IBM. IBM recommends application-oriented testing for performance predictions and offers the following reported benchmarks only as an initial indicator. Additional information on these tests is available from IBM Local Branch Office, IBM Authorized Reseller, or from one of the following locations.

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

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

TPC <http://www.tpc.org>

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

SPECint95	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_base_rate95	Geometric average of the eight SPEC rates from the SPEC integer tests (CINT95) with the restrictive compiler options.
SPECfp95	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_base_rate95	Geometric average of the eight 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 v.1.1 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).

The following Transaction Processing Council (TPC) benchmarks reflect the performance of the microprocessor, memory subsystem, disk subsystem, and some portions of the network.

TPC Benchmark C	Five OLTP transactions against nine databases so that 90 percent of the transactions finish ≤ 5 seconds.
TPC Benchmark D	Models 17 ad hoc queries representing decision support workload in a multiple user environment.
tpmC	The result of the TPC Benchmark C measured as the average transactions per minute during a valid twenty minute TPC-C run.
\$/tpmC	The TPC Benchmark C price-performance ratio reflects the estimated five year total cost of ownership for system hardware, software, and maintenance and is determined by dividing the estimate by the tpmC for the system.
QppD	This is the power metric of TPC-D and is based on a geometric mean of the 17 TPC-D queries, the insert test, and the delete test. It measures the ability of the system to give a single user the best possible response time by harnessing all available resources. QppD is scaled based on database size from 30 GB to 1 TB.
QthD	This is the throughput metric of TPC-D and is a classical throughput measure characterizing the ability of the system to support a multi-user workload in a balanced way. A number of query users is chosen, each of which must execute the full set of 17 queries in a difficult order. In the background, there is an update stream that runs a series of insert/delete operations. QthD is scaled based on the database size from 30 GB to 1 TB.
\$/QphD	This is the price/performance metric for the TPC-D benchmark where QphD is the geometric mean of Qppd and QthD. The price is the five year cost of ownership for the tested configuration and includes maintenance and software support.

Relative OLTP (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 system 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. Unless otherwise indicated for a system, ROLTP is

estimated only at the time the system is introduced. An IBM RS/6000 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.

1.6.1 SPEC95 Software Benchmark

SPEC95 is a forward step in the performance measurement of the core of a system. It covers the CPU, caches, memory, and compiler. The programs and data sets which 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 computationally-intensive workloads on different computer systems.

SPEC95 contains two suites of benchmarks:

- CINT95** Measures and compares computationally-intensive integer performance
- CFP95** Measures and compares computationally-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, the 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 metrics* (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.

1.6.2 Performance Summary

Table 1 contains the performance data for Models E30, F40, F50, and H50.

<i>Table 1. SPEC and LINPACK Performance</i>											
Model	CPU	Clock Rate (MHz)	L1 Cache (KB)	L2 Cache (MB)	SPEC int95	SPEC int_base95	SPEC fp95	SPEC fp_base95	LINPACK DP	LINPACK SP	LINPACK TPP
E30	604e	166	32/32	1.0	6.19	5.93	4.77	4.62	32.8	74.4	124.0
E30	604e	233	32/32	0.5	9.41	8.46	6.01	5.71	59.4	108.0	153.5
F40	604e/1	233	32/32	1.0	8.71	7.80	5.87	5.60	48.5	109.1	145.6
F50	604e/1	166	32/32	0.25	7.52	6.79	8.52	8.11	70.2	98.9	166.4
F50	604e/1	332	32/32	0.25	14.4	14.0	12.6	12.1	115.7	158.5	273.4
H50	604e/1	332	32/32	0.25	14.4	14.0	12.6	12.1	115.7	158.5	273.4

Table 2 on page 17 contains the commercial performance data for Models E30, F40, F50, and H50.

Table 2. Commercial Performance

Model	Processor	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	SPEC web96 ops/sec
E30	604e	166	32/32	1.0	3.7	—	—	—	—	—
E30	604e	233	32/32	0.5	4.7	—	—	—	—	—
F40	604e/1	233	32/32	1.0	3.7	78.0	69.9	52.0	50.1	—
F40	604e/2	233	32/32	1.0	5.2	151	132	95.2	89.5	—
F50	604e/1	166	32/32	0.25	8.2	67.5	61.0	76.5	72.8	—
F50	604e/2	166	32/32	0.25	14.9	135	121	149	143	—
F50	604e/3	166	32/32	0.25	21.0	—	—	—	—	—
F50	604e/4	166	32/32	0.25	27.1	267	241	283	267	2148
F50	604e/1	332	32/32	0.25	10.0	131	126	113	109	—
F50	604e/2	332	32/32	0.25	17.9	260	249	217	206	—
F50	604e/3	332	32/32	0.25	25.2	388	369	307	292	—
F50	604e/4	332	32/32	0.25	32.8	510	490	387	366	2755
H50	604e/1	332	32/32	0.25	10.0	131	126	113	109	—
H50	604e/2	332	32/32	0.25	17.9	260	249	217	206	—
H50	604e/3	332	32/32	0.25	25.2	388	369	307	292	—
H50	604e/4	332	32/32	0.25	32.8	510	490	387	366	2755

Table 3 contains the TPC-C data for the RS/6000 Model F50.

Table 3. RS/6000 Model F50 TPC-C Published Results

Model	Processor	Clock Rate (MHz)	L2 Cache (MB)	tpmC	\$/tpmC	Database	AIX Version	Date
F50	604e/4	166	.25	8,142.40	62.71	Sybase 11.5	4.2.1	2/09/98
F50	604e/4	332	.25	9,853.13	64.22	Sybase 11.5	4.2.1	2/20/98

Table 4 contains the TPC-D data for Models F50 and H50.

Table 4. TPC-D Published Results

Model	Processor	Clock Rate (MHz)	Power QppD @ 100 GB	Thruput QthD @ 100 GB	Price Perf. \$/QphD @ 100 GB	RDBMS	AIX Version	Date
F50	604e/4	166	833.2	366.8	714	DB2 UDB	4.2.1	3/31/98
F50	604e/4	332	1168.1	499.1	469	DB2 UDB	4.2.1	3/31/98
H50	604e/4	332	1168.1	499.1	489	DB2 UDB	4.2.1	3/31/98

1.7 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.

1.7.1 PowerPC Futures

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 PowerPC 630, or POWER3 architecture will provide a significant impact to the market by:
 - Providing SMP scalability not found with POWER2
 - Redefining the versatility of a desk side or departmental computing solution by providing outstanding float-point performance with solid integer performance
 - Providing an affordable entry-level 64-bit solution
- The RS64 (7017-S70 processor) will see accelerated development to increase its integer performance and overall throughput without a major impact on scalability.
- The PowerPC 604e will continue to reach toward 400 MHz as sub .25 micron levels and copper technologies are exploited.

1.7.2 Universal Serial Bus

The Universal Serial Bus (USB) is a plug and play peripheral connection for devices such as keyboards, mice, joysticks, scanners, printers, and digital cameras, to name a few. USB could replace the serial and parallel ports with a single bus. Hot swapping is supported, so your workstation will not need a reboot when swapping USB devices. The OS detects the peripherals and configures the necessary software.

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.

1.7.3 IEEE 1394

The IEEE 1394 standard (also known as FireWire) is new bus technology for the digital age. Over 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 (from DVD devices, for example) quickly, reliably, and at a low cost. The architecture is scalable and hot-pluggable. Industry consortiums are moving toward a 1-Gigabit data transfer rate.

Like the 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 6.

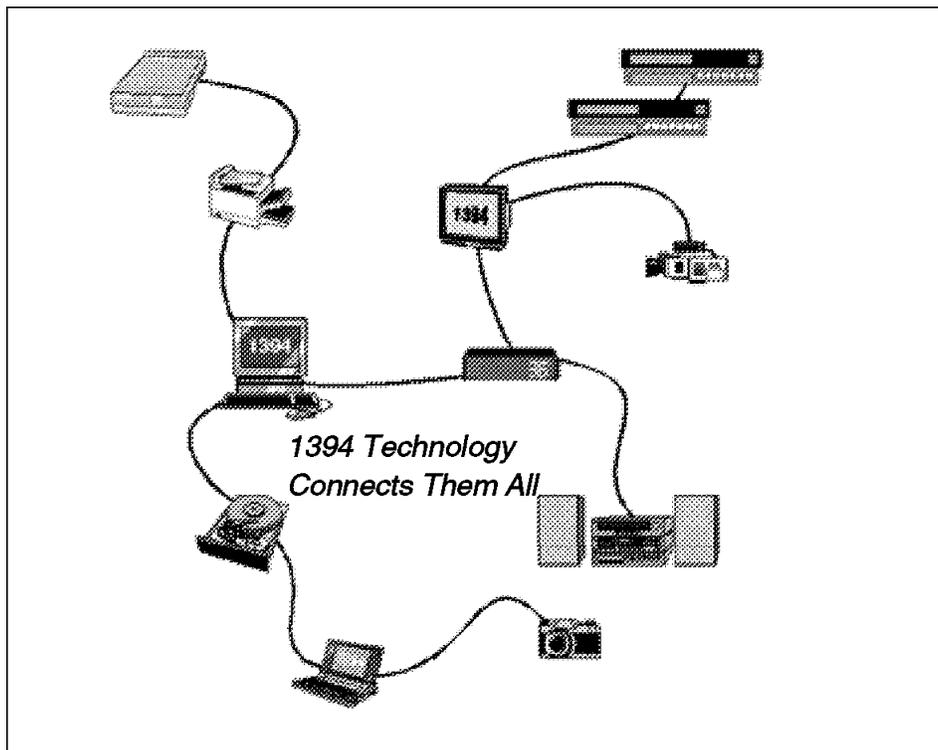


Figure 6. IEEE 1394 Connections

1.7.4 Storage Devices

The following descriptions are of storage devices that are under consideration for use in future RS/6000 systems.

SCSI Disks

IBM continues to enhance its storage technology. The arrival of 18.2 GB DASD, first available as external devices, will provide high volume storage options that allow your system to grow without increasing the footprint. With a platter speed of 7200 rpm and a seek time of 6.5 ms for read operations, these SCSI devices are easily integrated.

In the future, platter speeds of 10,000 rpm are possible for RS/6000 solutions.

Ultra-2 SCSI

Ultra-2 is the latest extension to the SCSI-3 specification which allows transfers to take place at 80 MBps on a wide bus, or 40 MBps 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).

Serial Storage Architecture

Serial Storage Architecture (SSA) will continue to be IBM's premier storage solution. New adapters with additional initiators will offer you additional configuration options.

DVD-ROM

Digital Versatile Disk (known originally as digital video disk, but now known just as DVD) represents a dramatic step forward in digital storage. It offers a combination of greater capacity than CD-ROM together with interactive video and audio function. The following discussion introduces the basic principals required when understanding DVD.

The appearance of the DVD is similar to a CD, but what you do not see is the technology that allows up to 26 times more data storage than a CD. This additional capacity allows faster data rate, incredible sound, and multiple languages, for example.

A major contribution to the storage capacity is made through reducing the size and space between the individual storage elements, allowing more data to be fit into a given area. Also, DVDs can have two layers of data on any side of the media. The two layers are separated by a translucent boundary which provides the read head two distinct areas to focus on.

Table 5 on page 21 shows the capacity achievable with this media.

<i>Table 5. DVD Capacity (GB)</i>		
Disk Type	8 cm Disc	12 cm Disc
Single-Sided, Single Layer	1.4	4.7
Single-Sided, Dual Layer	2.6	8.5
Double-Sided, Single Layer	2.9	9.4
Double-Sided, Dual Layer	5.3	17

CD has a maximum storage capacity of around 730 MB. DVDs can be produced with eight different capacities. They can have one or two layers of data, and can be single or double sided. Like CD, DVD can also be produced in two diameters; 8 cm and 12 cm. DVD players will be able to read all of the disc types, and existing CDs.

Chapter 2. Hardware Overview

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

- PowerPC Microprocessor Common Hardware Reference Platform
- Hardware Design
- SMP Overview

Keep in mind that specific details in the figures and supporting discussions may not apply to all the RS/6000 servers covered in this publication. However, the general architectures discussed in these sections help build an overall appreciation for RS/6000 hardware.

2.1 PowerPC Microprocessor Common Hardware Reference Platform

The PowerPC Microprocessor Common Hardware Reference Platform, previously known as the CHRP, is a superset of the previous PowerPC Reference Platform Specification (PRep) system architecture. To shorten this name, the industry has adopted the PowerPC Platform. The RS/6000 Model F50 and RS/6000 Model H50 (described in this publication) and the RS/6000 Model S70 are IBM's first CHRP-compliant machines. The RS/6000 Model E30 and RS/6000 Model F40 are not CHRP compliant, but contain several of the design points that are discussed in this section.

2.1.1 The PowerPC Platform Introduction

An introduction to the CHRP architecture will help you better understand the interactions between AIX and RS/6000 hardware. To correct the problems facing customers and developers, Apple, IBM and Motorola looked at various ways of combining the two hardware architectures into a common system architecture. In November 1994, Apple, IBM and Motorola agreed to develop a specification for a common hardware platform with the purpose of defining a system which will become the pervasive open industry standard from single user to server configurations. Finally, on November 13, 1995, the three companies announced the availability of the PowerPC Platform.

Apple, IBM, and Motorola have benefited from the design of this specification, but not as originally planned. It was thought that the architecture would allow multiple operating systems to run across multiple hardware platforms. Although both IBM and Apple have produced compliant machines, MAC OS is not supported on RS/6000, and AIX is not supported on all Apple systems.

PowerPC Platform Document

The PowerPC Platform document is a superset of the PowerPC Reference Platform Specification, Version 1.1, the Apple RISC Architecture (Power Macintosh) and IBM RISC server systems.

2.1.2 PowerPC Platform Achievements

The PowerPC Platform specification achieves the following:

- Creates an open industry standard to be used for the implementation of PowerPC-based systems. The architecture document is available to the industry and can be used by any hardware or software vendor to develop compliant products.
- 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 Standard for Boot (Initialization Configuration) Firmware, Core Requirements and Practices. Currently, the abstraction approach for some operating systems use platform description information discovered by a legacy boot process that is passed to the operating system in data structures. With these systems, operating systems and platforms will migrate to the Open Firmware boot process and device tree.
- Architects the control of power management by different operating systems. 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 which can evolve as technology changes. The creators of the architecture invite industry participation in evolving future versions of it.
- Architects a mechanism for error handling, error reporting, and fault isolation. The architecture allows the implementation of more robust systems if desired by the system developers.

PowerPC Platform on the Internet

A *Personal use Only* copy of the PowerPC Platform can be found at:

<http://www.rs6000.ibm.com/resource/technology/>

Morgan Kaufmann Publishers provides paper versions of the specification, ISBN 1-55860-394-8.

2.2 The Hardware Design

The designs shared by the servers discussed in this publication appear similar to Figure 7.

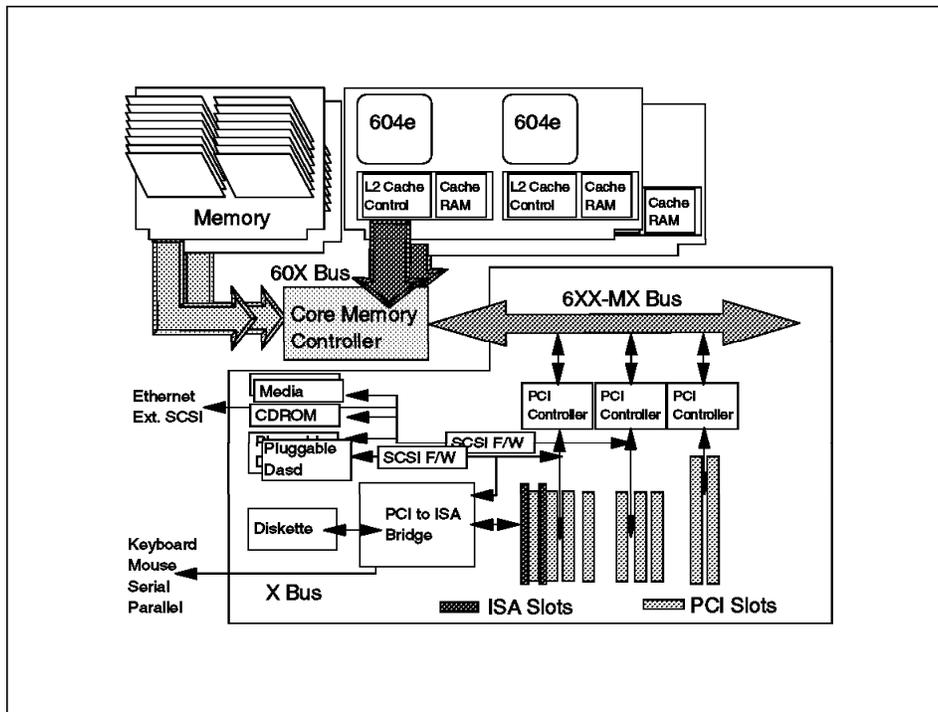


Figure 7. RS/6000 Logical Block Diagram (Model F50)

The processors and memory are located on removable cards. These cards fit into specialized slots on the system planar, or motherboard. A memory controller (in some designs containing an integrated PCI bridge) handles the traffic from the PCI adapters and other I/O with the processors and memory. 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.

Internal ISA, SCSI, and other devices are all interconnected to a PCI bus. Think of these internal devices as needing a PCI SCSI card or perhaps ISA support, but having the function integrated on the motherboard instead of requiring a separate adapter.

2.2.1 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 evolved over time. This is much like operating systems which 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.

The group defined PCI to support the high-performance basic system I/O devices, such as the graphics adapter, hardfile controller, or LAN adapter. In the original definition, these would be mounted on the planar and would communicate through the PCI bus. Current I/O buses (ISA, EISA and Micro Channel) would be used to attach pluggable features to configure the system for the desired use. The first release of PCI specification was made available in June of 1992.

The PCI-SIG soon realized that the PCI bus needed the capability to support connectors. For example, display controller evolution does not necessarily match planar development; so providing for an upgrade of the display controller became a requirement. The next release of the PCI Specification (Version 2.0 in April of 1993) included upgrade capability through expansion connectors.

According to PCI Specification Version 2.0, the PCI bus operates on 32- or 64-bits of data at a clock speed of 33 MHz (IBM's implementation of 64-bit slots uses a clock speed of 50 MHz). This yields a local bus performance of 132 MBps for 32-bit transfers and 264 MBps for 64-bit transfers. The next PCI Specification (Version 2.1) is expected to include a definition of 66 MHz PCI capability, increasing local bus performance to 528 MBps for 64-bit transfers.

Though each PCI bus is restricted to a maximum of four slots, the assembly of additional PCI bridges on the motherboard can allow a machine to be designed with more than four slots. Each PCI bridge adds another PCI bus, which in turn can handle up to four slots each.

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

High Data Transfer Speed: The high-speed data transfer is implemented by the following functions:

- Buffering and asynchronous data transfer

The PCI chip can support the processing and buffering of data and commands sent from the processor or from the peripherals in case the peripheral or the processor is not yet ready to receive the information.

- Burst mode transfer

Variable length linear or toggle mode bursting for both reads and writes improves write-dependant graphics performance.

- Caching

To reduce the access time, the PCI bus architecture supports caching of data which is frequently used.

- DMA

The Direct Memory Access (DMA) function is used to enable peripheral units to read from and write to memory without sending a memory request to the processor. This function is very useful for peripherals

that need to receive large amounts of data, such as video adapters, hard disks, and network adapters.

Processor Independence: Processor independence allows manufacturers to implement PCI buses on any computer. Any PCI-compliant peripheral will work on any PCI-compliant bus implementation.

Cross-Platform Compatibility: The key to cross-platform compatibility is processor independence. Until PCI, different systems used different buses, such as ISA, EISA, NuBus, and so forth. Now, different systems can use one bus.

Multi-bus Support: An important aspect to PCI-based system architecture is support for multiple PCI buses, operating transparently to existing software.

Plug and Play: PCI peripherals, following the PCI standard, load the appropriate set of installation, configuration, and booting information to the host CPU without user intervention. This provides a greater ease of use for the system integrator or end user.

Investment Protection: The PCI bus is designed for 64-bit addressing support.

2.2.2 The ISA Bus Architecture

The Industry Standard Architecture (ISA) is the most widely used system bus in the PC industry. Originally, there were no official definitions or standards for it. Later on, its specifications were defined by the Institute of Electrical and Electronics Engineers (IEEE) standards group.

The ISA bus, if implemented, allows a transfer rate of up to 8.3 MBps. 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 two bytes may be transferred during each transaction.

Since ISA is the most widely bus architecture used in the PC industry, it makes sense to provide users with the possibility to use hardware peripherals they already use with other systems, rather than having to expend additional money for these peripherals.

Moreover, supporting the ISA bus architecture, the system-provider has access to a wide spectrum of adapters and devices already available in the marketplace and does not have to wait for adapters to be built for a specific system bus. The provider has to ensure that the device driver for the specific operating system is also available.

The problem when connecting the processor to the ISA bus directly is that the processor's speed is reduced to match the slow ISA bus speed; thus the systems cannot take advantage of a fast processor.

Table 6 shows the bus specification for different architectures and compares them to the PowerPC processor's speed.

Specification	PowerPC	PCI	ISA(8)	ISA(16)	EISA
Processor Speed	166 233 332	-	-	-	-
Data Bus	64	64	8	16	16/32
Address Bus	32	64	20	24	24/32
Bus Clock	50 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

The solution to this problem is to use the PCI local bus 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.

2.2.3 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.

The 60X bus is designed as part of the 601, 603, and 604 series of processors. The 6XX bus is a variation that extends the function for 6XX processors and doubles the bus size. Several server designs use the 60X

bus only to the L2 cache and interconnect the processor/cache with a larger 6XX bus. The L2 cache controller can provide the translation of a 60X bus to a 6XX bus. The X5 chip is an example of this.

It is through this bus that the 2- to 4-way SMP's processors are attached, memory is accessed, and cache snooping is achieved at speeds up to 166 MHz.

2.2.4 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. Peak transfer rates beyond 400 MBps are possible but not yet implemented in any current RS/6000 hardware.

2.2.5 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 designed 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.

2.2.6 The Processor Subsystem

The PCI-based RS/6000 servers feature 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 superscaler 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 8 shows the PowerPC 604e microprocessor architecture.

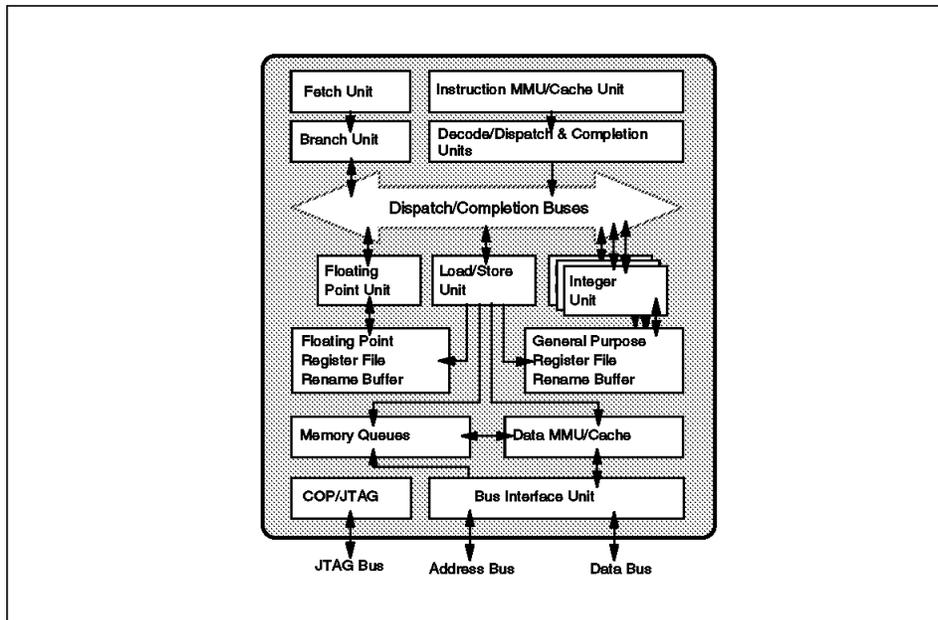


Figure 8. PowerPC 604e Microprocessor Logical Block Diagram

A 604e is defined by the following specifications:

- PowerPC 604e microprocessor running at:
 - 233 MHz on RS/6000 Model E30 (166 MHz available prior to July 24, 1998)
 - 233 MHz on RS/6000 Model F40
 - 166/332 MHz on RS/6000 Model F50

- 332 MHz on RS/6000 Model H50
- Up to 87 MHz memory bus clock
- 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 four petabytes (2^{52})
- Real memory support for up to four gigabytes
- ~5.1 million transistors
- Nap power management mode
- JTAG bus (IEEE 1149.1) for RAS

2.2.7 SMP Hardware Characteristics

Sharing resources is probably the main technical issue in the design of an 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.

Memory Hierarchy

In order to improve the hardware performance of a system (UP or SMP), 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 extremely fast but very small, and they have a high cost per bit. At the other end, you will find the disks, which are very slow but have a very low cost per bit. This allows a very high disk storage capacity.

In most UP or SMP implementations, you will find between these two ends a first level of cache (L1), which is a fast memory with a small capacity. The number of CPU cycles that are needed for the processor to load data from L1 depends on 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 also a very fast memory, faster than the main memory. It takes around one to ten cycles to load data from L2. The X5 cache controller uses special 10 ns memory to allow minimal latency. Its capacity is usually higher than L1 at around 256 KB to 4 MB.

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.

A cache is, in fact, a high-speed memory that holds a small subset of the main memory. Because of its short access time, it improves data *locality* (data is closer to the processor) and significantly increases system performance for frequently used instructions and data that are stored in the cache. When using caches, the increase in performance depends on the workload. For example, an L2 cache can be very beneficial with a commercial workload and not very beneficial with a technical workload. Figure 9 shows the relationship between proximity of memory to the processor versus the number of cycles required to access it.

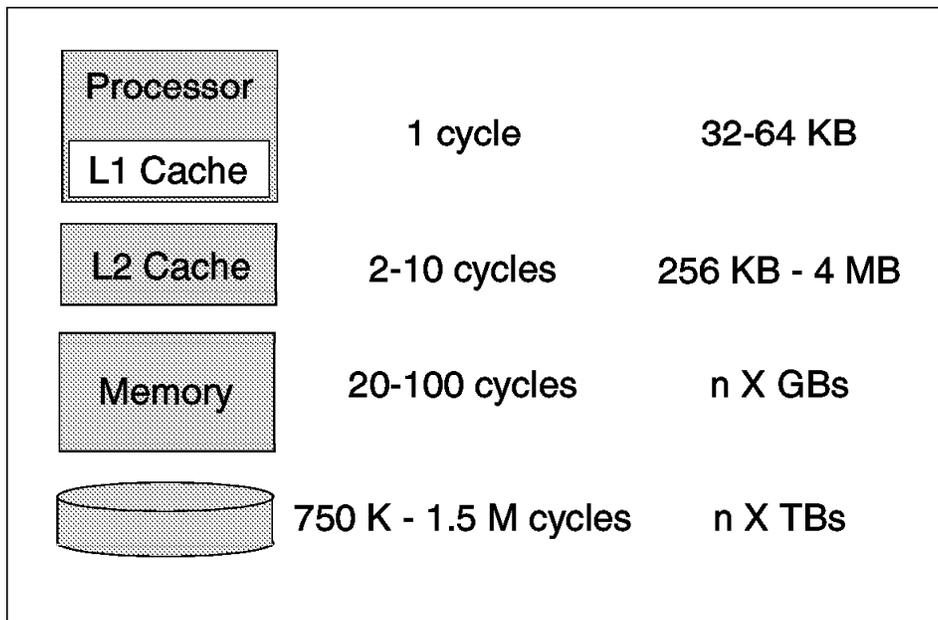


Figure 9. Memory Hierarchy

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.

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. Cache coherency 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 10.

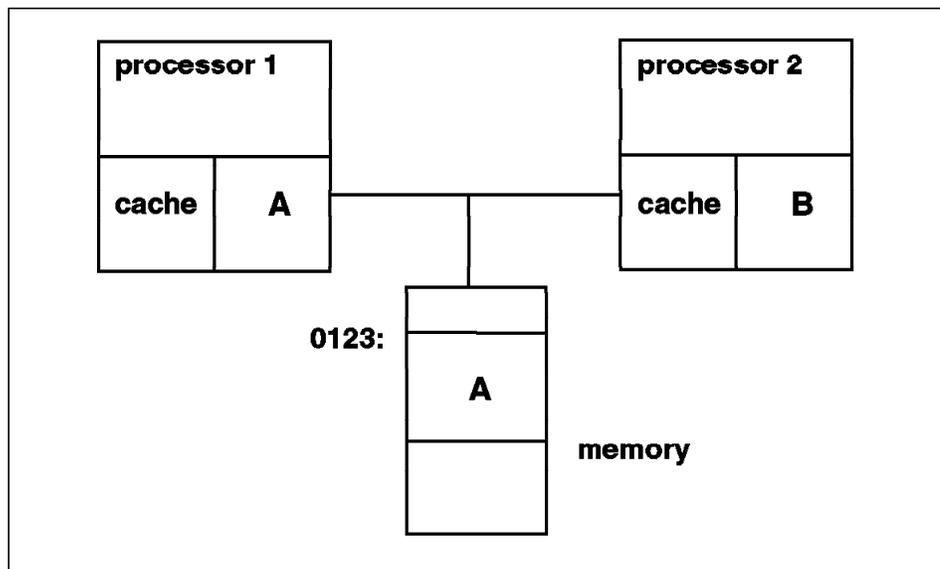


Figure 10. SMP Cache Coherency Problem

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 again load address 0x0123, 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.

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 better thing to do in an attempt to improve response time for a given request.

All PowerPC processors except the 603 have this extra logic; the POWER and POWER2 processors do not.

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

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 which is 64 bytes. This cache line is divided into two 32-byte sectors.

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.

This is best explained by means of an example. Assume there are two processors, processor 1 and processor 2, sharing the same memory, and there are different operations that can occur between these two processors. Use Figure 11 on page 36, Figure 12 on page 37, and Figure 13 on page 38 for the following steps.

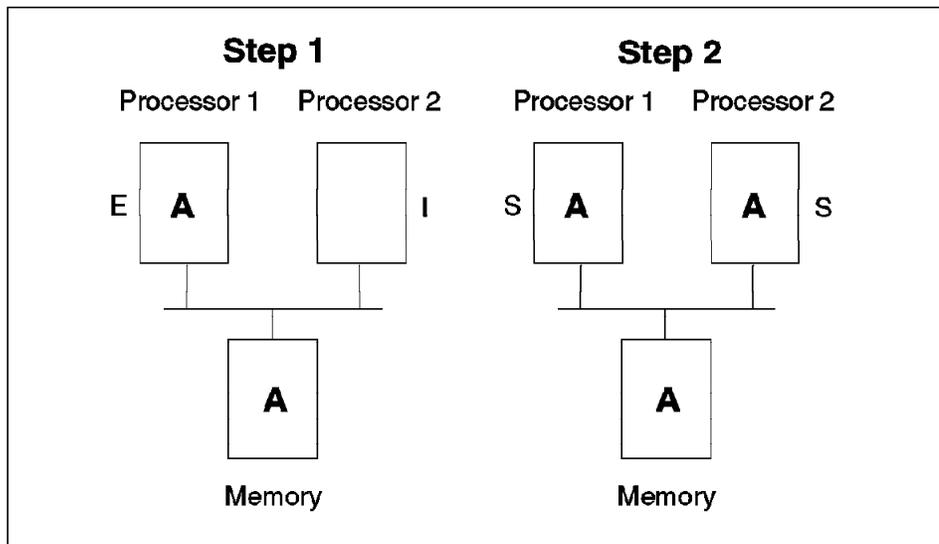


Figure 11. MESI Protocol — Steps 1 and 2

Step 1 — Processor 1 reads value A from memory and loads it in its cache. The corresponding cache sector is marked exclusive for processor 1 and invalid for processor 2.

Step 2 — Processor 2 needs the same memory address. Value A is read from memory, and the corresponding sector is marked shared in both processors.

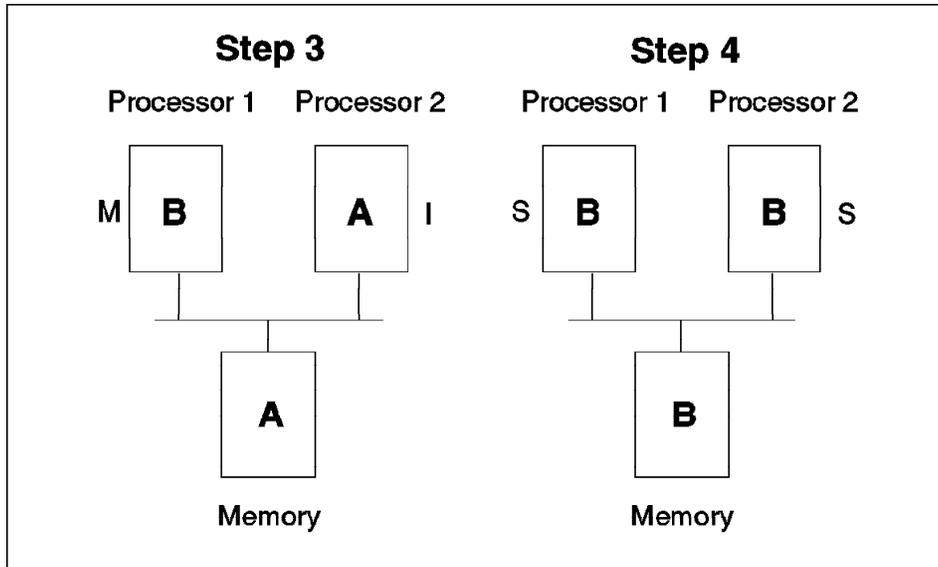


Figure 12. MESI Protocol — Steps 3 and 4

Step 3 — Processor 1 updates the address with a new value B. The snooping logic invalidates the processor 2 sector containing this address. The processor 1 cache sector is marked modified, the processor 2 cache sector is marked invalid, and the memory address is marked invalid.

Step 4 — Processor 2 again reads the same memory address. Since processor 1 holds the modified value, processor 2 loads the value from processor 1. If the operation is a Read With No Intent To Modify (RWNITM), the memory is updated during the transaction. If the operation is Read With Intent to Modify (RWITM), the memory is not updated. It makes sense not to update the memory since the address value is going to be modified.

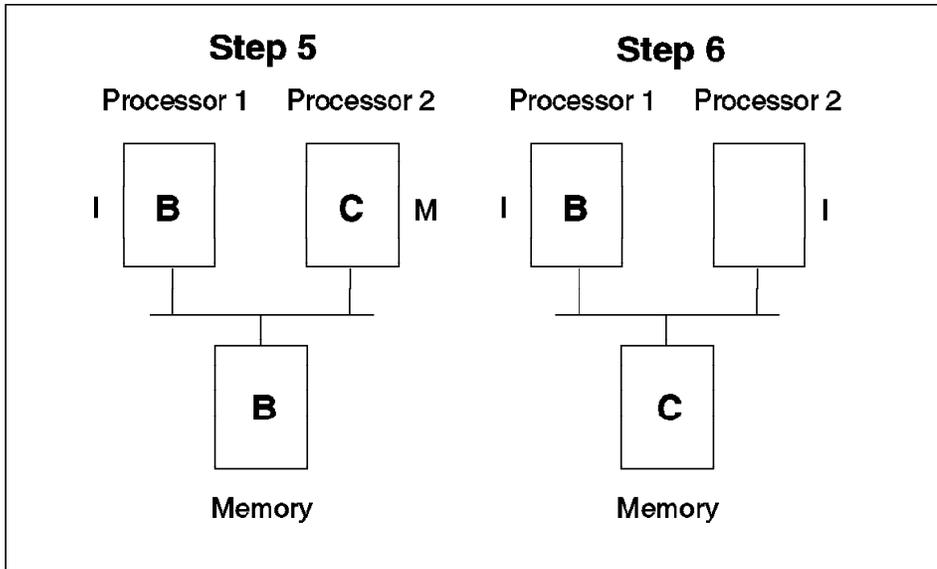


Figure 13. MESI Protocol — Steps 5 and 6

Step 5 — Processor 2 updates the same address. The snooping logic invalidates the cache sector on processor 1, and the memory address is marked invalid.

Step 6 — The Least Recently Used (LRU) sector is the one written back to the memory. In this example, the memory was updated with the value C.

You can see in this memory operation example that snooping in conjunction with the MESI protocol insures cache coherency. This solves one of the main SMP design issues.

L1/L2 Caches and Store Policy

In the IBM SMP implementation, L1 has a 64-byte cache line divided into two 32-byte sectors. L2 cache has a 32-byte cache line.

The store policy defines the way L1, L2, and the memory are kept coherent or up to date.

The store policy between L2 and the memory is write back. This means that when L2 is updated, the memory is not.

In fact, the memory is updated in several cases. It is updated in the case of a cache migration caused by a Read With No Intent To Modify (RWNITM) from another processor. For example, when a processor loads data from another processor and does not want to change that data. Logically, if the

processor wants to change that data, it is not necessary to update the memory since the fetched memory address is going to be changed.

The memory is also updated when the processor needs room in its cache. When this happens, the Least Recently Used (LRU) algorithm is used, and the LRU sector is written back to the memory. The memory is also updated in the case of a cache flush.

False Sharing

Figure 14 illustrates false sharing. Here, a whole cache unit of storage (a cache sector in the case of the IBM SMP system) is represented.

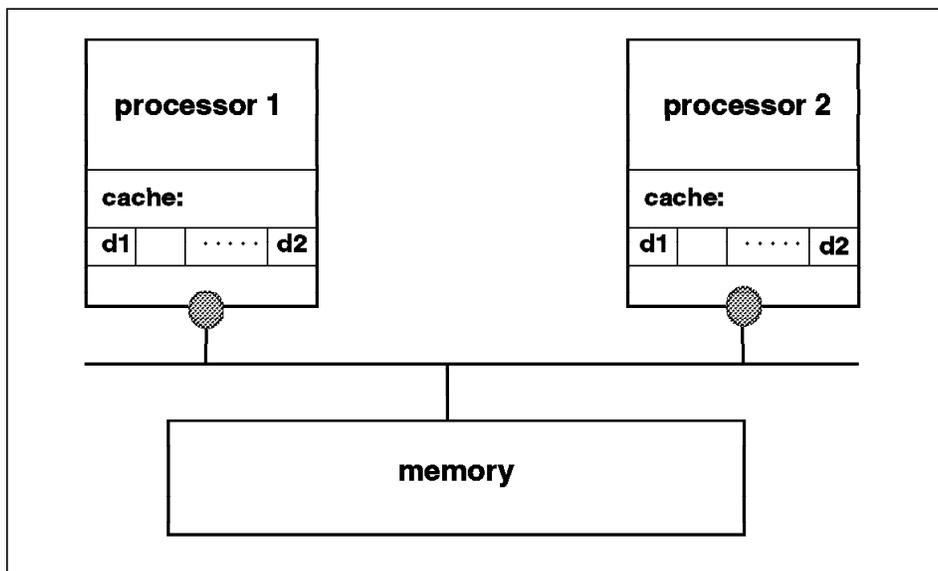


Figure 14. False Sharing

Suppose that 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. The IBM SMP works with a 32-byte cache sector, and the coherency mechanism is based on the cache sector.

Chapter 3. Detailed Product Descriptions

This chapter takes an in-depth look at the hardware packaging and features that comprise the RS/6000 7024-E30, 7025 F40 and F50, and 7026-H50 servers.

Note that similarities between models exist, especially in the 7025 and 7026 products. A complete understanding of the models is obtained when the differences are understood.

3.1 RS/6000 Model E30 Product Description

The Model E30 is a powerful uniprocessor workgroup server that is well-suited for network computing and interoperability. It can be used as either a stand-alone multiuser, departmental, transaction, or database server. The Model E30 has the connectivity to participate in most UNIX and PC networks.

It 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 E30 offers exceptional value for multiuser applications.

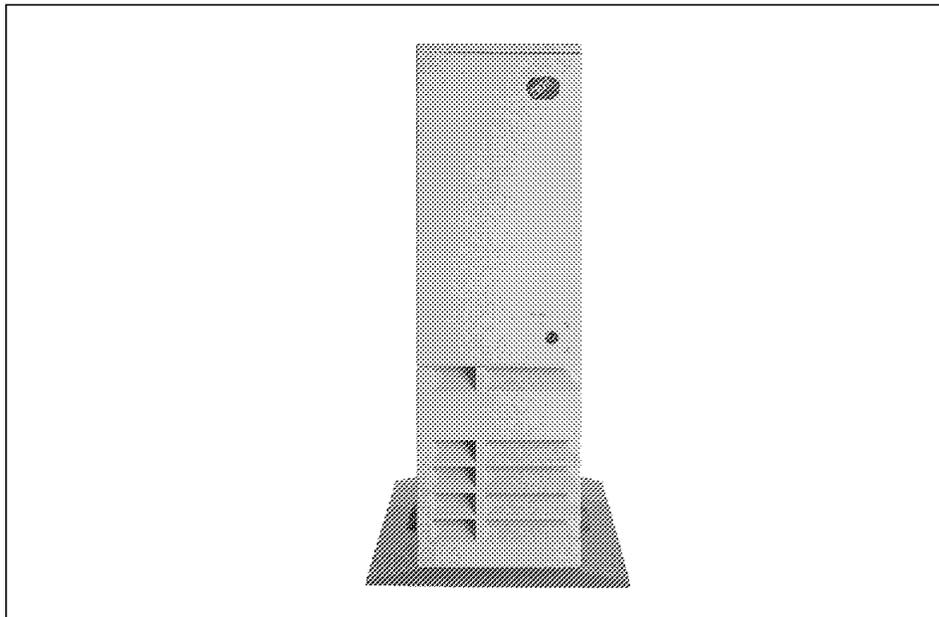


Figure 15. RISC System/6000 Model 7024-E30

The Model E30 offers outstanding performance capabilities due to PowerPC 604e technology running at 233 MHz (a 166 MHz feature is available until July 24, 1998).

The E30 has been designed for customer set up. The modular processor card makes it convenient to replace the microprocessor card as new PowerPC technology emerges.

The E30 is packaged in an 8 x 8 configuration (eight I/O slots, eight-disk/media bays, see Figure 16 on page 43), which allows up to a maximum of 54.6 GB of internal storage capacity and features a flexible I/O subsystem of 5 PCI, 2 ISA, and one shared PCI/ISA slots.

3.1.1 Standard Features

In this section, the standard features of the RS/6000 Model E30 and their most important characteristics are described.

Processor Subsystem

The PCI-based RS/6000 Model E30 server features the PowerPC 604e microprocessor at 233 MHz. The processor includes an L1 cache split into a 16 KB instruction and a 16 KB data cache.

The L2 cache is integrated onto the CPU board and is always present in the base system. The synchronous L2 cache is 512 KB in size.

Memory

The Model E30 comes standard with a 64 MB ECC memory DIMM in one memory slot. The system memory can be expanded up to 1 GB. For this reason, eight memory sockets are provided which accept 16 MB, 32 MB, 64 MB, or 128 MB 168-pin DIMMs.

No Select Feature for Memory

The standard 64 MB memory DIMM cannot be selected for memory upgrades in the initial machine order.

Bays

The system, as shown in Figure 16 on page 43, contains eight 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 20X CD-ROM drive with the convenience of loading a CD without a caddy. The tray-loading mechanism provides support for 12 and 8 cm disks.

The CD-ROM supports standard and XA formats. Its multi-session capability supports Photo-CD applications.

- 5.25-inch disk drive bay

It accommodates the standard 2.1 GB Ultra SCSI disk drive that can be upgraded in the initial machine order to a 2.1 GB Fast/Wide, 4.5 GB 1-inch Fast/Wide, or 4.5 GB Ultra SCSI disk.

The five empty bays can accommodate 3.5-inch or 5.25-inch drives such as disk, CD-ROM, tape drives, or other removable-media drives. Although these bays are 1.6-inch high, some only can be used with 1.0-inch high devices (see 3.1.4, "RS/6000 Model E30 Configuration Notes" on page 51).

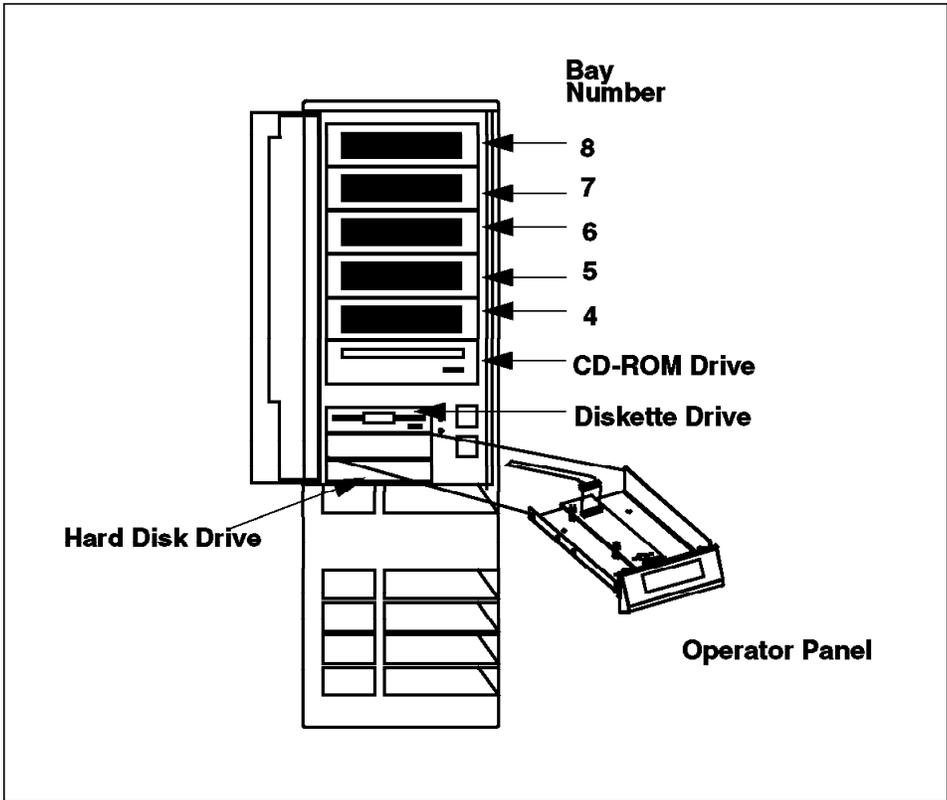


Figure 16. RS/6000 Model E30 Bays and Front View

Adapter Slots

Eight slots are available (free on base machine) which can accommodate the following adapter types:

- Five slots for PCI adapters
- Two slots for ISA adapters
- One shared slot that can accommodate either a PCI or an ISA adapter

The maximum adapter configuration is, therefore, 5 PCI and 3 EISA(ISA), or 6 PCI and 2 EISA(ISA). The CPU does not occupy one of these slots. The I/O motherboard has separate slots for both the CPU card and the optional service processor card.

Integrated SCSI-2 Adapter

The system comes standard with an integrated PCI-based internal SCSI-2 Fast/Wide controller that can be used for internal SCSI-2 devices only. You can have all narrow, all Fast/Wide, or a combination of both SCSI-2 types on the internal SCSI-2 bus.

Standard I/O Ports

The system provides the following standard I/O ports:

- Keyboard port** An optional keyboard, such as the IBM Enhanced Keyboard, can plug into this port
- Mouse port** A port provided to support an optional three-button mouse
- Serial port** Two 9-pin D-shell serial connectors for external RS232 access
- Parallel port** One 25-pin parallel port that maintains complete compatibility with ISA, EISA, and Micro Channel parallel ports.

AIX Operating System

The RS/6000 Model E30 includes a one or two user, full-function AIX Version 4.1.5, Version 4.2.1, or Version 4.3 package at no additional charge.

3.1.2 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 E30.

Memory

Either parity or ECC memory DIMMs can be used and mixed together. This allows the customer to easily upgrade the processor cards without concern over which kind of memory is installed.

Service Processor

To provide greater availability, there is an optional service processor that improves your ability to remotely monitor your system. With this feature, you can remotely restart the system, place it off line, or configure it to make an automatic service call.

Internal Disk Drives

Six internal SCSI-2 disk drives can be accommodated in the internal bays of the system. The disk drives that can be installed and configured are described in Table 7.

Only disk drives with a 1-inch height (half-height) are supported for internal use (bays other than 3-8). Disk drives with 1.6-inch height are not supported because they may result in overheating and data loss.

<i>Table 7. Optional Disk Drives on Model E30</i>				
Formatted Capacity	9.1 GB	4.5 GB	2.2 GB	2.1 GB
Feature Number	# 2908	# 2900	# 3033	# 3095
SCSI Bus Rate MBps	40	40	20	40
Average Seek Time ms	6.7	7.9	8.5	8.5
RPM	7200	7200	7200	5400
Average Latency ms	4.17	4.17	4.17	5.6
Interface	SCSI-3 Fast 20	SCSI-3 Fast 20	SCSI-2 F/W	SCSI-3 Fast 20
Data Width	16 bit	16 bit	16 bit	16 bit
Note: Maximum Internal Disk Drives: 6				

Internal Tape Drives

Two tape drives can be installed in an internal bay of the system, either a 4 GB/8 GB or a 12 GB/24 GB 4 mm internal tape drive.

A maximum of two internal tape drives can be installed. Due to overheating problems that could lead to data errors, they have to be mounted in bays 6 and 8, leaving bay 7 empty. If one tape drive is installed, mount it in bay 8. Table 8 on page 46 lists the optional tape drives available.

<i>Table 8. Optional Tape Drives on Model E30</i>		
Tape Capacity	4.0/8.0 GB	12/24 GB
Feature Number	# 6142	# 6159
Cartridge	IBM 4 mm DDS	IBM 4 mm DDS
Data Compression	Selectable	Selectable
Maximum Effective Transfer Rate	800 KB/sec	2.2 MB/sec
Interface	SCSI-2 (8 bit)	SCSI-2 (8 bit)
Bays	1 half-height	1 half-height
Note: Maximum Internal Tape Drives: 2		

Internal CD-ROM Drives

A maximum of two internal CD-ROM drives can be installed in the E30 system. The CD-ROM provides twenty-speed performance with the convenience of loading CDs without a caddy. Table 9 lists the available CD-ROM drives.

<i>Table 9. CD-ROM Drive on Model E30</i>		
Speed	20X	
Feature Number	# 2619	
Interface	SCSI-2 (8 bit)	
Interface Speed (average)	3000 KB/sec	
Interface Speed (burst)	5.0 MBps	
Bays	1 half-height	
Note: Maximum CD-ROM Drives: 2		

Graphics Adapter

There are two available graphics adapters available for the RS/6000 Model E30. The GXT110P is an 8-bit, 256 color adapter that attaches to one 32-bit PCI slot. The GXT250 contains 3 MB of VRAM (one more than the GXT110P), three hardware color maps, and additional hardware acceleration over the GXT110P. It is suitable for desktop publishing, X Window applications, 2D mechanical drafting, and electrical CAD, to name a few.

Table 10 on page 47 shows the supported monitors.

<i>Table 10. Supported Monitors for Model E30</i>			
Description	Model	Feature Code	Viewable Size (Screen Size)
P70 Color Monitor (Northern hemisphere model available until July 21, 1998, check other models for availability)	6554-603	3613	15.9"
P200 Color Monitor (Northern hemisphere model available until July 21, 1998, check other models for availability)	6555-703	3614	19.1"
P201 Color Monitor (Northern hemisphere model available until July 21, 1998, check other models for availability)	6555-803	3615	19.1"
G52 Monitor	6543-G41	3617 3618 3619	13.6"
P72 Monitor	6556-03N	3620	16.0" (17")
P92 Monitor	6557-03N	3621	17.9" (19")
P202 Monitor	6558-03N	3622	19.8" (21")
Note: Maximum graphics adapters: 1			

PCI SCSI-2 Adapters

The following PCI SCSI-2 adapters are supported:

- PCI SCSI-2 Fast/Wide RAID Adapter (# 2493)
- PCI SCSI-2 Single-ended Ultra SCSI Adapter (# 6206)
- PCI SCSI-2 Differential Ultra SCSI Adapter (# 6207)
- PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)
- PCI SCSI-2 Fast/Wide Differential Adapter (# 6209)

PCI SCSI-2 Fast/Wide RAID Adapter (# 2493): The SCSI-2 Fast/Wide PCI RAID adapter implements RAID level 0, 1, and 5 support for SCSI-2 attached disks. Configuration management, RAID algorithms, and error recovery are handled by an onboard 403 PowerPC. Three internal 16-bit SCSI-2 SE connectors for three buses are provided, one of which is shared with an external connector. The maximum number of addressable device IDs

enabled by the adapter is 15 per bus or 45 devices per adapter. Cabling and housing restrictions may limit the actual total number of supported devices attached to the internal or external connector.

PCI SCSI-2 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-2 protocol. This adapter has a maximum data transfer rate of 40 MBps which is two times the maximum data transfer rate of Fast/Wide adapters. This adapter requires one PCI slot.

Note: Boot support requires a minimum of AIX Version 4.2.1. Earlier releases of AIX will not be able to IPL from devices attached to this adapter.

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 MBps. This adapter allows connection to external differential SCSI-2 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 device. An example of an external SCSI device include the 7137 Disk Array.

PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208): The PCI SCSI-2 Fast/Wide Single-Ended adapter provides a single-ended SCSI-2 Fast/Wide interface that can burst data to devices on the SCSI bus at 20 MBps. This adapter can provide both internal and external device support, connected to the same SCSI bus.

Machines with AIX Version 4.2.0 will require the addition of APAR Ix60857 to operate optimally. This APAR can be obtained by FixDist.

PCI SCSI-2 Fast/Wide Differential Adapter (# 6209): The PCI SCSI-2 Fast/Wide Differential adapter is a PCI adapter that provides a differential Fast/Wide SCSI-2 interface. It can burst data to external devices on the SCSI bus at 20 MBps.

Machines with AIX Version 4.2.0 will require the addition of APAR Ix60857 to operate optimally. This APAR can be obtained by FixDist.

The adapter occupies one PCI bus slot and a maximum of six adapters can be installed in the Model E30.

Table 11 on page 49 shows the maximum number of both internal and external devices that can be attached to a SCSI-2 adapter together with related cables. By *multi-device*, we mean an external enclosure that may contain more than one SCSI device.

<i>Table 11. SCSI Adapter Connections on Model E30</i>					
Adapter	Maximum number of devices			Cables Feature Code	
Feature Code	Internal	External		Internal	External
		Single	MultiDevice		
Integrated	7	0	0	Not required	Not allowed
2408/6208 F/W Single Ended	4	4	0	2442	2111, 2113, 2115, 2117
2408/6208	0	4	0	Not allowed	2111, 2113, 2115, 2117
2408/6208	0	0	1	Not allowed	2111, 2113, 2115, 2117
2409/6209 F/W Differential	0	25 meter cable maximum 15 SCSI addresses maximum		Not allowed	2112, 2114, 2116
Note: Maximum SCSI adapters: 6					

Communication Adapters

The following communication adapters are supported:

- Auto LANStreamer Token-Ring PCI Adapter (# 2979)
- PCI Ethernet Adapters
 - PCI Ethernet BNC/RJ-45 Adapter (# 2985)
 - PCI Ethernet AUI/RJ-45 Adapter (# 2987)
 - 10/100 Mbps Ethernet (# 2968)
 - Fast Etherlink XL (# 2986)
- ATM Adapters
 - TURBOWAYS 155 PCI UTP ATM Adapter (# 2963)
 - TURBOWAYS 155 PCI MMF ATM Adapter (# 2988)
 - TURBOWAYS 25 PCI ATM Adapter (# 2998)
- ISA X.25 Adapter (# 2961)
- ISA 4-Port Communications Controller (# 2701)
- ISA 2-Port Multiprotocol PCI Adapter (# 2962)

- FDDI Adapters
 - SysKonnnect SK-NET FDDI-LP SAS (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS (# 2743)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)

3.1.3 External Subsystems

The following is a list of the external subsystems which are officially supported by the RS/6000 Model E30.

Disk Subsystems

- 7131 Model 105 SCSI Multi-Storage Tower
- 7137 Models 412, 413, and 414 RAID Disk Array Subsystem
- 7203 Model 001 External Portable Disk

Tape Drive Subsystem

- 7207 Model 012 — 1.2 GB 1/4 inch Tape drive
- 7206 Model 005 — 4 GB 4 mm Tape Drive
- 7208 Model 011 — 5 GB 8 mm Tape Drive

Tape Libraries

- 7332 Model 005 — 4 mm Library
- 7331 Model 205 — 8 mm Library

LAN Subsystems

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

Optical Drives/Libraries

- 7210 Model 010 External CD-ROM Drive
- 3995 Models A63, 063, 163 Optical Library Dataservers

ASCII Terminals

- IBM 3151, 3153, 3161, 3162, 3163, 3164
- DEC VT100, VT220, VT320, VT330
- WYSE 30, 50, 60, 350

Modems

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

Printer/Plotters

Any AIX 4.1.5 or above supported serial, parallel, or LAN-attached printer/plotter.

3.1.4 RS/6000 Model E30 Configuration Notes

When planning or implementing an RS/6000 Model E30, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

SCSI Adapters

The limitations regarding SCSI adapters and the number of attachable devices are shown in Table 11 on page 49.

Internal Disk Drives

Internal disk drives with a 1.6-inch height are supported in bays 3-8. Disk drives with 1.6-inch height (4.5 GB disks) are not supported in other empty bays. Installation of any full-height internal disk in bays other than 3-8 may result in overheating the system, and data consistency is not guaranteed.

Internal Tape Drives

The 4 GB/8 GB 4 mm internal tape drive should be mounted in bay number 8 only. This insures adequate air flow, as required for proper operation. A second tape drive may be installed in bay number 6 if bay number 7 remains empty. Mounting in another bay may cause overheating and could lead to data errors.

An 8 mm internal tape drive cannot be installed in the RS/6000 Model E30.

Graphics Adapter

There can be a maximum of two graphics adapters installed on the system.

Service Processor Limitations

PCI slots 4 and 5 must contain *short* adapters if the service processor is installed. These adapters cannot exceed 7.5 inches (190 mm).

Memory Limitations

The 128 MB DIMM cannot be installed in slot 1 due to physical limitations.

3.2 RS/6000 Model F40 Product Description

The RS/6000 7025-F40, a desk side 1- or 2-way symmetric multiprocessor (SMP), is one of IBM's most expandable PCI workgroup servers. It comes standard with a single 233 MHz PowerPC 604e. Converting from a 1-way to a 2-way system is as simple as plugging in an additional processor card. Nine available I/O expansion slots and up to 22 disk and media bays (10 standard, 12 optional) provides outstanding expansion room. Eighteen of the disk drive bays are hot-swappable. The F40 supports an optional service processor for improved availability and serviceability.

Figure 17 shows an RS/6000 Model F40.



Figure 17. RS/6000 Model F40

3.2.1 Model F40 Highlights

The most remarkable features of the RS/6000 Model F40 are described in this section.

Exceptional Expandability

With expansion capability of up to 1 GB of memory, 172.8 GB of internal disk capacity, and nine expansion slots, the Model F40 can meet the needs of your growing business. In a small package slightly more than two feet (620 mm) high, this powerful workgroup server can easily fit into your department or work area.

Outstanding Reliability

The Model F40 provides the reliability and serviceability features found in higher-priced systems, including:

- Error Checking and Correcting (ECC) memory, which corrects single-bit errors and detects two-bit errors.
- Optional hot-swappable disks, which can help improve system availability and serviceability by allowing disks to be added or changed without shutting down the system.
- Optional service processor, which can help improve system availability by providing better problem determination as well as remote system support and management capability.
- Internal RAID-5 support, which provides disk redundancy for internal (and external) devices.

Powerful Graphics Capabilities

The Model F40 offers comprehensive graphics support from entry 8-bit 2D to 24-bit true-color 2D through high-performance 3D graphics accelerators. In addition, IBM's multithreaded graphics libraries exploit the SMP architecture to provide excellent performance for applications such as MCAD. The F40 also has an integrated tablet port that satisfies many graphics application requirements, such as CAD, without taking up a slot.

Easy Processor Upgrade

Converting from a 1-way to a 2-way SMP system is as simple as plugging in a second PowerPC 604e 233 MHz processor card.

Industry Common I/O and Storage

The Model F40 incorporates popular and economical PCI, ISA, SCSI-2, and Ultra SCSI technologies that support many widely-used I/O and storage peripherals. The dual-PCI buses (32- and 64-bit) help provide higher I/O throughput.

High Performance Storage Options

Ultra SCSI and Serial Storage Architecture (SSA) storage options offer improved data transfer rates over SCSI-2.

Internal RAID-5

Cost-effective RAID-5 is available for internal attachment thereby reducing the need for expensive RAID subsystems.

Investment Protection

An upgrade path to the next larger SMP workgroup server, the Model F50, is available.

3.2.2 Standard Features

In this section, the standard features of the RS/6000 Model F40 are described, and their most important characteristics are pointed out.

Processor Subsystem

The base RS/6000 Model F40 server configuration features one PowerPC 604e microprocessor at 233 MHz. The F40 is upgradable to a 2-way symmetric multiprocessing (SMP) system. The processor includes an L1 cache composed of a 16 KB instruction and a 16 KB data cache. Each processor has a 1 MB L2 cache.

Memory

The Model F40 comes standard with two 32 MB ECC memory EDO DRAM DIMMs with a cycle time of 60 ns. They use two memory slots. The system memory can be expanded up to 1 GB. Eight memory sockets are provided which accept 16 MB, 32 MB, 64 MB, or 128 MB DRAM DIMMs. The total quantity of each size of memory card must be an even number.

Bays

The standard system contains four media bays and one disk drive carrier, known as a *6-Pack*, which accommodates six disk drive bays. A 6-Pack is a tray providing an interface between the standard disk drive and a backplane connector. The 6-Pack can be either SCSI, SCSI hot swap, or SSA. Figure 18 on page 55 shows the structure of the system.

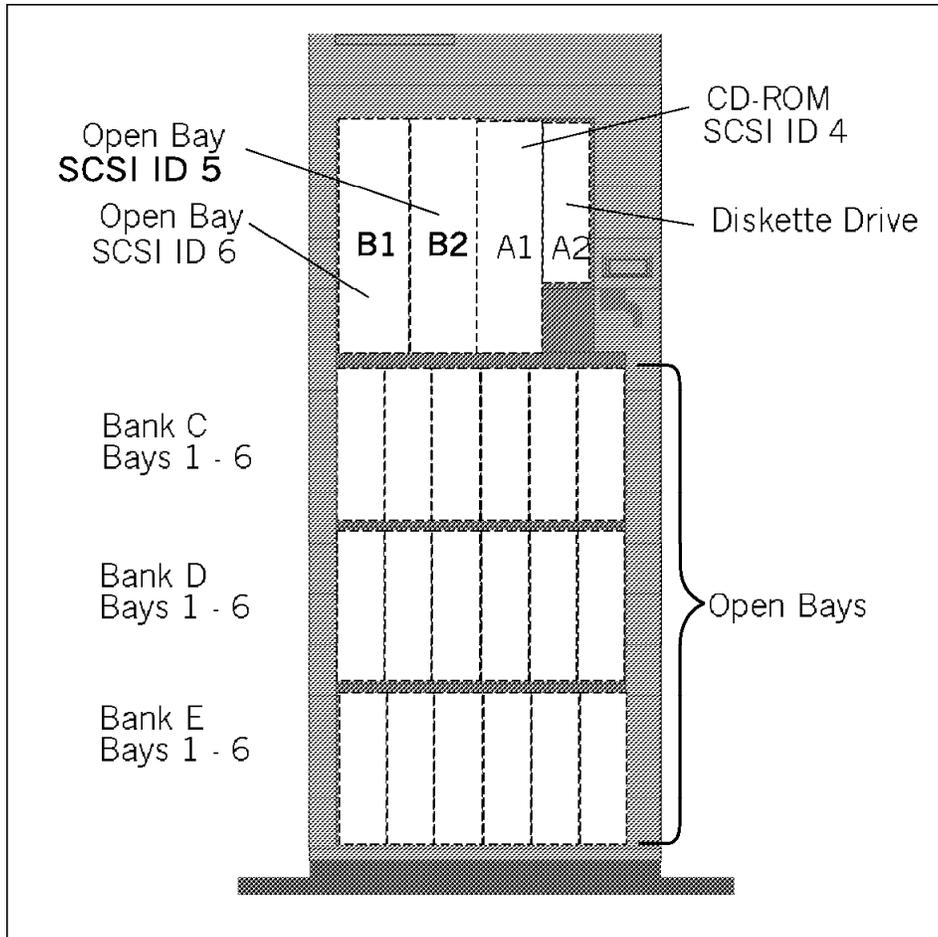


Figure 18. RS/6000 Model F40 Bays

The standard bays are composed of four media and six SCSI disk bays which accommodate the following drives:

- Two of the media bays accommodate the following standard devices:
 - 3.5-inch diskette drive bay that accommodates a standard 1.44 MB diskette drive.
 - 5.25-inch CD-ROM drive bay that accommodates a 20X SCSI-2 CD-ROM drive. The tray-loading mechanism provides support for 12 cm and 8 cm disks.
- The additional two media bays are empty on standard configuration. They can accommodate CD-ROM drives, tape drives or 4.5 GB Ultra SCSI disks.

- Depending on the base 6-Pack (SCSI or SSA), the base disk is one of the following:

SCSI One hot swap disk bay accommodates the standard 2.2 GB SCSI-2 F/W disk drive. This disk can be upgraded to a 4.5 GB or 9.1 GB Ultra SCSI disk drive in the initial order. The 2.2 GB SCSI F/W and 9.1 GB 1-inch Ultra SCSI disks (1 inch, half-height) require one hot swap bay; the 4.5 GB and 9.1 GB Ultra SCSI disk drives (1.6 inch, full-height) occupy two.

SSA The boot disk must be SCSI; so the 4.5 GB Ultra SCSI Media Bay disk must be ordered. Two additional SSA disks (4.5 GB or 9.1 GB) are required when an SSA 6-Pack is used. Because of the disk's height, either six 4.5 GB disks or three 9.1 GB disks fit into one 6-Pack.

- The additional available five hot swap disk bays are empty in standard system configurations.
- For economical reasons, it is possible to keep the non hot swap 6-Pack. In this case, only 4.5 GB disks can be ordered in conjunction with a media bay disk drive mounting kit for each disk.

Integrated SCSI-2 Fast/Wide Adapter

The system comes standard with two integrated PCI-based SCSI-2 Fast/Wide controllers. One of them can be used only for internal, single-ended (Fast/Wide, narrow, or both) SCSI-2 devices (maximum of seven). The other one is externalized, onto which up to six internal and two external devices can be connected.

Adapter Slots

Nine slots are available which can accommodate the following adapters types:

- Five slots for PCI adapters (32-bit)
- Two slots for PCI adapters (64-bit)
- Two slots that can accommodate either a PCI or an ISA adapter

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 for external RS232 access.

Parallel port One 25-pin parallel port that maintains complete compatibility with machines containing ISA, EISA, and Micro Channel parallel ports.

Communication Port

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

Audio Adapter

The Ultimedia Integrated Audio Adapter is capable of simultaneous playback and capture of audio, with sample rates ranging from 5513 to 48000 samples per second. The data can be mono or stereo. Pulse code modulation (PCM) with either mu-law (14-bit input samples mapped to 8-bit values) or a-law (13-bit input samples mapped to 8-bit values) encoding is supported. Microphone, headphone, line-in, and line-out ports are provided.

Figure 19 shows the rear view of the system.

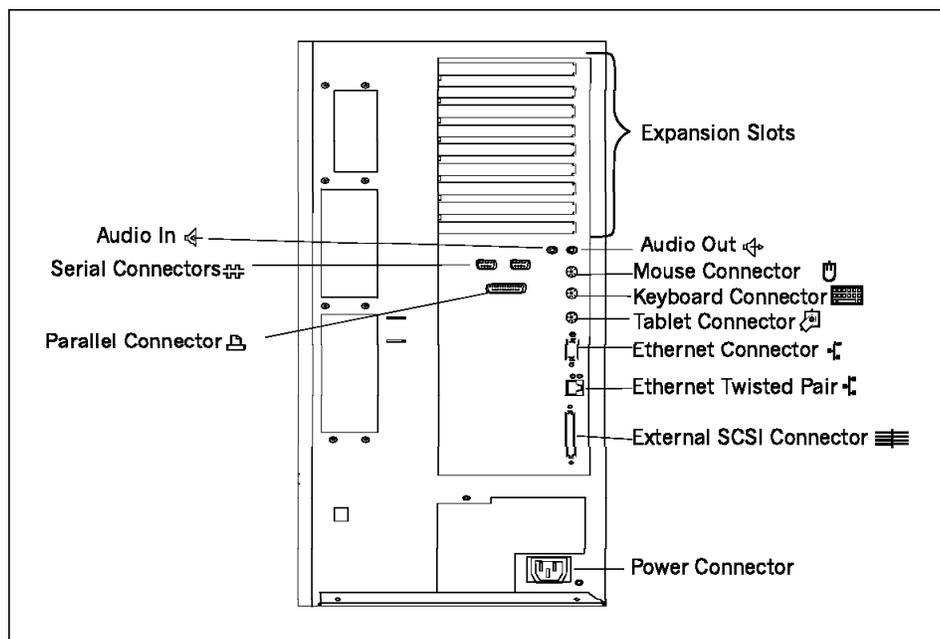


Figure 19. RS/6000 Model F40 Rear View

AIX Operating System

The Model F40 includes a one or two user, full-function AIX Version 4.1.5, Version 4.2.1 or Version 4.3 package at no additional charge.

3.2.3 Supported Optional Features

In this section, *optional features* are defined only as internal optional devices. Adapters are described in more detail in Chapter 5, "Adapters, Overview and Configuration" on page 147, and storage subsystems in Chapter 4, "Internal Storage Features" on page 91.

Additional Processor

The F40 is upgradable to a 2-way SMP system by adding a second 233 MHz PowerPC 604e.

Memory

Memory DIMMs must be installed in pairs. The base system memory configuration is 64 MB, consisting of two 32 MB EDO ECC DIMMs. A select feature can be ordered in pairs to replace this base memory with 32 MB, 64 MB, or 128 MB EDO ECC DIMMs. Additional DIMMs in capacities of 16 MB, 32 MB, or 128 MB may be ordered in pairs up to a total of six.

Service Processor

To provide greater availability, there is an optional service processor that improves your ability to remotely monitor your system. With this feature you can remotely restart the system, place it off line, or have it make an automatic service call. For more information, refer to Section 6.7, "Service Processor" on page 226.

Optional Disk Bays

As described in 3.2.2, "Standard Features" on page 54, the RS/6000 Model F40 features four media bays, six standard, and twelve optional half-height standard SCSI or hot-swappable disk bays. To add six half-height hot-swappable disk bays to the system, additional hardware is required. It can be ordered as a hardware feature named SSA 6-Pack, SCSI Hot Swap 6-Pack, or SCSI 6-Pack, depending on which optional bays are going to be installed. In addition, a power supply and one connection cable for each 6-Pack must be ordered.

Internal Disk Drives

Additional disk drives can be mounted in either SSA or SCSI 6-Packs. SCSI and SSA 6-Packs can be mixed, but the first 6-Pack must be SCSI or the select feature # 6527 must be added to replace the base SCSI backplane with an optional SSA backplane.

Internal SCSI or SSA disk drives in varied capacities in the industry standard 3.5-inch form factor can be installed on the system. Available internal disk drives have 2.2 GB, 4.5 GB, and 9.1 GB capacities.

Both 1-inch height (half-height) and 1.6-inch (full-height) disk drives are supported, but while half-height disks require only one hot-swappable bay, the full-height ones occupy two. For this reason, the maximum number of full-height disk drives that can be installed is nine (two bays for each drive), and the maximum number of half-height disk drives is 18 (one bay for each drive).

For more information, refer to 4.2.3, “Internal SSA Devices” on page 118, and to “Disk Drives” on page 98.

Internal Tape Drives

All four types (4 GB/8 GB 4 mm, 12 GB/24 GB 4 mm, 5 GB/10 GB 8 mm, or 20 GB/40 GB 8 mm) of internal tape drives can be installed in an internal non-hot-swappable media bay.

Refer to Table 17 on page 100 for more information.

Internal CD-ROM Drives

The CD-ROM provides 20X performance with the convenience of loading CDs without a caddy. Refer to “CD-ROM Drives” on page 99 for details.

SCSI Adapters

The following five PCI SCSI adapters are supported:

- Ultra SCSI Adapter (# 6206)
- SCSI-2 Fast/Wide PCI-Bus Adapter (# 6208)
- Ultra SCSI Differential Adapter (# 6207)
- SCSI-2 Differential Fast/Wide Adapter (# 6209)
- SCSI-2 Fast/Wide RAID Adapter (# 2493)

Table 12 on page 60 shows the maximum number of both internal and external devices that can be attached to a SCSI-2 adapter and related cables for the RS/6000 Model F40.

Table 12. SCSI Adapter Connections on Model E30

Controller	Maximum number of devices			Cables Feature Code	
Feature Code	Internal	External		Internal	External
		Single	MultiDevice		
Integrated interface 1	3 Media (1 Standard) 6 Disk (1 Standard)	No External Support	No External Support	Standard	Not allowed
Integrated interface 2 (External multi-devices)	6	2	0	2443 or 2444	2111, 2113, 2115, 2117
Integrated interface 2 (External multi-devices)	0 (Cannot mix single and multi devices)	0 (Cannot mix single and multi devices)	1	Not allowed	2111, 2113, 2115, 2117
6208 Without external multi-devices	6	4	0 (Cannot mix single and multi devices)	2447, 6552, 6536	2111, 2113, 2115, 2117
6208 With external multi-device	0	0 (Cannot mix single and multi-device)	1	Not allowed	2111, 2113, 2115, 2117
6209	0	25-meter cable maximum 15 SCSI addresses maximum		Not allowed	2111, 2113, 2115, 2117

SCSI-2: The Model F40 will typically be booted from a hard disk which requires a SCSI-2 disk and controller. The base system configuration provides a single, non-hot-swappable SCSI 6-Pack connected to the first integrated SCSI controller using a base SCSI cable. This cable also connects the SCSI media devices. A single 2.2 GB SCSI-2 F/W disk (# 9398) is standard. An optional hot-swappable SCSI 6-Pack backplane (# 6520) can be selected if desired.

An option to disconnect the first SCSI 6-Pack from the media devices is available. This option, using feature # 2443, provides a cable and a terminator which connect the second integrated SCSI controller to the first 6-Pack and terminates the base cable after the media devices. If this option is implemented, up to two additional SCSI 6-Packs (# 6537 or # 6523) can be added. Each additional 6-Pack must be driven by an additional PCI SCSI adapter (# 6208) and connected using cable # 2444.

Note

An auxiliary power supply (# 6549) must be added to the F40 to support the second and third 6-Packs, if installed.

If there is no requirement to disconnect the first SCSI 6-Pack from the media devices, an optional second SCSI 6-Pack (# 6537 or # 6523) may be ordered and attached to the second integrated SCSI adapter using cable # 2444. An optional third SCSI 6-Pack may then be driven by PCI SCSI adapter (# 6208) connected using cable (# 2444).

In certain configurations, it may be desirable to locate the F40's boot disk in a media bay. Feature # 3005 is a select option that replaces the base 2.2 GB SCSI-2 F/W hard disk (# 9398) with a 4.5 GB Ultra SCSI enabled hard disk mounted in a media bay. Feature # 6509, media bay mounting hardware, allows a second 4.5 GB Ultra SCSI-enabled hard disk (# 2900) in the remaining media bay to provide boot disk mirroring, if desired, by using the RAID mirroring function contained within AIX. Hard disks located in the media bays are not hot-swappable. When the F40's two available media bays are used in this manner, any other optional devices, such as a tape back-up unit, must be attached externally to the system.

SCSI RAID: The IBM PCI RAID Adapter (# 2493) can be utilized to control the SCSI 6-Packs and enable SCSI RAID 0, 1, and 5 under the covers of the F40. The RAID adapter can control up to three SCSI 6-Packs. Each 6-Pack requires cable feature # 2447 for connection to the PCI SCSI RAID adapter. If SCSI RAID is implemented on the first 6-Pack, the SCSI-2 boot disk must be moved to a media bay using feature # 3005.

Internal RAID Indicator (# 9995) is utilized to indicate that attachment of the backplanes to the RAID adapter is desired. If feature # 9995 is not present, it is assumed external devices only are to be attached to the RAID adapter.

Ultra SCSI: The IBM PCI Ultra SCSI Adapter (# 6206) can be utilized to control the SCSI 6-Packs and provide Ultra SCSI speed when used with the 4.5 GB and 9.1 GB Ultra SCSI-enabled hard disks. Each Ultra SCSI-enabled 6-Pack requires a dedicated PCI Ultra SCSI adapter. Ultra SCSI is not supported on the Model F40's integrated SCSI-2 F/W adapters. If Ultra SCSI is implemented on the first 6-Pack, it must be disconnected from the media bays. Feature # 2446 provides a cable for connection of the first SCSI 6-Pack to the Ultra SCSI adapter and a terminator for the cable from the media devices when disconnected from the first SCSI 6-Pack. Cable feature # 2447 is used to connect the second and third 6-Packs to their dedicated Ultra SCSI adapters.

Note

Ultra SCSI enabled disks perform at SCSI-2 speed if the SCSI 6-Packs are connected to the integrated or non-integrated SCSI Fast/Wide adapters, rather than the PCI Ultra SCSI adapters.

Ultra SCSI RAID: If an Ultra SCSI RAID configuration is desired, the PCI Ultra SCSI Differential adapter (# 6207) can be used to support external Ultra SCSI RAID subsystems. Ultra SCSI RAID is supported in an under-the-covers configuration using the striping (RAID 0) and mirroring function (RAID 1) of AIX and using the PCI Ultra SCSI SE adapter (# 6206).

Serial Storage Architecture (SSA) and SSA RAID

The Model F40 supports both external and under-the-covers SSA disk subsystems. The SSA Multi-Initiator/RAID EL Adapter (# 6215) can be used to control up to three 6-Packs of SSA disks. When implementing SSA internally, SSA-enabled hard disks (# 3071 and # 3072) and 6-Pack backplanes (# 6538 and # 6539) are required. SSA disks and backplanes are hot swap capable.

If the F40 is booted from hard disk, SCSI attachment is required. A boot disk in the first 6-Pack may be used for this purpose, leaving the second and third 6-Packs available for internally attached SSA drives. A maximum of one SCSI 6-Pack is allowed in configurations including both SSA and SCSI 6-Pack backplanes.

When SSA is desired on the first 6-Pack, feature # 3005 must be ordered to move the system boot disk to a media bay, and the first SCSI 6-Pack is replaced with an SSA backplane using select feature # 6538 (or # 6527 for initial orders only). Optional SSA backplane feature # 6539 is used to implement SSA on the second and third 6-Packs.

Note

An auxiliary power supply (# 6549) must be added to the Model F40 when the second and third 6-Packs are installed.

To optimize performance when more than one SSA adapter (# 6215) is required, it is recommended that each adapter be assigned to a slot supported by an individual PCI bus. For example, the first SSA adapter could use slots 1 or 2; a second PCI SSA adapter could use slots 3, and a third PCI SSA adapter could use slots 4 through 9.

Two SSA cable options are available. The first cable option (# 2450) provides internal cables to internal ports on the PCI SSA adapter to allow connection of up to three SSA 6-Packs (up to fifteen SSA 4.5 GB drives), all under-the-covers of the model F40.

The second cable option (# 2453) provides for the connection of up to three SSA 6-Packs under the covers and exits the rear cover through a bulkhead connection to the external ports of the PCI SSA adapter (# 6218). This cable option will allow up to 30 additional external SSA devices, for a maximum of 48 total devices, to be connected to the SSA loop at the point where it exits the rear cover.

A minimum quantity of two SSA hard disks is required. The first two drives must be located in bays one and five of each 6-Pack ordered. A full 6-Pack can be configured as well. Should the desired capacity fall short of fully populating a 6-Pack, an SSA blank disk drive module is required for each SSA backplane connector that is not occupied with an SSA disk. These blank disk modules will be included for initial plant orders. The SSA Blank Disk Drive Module feature (# 2452) must be ordered separately for MES installation of SSA disk subsystems. The SSA blank disk drive module (# 2452) must be installed first if it is being used in conjunction with a 1.6-inch 9.1 GB SSA hard disk (# 3072).

Graphics Adapters

The Model F40 has comprehensive AIX graphics support, enabling it to run 2D business graphics through complex high-function 3D graphics applications. Coupled with its expansion potential, the F40 can serve a high-function workstation for environments that handle large volumes of 2D and 3D graphics data.

For 2D graphics applications, the Model F40 offers a number of graphics accelerators. The POWER GXT120P 2D accelerator meets the graphics needs of many commercial applications and provides an entry point for workstation users. It is also a cost-effective choice for server environments that require a graphical user console. For high levels of 2D graphics function and performance, you can choose the 8-bit POWER GXT250P or the 24-bit true-color GXT255P graphics accelerator.

For advanced 3D graphics function, the Model F40 offers several options that vary in their speed of processing graphics data and cost. The POWER GXT500P and GXT550P graphics accelerators provide outstanding 3D accelerated function at an affordable price and are well suited for wireframe and entry solid-modeling applications. The higher performing POWER GXT800P graphics accelerator has optional hardware texture mapping, which is ideal for solid modeling and visualization. These graphics

accelerators utilize IBM's multithreaded PHIGS and OpenGL libraries, which exploit the F40's SMP architecture to provide excellent performance.

For complex, high-function 3D graphics applications, the POWER GXT1000 graphics accelerator can be attached either internally or externally (the GXT1000 will no longer be offered after July 24, 1998). Attached internally, this accelerator provides more effective 3D graphics based on price/performance. The maximum number of disk bays is reduced to six, due to the space requirements of the internal GXT1000 adapter.

For more information regarding the RS/6000 graphics offering, refer to 5.2, "Graphics Adapters" on page 158.

Multimedia Adapters

The optional Ultimedia Video Capture Adapter (# 2638) supports video monitoring from standard PAL, NTSC, or SECAM sources, such as video cameras, VCRs, and laser disc players. The adapter also supports full-screen monitoring of incoming video (with a graphics adapter with scaling support like the GXT120 adapter) and allows capture of still images and raw video to system memory or disk. Device and user interface support is provided through Ultimedia Services software.

Communication Adapters

The following communication adapters are supported:

- Auto LANStreamer Token-Ring PCI Adapter (# 2979, until July 24, 1998)
- IBM PCI Token-Ring Adapter (# 2920)
- PCI Ethernet Adapters
 - PCI Ethernet BNC/RJ-45 Adapter (# 2985)
 - PCI Ethernet AUI/RJ-45 Adapter (# 2987)
 - 10/100 Mbps Ethernet (# 2968)
 - Fast Etherlink XL (# 2986)
- ATM Adapters
 - TURBOWAYS 155 PCI UTP ATM Adapter (# 2963)
 - TURBOWAYS 155 PCI MMF ATM Adapter (# 2988)
 - TURBOWAYS 25 PCI ATM Adapter (# 2998)
- ISA X.25 Adapter (# 2961)
- ISA 4-Port Communications Controller (# 2701)
- ISA 2-Port Multiprotocol PCI Adapter (# 2962)

- FDDI Adapters
 - SysKonnnect SK-NET FDDI-LP SAS (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS (# 2743)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)

More detailed technical information about these PCI and ISA adapters can be found in 5.1, “Communication Adapters” on page 147.

Model Conversion to a Model F50

Model F40 1- or 2-way SMP can be converted to a Model F50 1-, 2-, 3-, or 4-way SMP. For a detailed description about the Model F50, refer to 3.3, “RS/6000 Model F50 Product Description” on page 66.

The upgrade to the Model F50 contains the Model F50 base configuration, less the boot hard disk and 20X CD-ROM, with the customer’s existing serial number:

- A 1-way 166 MHz PowerPC 604e processor card with 32 KB instruction and 32 KB data Level 1 cache.
- 256 KB ECC Level 2 cache per processor
- 128 MB ECC Synchronous DRAM
- Service processor
- 1.44 MB, 3.5-inch diskette drive
- Five 32-bit PCI slots, two 32/64-bit PCI slots, and two shared PCI/ISA slots
- Dual integrated SCSI-2 Fast/Wide controllers
- One SCSI hard disk backplane providing for up to six available hard disk bays. One media bay filled by base diskette drive and two available 5.25-inch x 1.6-inch bays for media or boot disk mirroring. (One will be filled with the CD-ROM transfer from F40)
- Integrated Ethernet with thick (10Base5) and twisted pair (10BaseT) output connectors
- Keyboard and mouse ports
- Three serial ports and one parallel port
- 930-watt base power supply
- New labels
- Installation instructions, new diagnostics, new publications

The existing F40 disk, media devices, and PCI/ISA adapters may be transferred to the Model F50 system since most are supported.

Note

The F40 processor cards, memory, SCSI backplanes, service processor, power supplies, media device feature codes # 6137 and # 2630, most high-end graphics adapters and the video capture adapter, microphone, and speakers are not supported on the F50 system.

The following features may be transferred to the F50:

- All hard disks
- 20X CD-ROM
- Most PCI/ISA communications and storage adapters
- Graphics adapter limited to GXT110P, GXT120P, and GXT800P

Due to the F50's requirements, AIX Version 4.2.1 or AIX 4.3 must be ordered if not already available.

3.2.4 RS/6000 Model F40 Configuration Notes

When planning or implementing an RS/6000 Model F40, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

- 9.1 GB Disk drives (1.6 inch height) require two hot-swappable bays.
- There can be a maximum of two graphics adapters installed in the system.
- A SCSI boot disk must be configured in the media bay if all three 6-Packs use SCSI RAID or SSA disks.

3.3 RS/6000 Model F50 Product Description

The RS/6000 Model F50 server is the first of a new generation of 1- to 4-way symmetric multiprocessing (SMP) workgroup and enterprise servers that provide for unprecedented commercial price/performance. Utilizing the advanced PowerPC 604e processor and X5 Level 2 (L2) cache system, the Model F50 is well suited for customers who want commercial solutions to address mission-critical online transaction processing as well as collaborative computing comprised of Lotus Notes, Domino Server, Internet, intranet, extranet, and groupware application solutions.

The Model F50 can be used as a stand-alone, multiuser application server and has the connectivity to participate in most currently installed UNIX and PC networks. It can also be configured as a database or transaction server with your choice of SCSI-2, Ultra SCSI, or Serial Storage Architecture (SSA) RAID disk by taking advantage of the storage expansion capability of the system. It is designed to allow for system growth in memory, disk, media devices, and 1- to 4-way processor scalability as needs grow over time.

Figure 20 shows an RS/6000 Model F50.

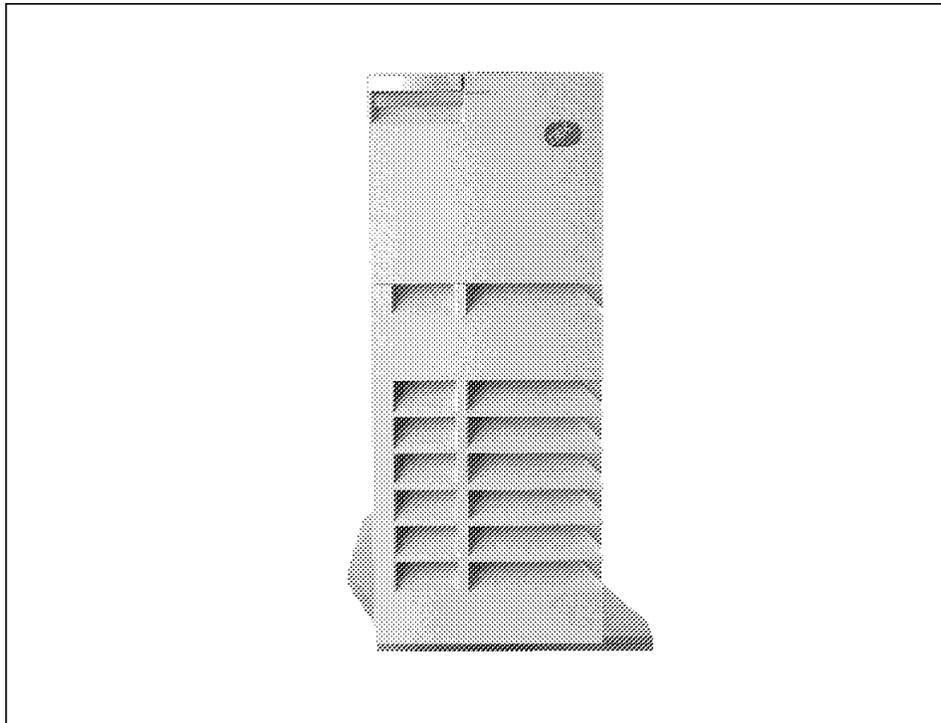


Figure 20. RS/6000 Model F50

3.3.1 Model F50 Highlights

The most remarkable features of the RS/6000 Model F50 are described in the following section.

- 1- to 4-way SMP desktop server with superior commercial price/performance.
- Offers affordable upgrades, including additional processor speeds and memory.

- Cost-effective RAID-5 is available for internal attachment, lessening the need for expensive RAID systems.
- First RS/6000 server to use Synchronous DRAM (SDRAM) technology to enhance system performance.
- Can be configured with SCSI-2, Ultra SCSI, or SSA RAID disk.
- Outstanding reliability through ECC memory, ECC L2 cache, plus a standard service processor.
- Hot-swappable disk bays give improved system availability, smooth growth, and data portability.
- Designed for ease of scalability and expandability including the ability to add media devices, system disk and backplanes, system processor cards, dual inline memory modules (DIMMs) and memory cards, as well as I/O adapters without changing the standard power supply.
- Adheres to open industry standards as well as industry standards in general. The I/O subsystem is PCI/ISA-based and also supports the industry interconnect standards of SCSI, Ultra SCSI, and SSA.
- Preconfigured, tailorable solutions for ease of ordering, installation, and set up.

3.3.2 Model F50 System Architecture

The F50 design incorporates high-speed processor and memory buses, a more efficient 8-way associative Level 2 (L2) cache design, a high-performance memory-I/O controller, and fast Synchronous Dynamic Random Access Memory (SDRAM). The memory I/O controller has a separately clocked, 64-bit mezzanine I/O bridge bus to which three independent PCI buses are connected. This unique design partitions system logic into a high-speed, processor-memory portion and a lower-speed industry standard I/O portion. This allows for greater throughput on the SMP system bus and lower manufacturing costs because standard PCI and ISA devices can be utilized.

Processors and L2 Cache

The F50 contains either 166 MHz PowerPC 604e or 332 MHz PowerPC 604e processors, each with its own 256 KB L2 cache. The X5 cache controller chip serves as a bus converter between the 64-bit 166 MHz processor bus and the 128-bit 83 MHz SMP system bus. The processor and system buses have a peak transfer capacity of 1330 MB per second.

The cache controller implements an eight-way, dual directory, Set Associative cache using SDRAM. When instructions or data are stored in cache, they are grouped into sets of eight 64-byte lines. The X5 maintains

an index to each of the eight sets. It also keeps track of the tags used internally to identify each cache line. Dual tag directories allow simultaneous processor requests and system bus snoops, thus reducing resource contention and access time.

Snooping (checking to see if instructions or data required by one processor can be found in another CPU's L1 or L2 cache) is one way the F50 maintains cache coherency across all processors. Much of the snooping is performed in the X5 controllers since they maintain awareness of L1 cache as well as L2. Some snooping also occurs in the memory-I/O controller. The hit rate on the 256 KB X5 cache controllers is comparable to a 1 MB direct mapped cache using 32-byte cache lines.

Figure 21 shows a logical diagram of the internal architecture.

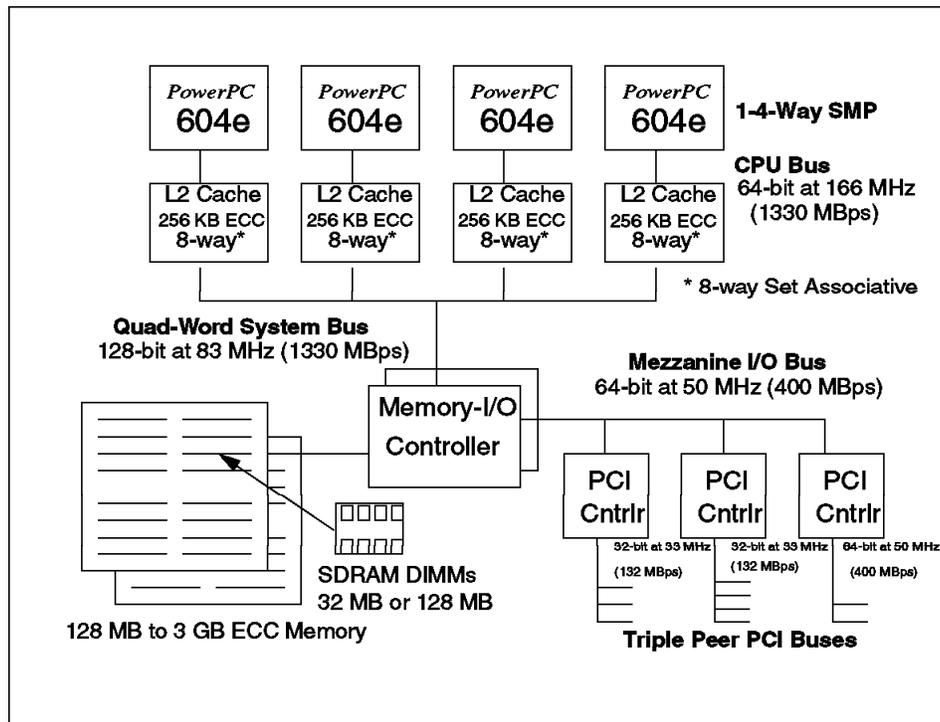


Figure 21. RS/6000 Model F50 Architecture

Memory-I/O Controller

The memory-I/O controller is capable of providing sustained memory bandwidth of over 1 GB per second and I/O bandwidth of up to 400 MB per second. The memory controller supports up to two memory cards, with up

to eight increments of SDRAM on each card. Each increment represents a pair of DIMMs. Two types of memory are supported; one using 32 MB DIMMs and the other using 128 MB DIMMs. Memory must be installed in pairs. Each module provides 72 bits of data (64 data bits plus 8 bits of Error Correction Code (ECC)) which, when paired with another module, forms the 144-bit memory interface.

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.

MESI Protocol

To maintain coherency between processors, the F50 uses an enhanced MESI (modified, exclusive, shared, invalid) protocol. MESI is explained in "MESI Protocol" on page 35.

This allows for cache-to-cache transfers (intervention) on both modified and shared data found in the L1 and L2 caches without having to go through the memory subsystem. The data can be transferred directly from one cache to another, thus improving system efficiency.

The memory I/O controller also acts as a snooping agent. This is accomplished by keeping track of the latest reads from I/O in an on-chip directory. This directory is searched whenever a processor request is issued on the SMP bus. This saves time because the request does not have to be forwarded to I/O to be snooped. When I/O issues a request, the controller places the requested address on the SMP bus to be snooped.

PCI Mezzanine Bridge

The memory I/O controller implements a 64-bit plus parity multiplexed address and data bus for attaching several PCI I/O bridges. This bus runs independently from the SMP system and the SDRAM buses. The bridge bus is clocked at 50 MHz, and there are two 32-bit and one 64-bit PCI I/O buses. The peak bandwidth of this bus is 400 MB per second. Multiple requests from each I/O bridge and the processors may be active at any time. The

PCI bridges contain over 700 bytes of internal data buffering and will support either a 32-bit or a 64-bit PCI data bus configuration.

3.3.3 Standard Features

Most standard features of the Model F50 are the same as those of the RS/6000 Model F40. For that reason, only the different components of the Model F50 are described in the following section.

Processor Subsystem

The Model F50 server base offering includes a 166 MHz PowerPC 604e processor with 32 KB of both instruction and data Level 1 (L1) cache and an ECC X5 L2 cache comprised of 256 KB per processor in an 8-way set associative implementation.

Its SMP design allows easy expandability through the addition of a second, third, or fourth 166 MHz or 332 MHz processor (of like speed) each with its own L1 and L2 cache.

Memory

The base configuration features a 128 MB of Error Checking and Correcting (ECC) SDRAM. It consists of four 32 MB 200-pin DIMMs with a cycle time of 10 ns. A maximum of 3 GB of memory is supported and addressable when both memory cards are used. A maximum of 4 GB could physically be installed, but is not addressable nor supported by the hardware.

Reliability, Availability, Serviceability Features

The Model F50 integrates key reliability, availability, and serviceability (RAS) attributes including:

- ECC on system memory and L2 cache which corrects environmentally induced single-bit intermittent memory failures as well as single-bit hard failures. With ECC, the majority of L2 cache and memory failures will not impact system operation. ECC also provides double-bit memory error detection which protects data integrity in the event of a double-bit memory failure.
- The system, I/O, and PCI buses are designed with parity error detection.
- Hot-swappable disk drives so users can take disks offline or put them back into service without interrupting the operation of their system.
- RAID 0, 1, and 5 for internal devices (optional through an adapter) which provides data integrity and fault tolerance in the event of disk drive failure.
- Online (concurrent) and remote (LAN or modem) system diagnostics.

- AIX's Journaled File System maintains file system consistency and prevents data loss when the system is abnormally halted, for example, due to a power failure.
- Service Director for RS/6000 (available at no additional charge if your RS/6000 processor is covered by an IBM Warranty or IBM Maintenance Service Agreement) monitors and analyzes system errors and if needed, can automatically place a service call to IBM without customer intervention. It also performs problem analysis on a subset of hardware-related problems and can automatically report the results to IBM Service. This can reduce the effect of business disruptions due to unplanned system outages and failures.

Service Processor

A service processor is included as standard and provides an immediate means to diagnose, check status, and sense operational conditions of a remote server, even when the main processor is inoperable. It enables or contributes to the following RAS functions:

- Integrated system environmental monitoring/alerting. This includes AC/DC voltage, fan speed, and temperature sensing.
- Early Power Off Warnings (EPOW) and error log analysis and alert.
- Auto-dial out through the use of Service Director, in order to call the IBM Service Center. The objective is to predict a potential component failure and automatically dispatch IBM Service to take preventative maintenance measures to avoid an outage.
- Programmable by the system administrator to reboot after a power loss, hardware checkstop failures, machine check interrupt, and operating system hang or failure.

Bays

The standard bays are composed of four media and six hot swap SCSI disk bays which accommodate the following drives:

- Two of the media bays accommodate the following standard devices:
 - 3.5-inch diskette drive bay that accommodates a standard 1.44 MB diskette drive.
 - 5.25-inch CD-ROM drive bay that accommodates a 20X SCSI-2 CD-ROM drive. The tray-loading mechanism provides support for 12 cm and 8 cm disks.
- The additional two media bays are empty on standard configuration. They can accommodate CD-ROM drives, tape drives, or 4.5 GB Ultra SCSI disks.

- Depending on the base 6-Pack (SCSI or SSA), the base disk is one of the following:
 - SCSI** One hot swap disk bay accommodates the standard 4.5 GB SCSI-2 F/W disk drive. This can be upgraded to a 9.1 GB Ultra SCSI disk drive in the initial order.
 - SSA** At least the boot disk must be SCSI; so the 4.5 GB Ultra SCSI Media Bay disk must be ordered. Two additional SSA disks (4.5 GB or 9.1 GB) are required when an SSA 6-Pack is used.
- The additional available five hot swap disk bays are empty in system standard configurations.

Figure 22 shows the bay configuration of the system.

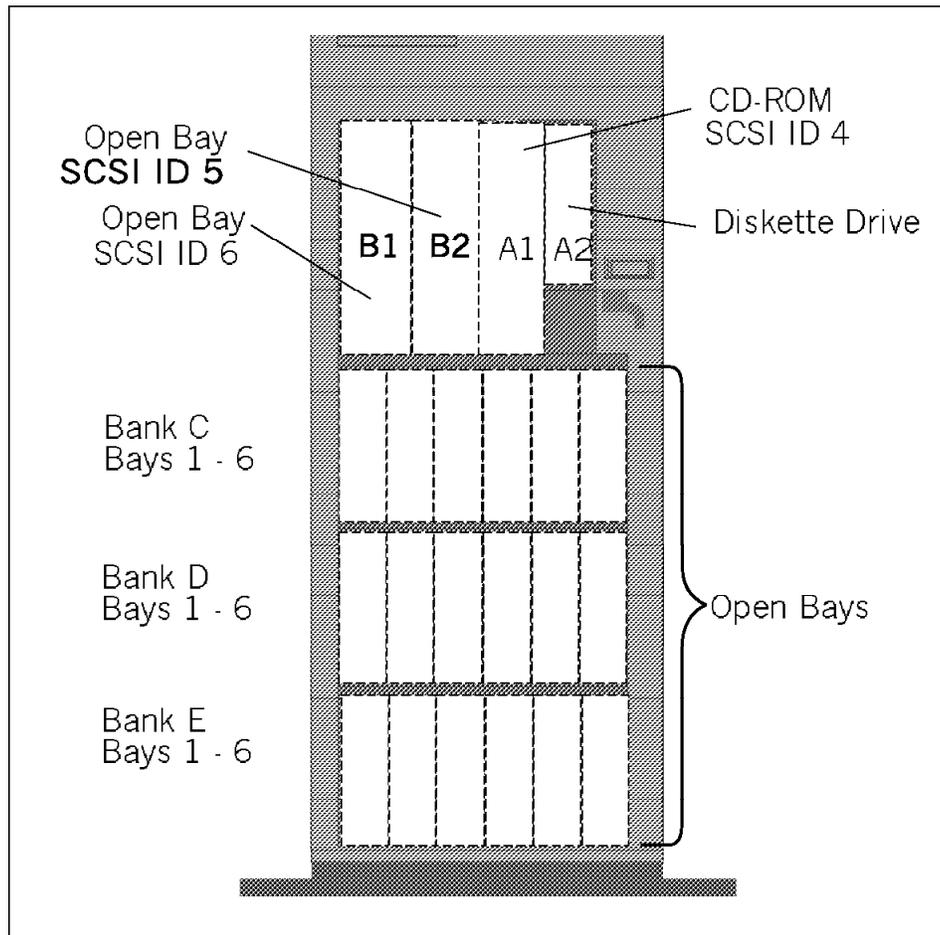


Figure 22. RS/6000 Model F50 Bays

Integrated SCSI-2 Fast/Wide Adapter

The system comes standard with two integrated PCI-based SCSI-2 Fast/Wide adapters. One of them can be used only for internal, single-ended (Fast/Wide, narrow, or both) SCSI-2 devices (maximum of seven). The other one is externalized, on to which up to six internal and two external devices can be connected.

Adapter Slots

Nine adapter slots are provided in the following configurations:

- Five slots for PCI adapters (32-bit)
- Two slots for PCI adapters (64-bit)
- Two slots that can accommodate either a PCI or an ISA adapter

The bus implementation is a tri-peer PCI bus architecture resulting in an aggregate I/O bandwidth of up to 400 MBps. All adapters have undergone testing to assure system compatibility.

Standard I/O Ports

The system provides the following standard I/O ports:

- Keyboard port** The Quiet Touch Keyboard, an optional feature, may be plugged into this port.
- Mouse port** A port is provided to support an optional three-button mouse.
- Serial port** Three 9-pin D-shell serial connectors for external RS232 access.
- Parallel port** One 25-pin parallel port that maintains complete compatibility with ISA, EISA, and Micro Channel parallel ports.

Communication Port

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

Figure 23 on page 75 shows the rear view of the system.

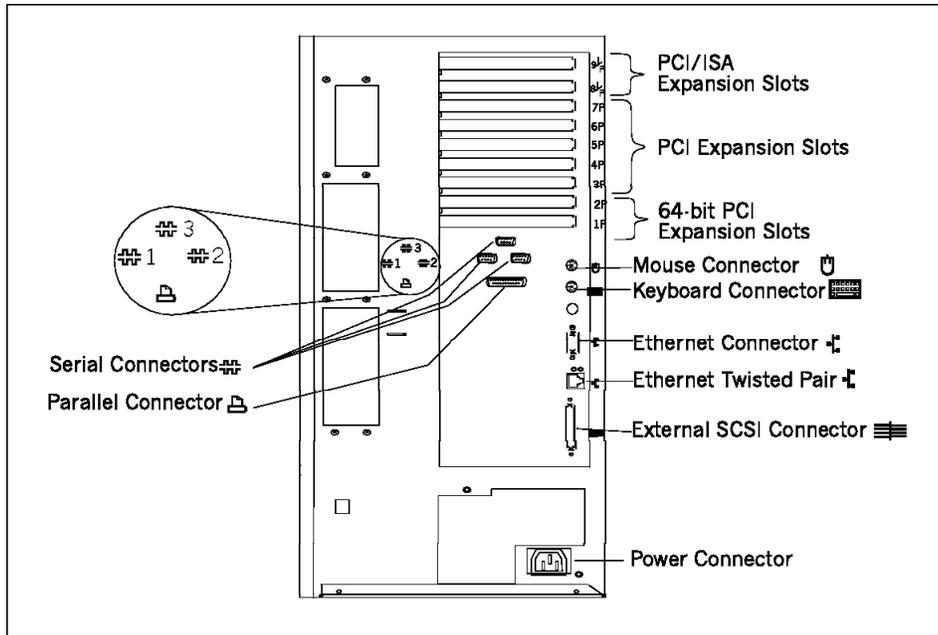


Figure 23. RS/6000 Model F50 Rear View

AIX Operating System

The Model F50 includes a one or two user, full-function AIX Version 4.2.1 or Version 4.3 package at no additional charge.

3.3.4 Supported Optional Features

In this section, *optional features* are defined only as internal optional devices. Adapters are described in more detail in Chapter 5, "Adapters, Overview and Configuration" on page 147, and storage subsystems are described in Chapter 4, "Internal Storage Features" on page 91.

Additional Processor

A second, third or fourth 166 MHz or 332 MHz processor can be added. All processors must have the same clock rate; they cannot be mixed.

Memory

Either 64 MB (two 32 MB) or 256 MB (two 128 MB) SDRAM DIMMs (200 pin, 10 ns) can be added. For using more than 16 DIMMs, a second memory expansion card must be added. Due to hardware addressing restrictions, the maximum memory is limited to 3 GB.

Optional Disk Bays

As described in 3.3.3, “Standard Features” on page 71, the RS/6000 Model F50 features four media bays, six standard, and twelve optional half-height standard hot-swappable disk bays. To add six half-height hot-swappable disk bays to the system, additional hardware is required. It can be ordered as a hardware feature named SSA 6-Pack or SCSI Hot Swap 6-Pack, depending on which optional bays are going to be installed. In addition, a power supply and one connection cable for each 6-Pack must be ordered.

Internal Disk Drives

Additional disk drives can be mounted in either SSA or SCSI hot swap 6-Packs. SCSI and SSA 6-Packs can be mixed, but the first 6-Pack must be SCSI.

Internal SCSI or SSA disk drives in varied capacities of the industry standard 3.5-inch form factor can be installed on the system. Internal disk drives available have 4.5 GB or 9.1 GB capacities.

Both 1-inch height (half-height) and 1.6-inch (full-height) disk drives are supported, but while half-height disks require only one hot-swappable bay, the full-height ones occupy two. For this reason, the maximum number of full-height disk drives that can be installed is nine (two bays for each drive), and the maximum number of half-height disk drives is 18 (one bay for each drive).

For more information, refer to 4.2.3, “Internal SSA Devices” on page 118 and “Disk Drives” on page 98.

Internal Tape Drives

All four types (4 GB/8 GB 4 mm, 12 GB/24 GB 4 mm, 5 GB/10 GB 8 mm, or 20 GB/40 GB 8 mm) of internal tape drives can be installed in an internal, non-hot-swappable media bay.

Refer to Table 17 on page 100 for more information.

Internal CD-ROM Drives

The CD-ROM provides 20X performance with the convenience of loading CDs without a caddy. Refer to “CD-ROM Drives” on page 99 for details.

SCSI Adapters

The following five PCI SCSI adapters are supported:

- Ultra SCSI Adapter (# 6206)
- SCSI-2 Fast/Wide PCI-Bus Adapter (# 6208)

- Ultra SCSI Differential Adapter (# 6207)
- SCSI-2 Differential Fast/Wide Adapter (# 6209)
- SCSI-2 Fast/Wide RAID Adapter (# 2493)

Table 12 on page 60 shows the maximum number of both internal and external devices that can be attached to a SCSI-2 adapter and related cables for the RS/6000 Model F50. For additional information, refer to 4.1.8, “SCSI Adapters” on page 100 and to “SCSI Adapters” on page 59.

Serial Storage Architecture (SSA) and SSA RAID

The Model F50 supports both external and under-the-covers SSA disk subsystems. The SSA Multi-Initiator/RAID EL Adapter (# 6215) can be used to control up to three 6-packs of SSA disks. When implementing SSA internally, SSA-enabled hard disks (# 3071 and # 3072) and 6-pack backplanes (# 6538 and # 6539) are required. SSA disks and backplanes are hot swap capable.

If the F50 is booted from hard disk, SCSI attachment is required. A boot disk in the first 6-pack may be used for this purpose, leaving the second and third 6-packs available for internally attached SSA drives. A maximum of one SCSI 6-pack is allowed in configurations, including both SSA and SCSI 6-pack backplanes.

When SSA is desired on the first 6-pack, feature # 3005 must be ordered to move the system boot disk to a media bay, and the first SCSI 6-pack is replaced with an SSA backplane using select feature # 6538. Optional SSA backplane feature # 6539 is used to implement SSA on the second and third 6-packs.

Refer to “Serial Storage Architecture (SSA) and SSA RAID” on page 62 and to 4.2.4, “SSA Adapter” on page 119, for more information.

Graphics Adapters

For 2D graphics applications, the Model F50 offers the POWER GXT120P 2D accelerator, which meets the graphics needs of many commercial applications and provides an entry point for workstation users. It is also a cost-effective choice for server environments that require a graphical user console.

For advanced 3D graphics function, the Model F50 offers the POWER GXT800P graphics accelerator, with optional hardware texture mapping, which is ideal for solid modeling and visualization. These graphics accelerators utilize IBM’s multithreaded PHIGS and OpenGL libraries, which exploit the F50’s SMP architecture to provide excellent performance.

It is possible to use a GXT800P adapter from another RS/6000 models such as the 43P-140 or F40, but the firmware level must be checked first. Before taking the adapter out of the machine, you can determine the level with the command `lscfg -l smint0`. You will get the base card part number. If the part number is 07L7111, the card can be used in the F50. If it is 93H2789, the card needs a firmware update. All other part numbers are not supported in the F50.

For details, refer to 5.2, “Graphics Adapters” on page 158.

Note

No multimedia features (audio or video) are supported in the Model F50.

Communication Adapters

The following communication adapters are supported:

- Auto LANStreamer Token-Ring PCI Adapter (# 2979, until July 24, 1998)
- IBM PCI Token-Ring Adapter (# 2920)
- PCI Ethernet Adapters
 - PCI Ethernet BNC/RJ-45 Adapter (# 2985)
 - PCI Ethernet AUI/RJ-45 Adapter (# 2987)
 - 10/100 Mbps Ethernet (# 2968)
 - Fast Etherlink XL (# 2986)
- ATM Adapters
 - TURBOWAYS 155 PCI UTP ATM Adapter (# 2963)
 - TURBOWAYS 155 PCI MMF ATM Adapter (# 2988)
 - TURBOWAYS 25 PCI ATM Adapter (# 2998)
- ISA X.25 Adapter (# 2961)
- ISA 4-Port Communications Controller (# 2701)
- ISA 2-Port Multiprotocol PCI Adapter (# 2962)
- FDDI Adapters
 - SysKonnnect SK-NET FDDI-LP SAS (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS (# 2743)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)

More detailed technical information about these PCI and ISA adapters can be found in Section 5.1, "Communication Adapters" on page 147.

3.3.5 Configuration Notes

When planning or implementing an RS/6000 Model F50, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

- 9.1 GB Disk drives (1.6-inch height) require two hot-swappable bays.
- There can be a maximum of two graphics adapters installed in the system.
- A SCSI boot disk must be configured in the media bay if all three 6-Packs use SCSI RAID or SSA disks.

3.4 RS/6000 Model H50 Product Description

The H50 is the first IBM rack-drawer version of a new generation of 1- to 4-way symmetric multiprocessing (SMP) enterprise servers providing unprecedented commercial price/performance in its class.

The RS/6000 Model H50 base offering includes a 332 MHz 604e PowerPC processor with X5-Cache system, 128 MB of error checking and correcting (ECC) synchronous DRAM (SDRAM) (expandable to 3 GB), an ECC Level 2 cache comprised of 256 KB per processor in an 8-way set associative implementation, a 4.5 GB high-performance disk drive, a 20X maximum speed CD-ROM, and a 1.44 MB 3.5-inch diskette drive. The H50 is packaged in a 9 X 16 (9 I/O slots, up to 16 disk/media bays) 8-EIA rack drawer. Unlike previous RS/6000 rack-mounted systems, the Model H50 is black, and it is recommended for use in the Model S00 rack, also black.

The RS/6000 Model H50 includes a modular disk subsystem that allows fast, easy addition and replacement of drives. It has a hot-swappable storage maximum capacity of 118.2 GB and a flexible I/O subsystem of five PCI (32-bit), two PCI (64-bit) and two shared PCI/ISA slots. The I/O channels supported are the industry-common, high-performance PCI bus in addition to the traditional ISA bus.

Figure 24 on page 80 shows an RS/6000 Model H50.



Figure 24. RS/6000 Model H50 Installed in the S00 Rack

3.4.1 Model H50 Highlights

The new generation Enterprise Server Model H50 offers some of the most advanced features in the industry:

- SCSI and Ultra SCSI
- *Under-the-Covers* SSA RAID capability
- High-capacity 4.5 GB and 9.1 GB disks
- Two integrated SCSI Fast/Wide controllers
- Three PCI busses assure very high I/O performance for the nine PCI slots
- A standard service processor for enhanced reliability and serviceability
- Redundant power and cooling capability for enhanced reliability
- Packaged in an 8-EIA rack drawer: a 9-slot (7 PCI/2 PCI or ISA), 16-bay configuration (3 media/13 hard disk)

3.4.2 Model H50 System Architecture

The H50 takes advantage of the same system architecture as the Model F50. There are small variations, such as an additional port, and optional second power supply. For a summary of the main characteristics, refer to Section 3.3.2, “Model F50 System Architecture” on page 68.

3.4.3 Standard Features

Most standard features of the RS/6000 Model H50 are the same as those of the RS/6000 Model F40 and RS/6000 Model F50. For that reason, only the different components of the RS/6000 Model H50 are described in detail in the following section.

Processor Subsystem

The Model H50 base system includes one PowerPC 604e microprocessor operating at 332 MHz and couples an X5 8-way set associative ECC Level 2 cache of 256 KB for each processor. The X5 Level 2 cache and microprocessor communicate at a bus speed of 166 MHz.

Its SMP design allows easy expandability through the addition of a second, third, or fourth 332 MHz processor each with its own L1 and L2 cache. Either additional 1- or 2-way cards can be selected.

Memory

The base configuration features a 128 MB of Error Checking and Correcting (ECC) SDRAM. It consists of four 32 MB 200-pin DIMMs with a cycle time of 10 ns. A maximum of 3 GB of memory is supported and addressable when both memory cards are used. A maximum of 4 GB could physically be installed, but is not addressable nor supported by hardware.

Reliability, Availability, Serviceability (RAS) Features

The Model H50 features the same RAS features as the F50 with the addition of optional dual power supplies. For more information about the RAS features, refer to “Reliability, Availability, Serviceability Features” on page 71 and to “Service Processor” on page 72.

Bays

The H50 provides 16 media and disk bays. Twelve are hot-swappable disk bays as shown in Figure 25 on page 82.

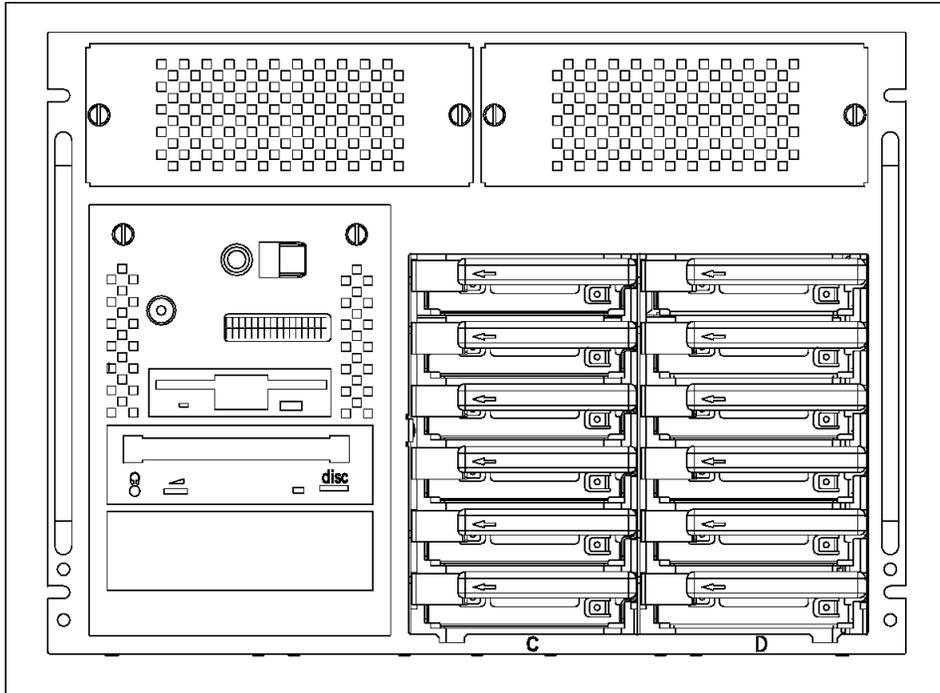


Figure 25. RS/6000 Model H50 Bays

Two media bays are filled by the base CD-ROM and diskette drive. One media bay is available for 5.25-inch x 1.6-inch media drives or a second disk if boot disk mirroring is used. One additional internal bay is provided to hold the 4.5 GB SCSI IPL disk if the optional SSA or SCSI RAID configuration is selected.

One of the disk bays is filled with a 4.5 GB hot-swappable disk drive; five remain available, and six optional hot-swappable disk drive bays can be ordered.

Integrated SCSI-2 Fast/Wide Adapter

Two integrated SCSI-2 Fast/Wide adapters are provided. Refer to “SCSI Adapters” on page 85 for details.

Slots

Nine adapter slots are provided of the following configuration:

- Five slots for PCI adapters (32-bit)
- Two slots for PCI adapters (64-bit)
- Two slots that can accommodate either a PCI or an ISA adapter

Standard I/O Ports

The system provides the following standard I/O ports:

- Keyboard port** The Quiet Touch Keyboard is an optional feature that may plug in this port.
- Mouse port** A port for an optional three-button mouse is provided.
- Serial port** Three 9-pin D-shell RS232 connectors are provided.
- Parallel port** One 25-pin parallel port maintains compatibility with ISA, EISA, and Micro Channel parallel ports.
- Power Controller** The Power Controller Interface (PCI) allows the H50 to control external devices and racks. It does not allow the H50's power to be controlled by a PCI attached device. To remotely control the power, use the service console functions.

Communication Port

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

Figure 26 shows the rear view of the H50.

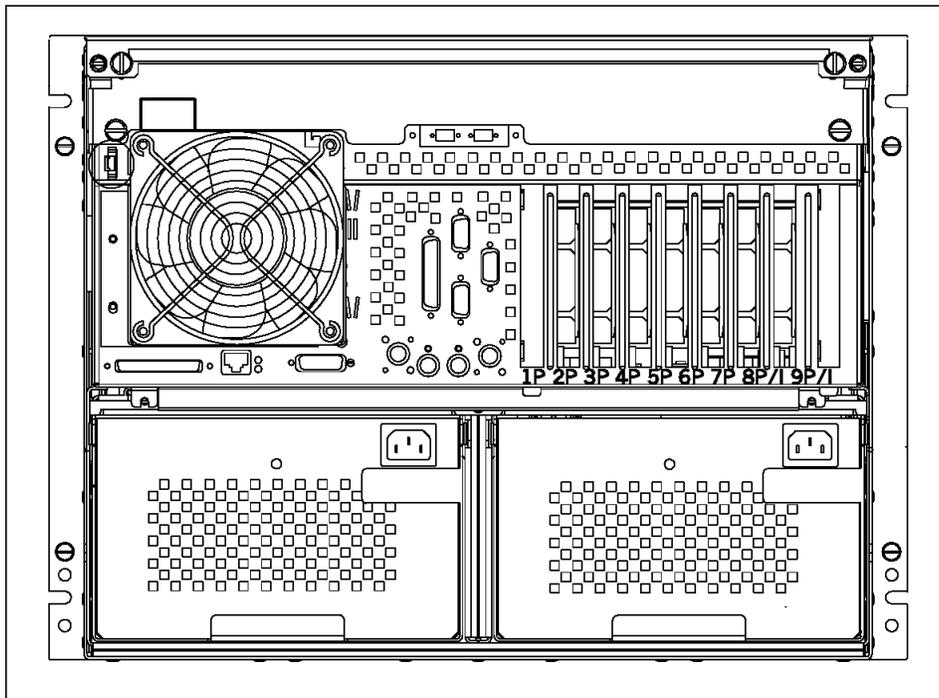


Figure 26. RS/6000 Model H50 Rear View

AIX Operating System

The Model H50 includes a one or two user, full-function AIX Version 4.2.1 or Version 4.3 package at no additional charge.

3.4.4 Supported Optional Features

In this section, *optional features* are defined only as internal optional devices. Adapters are described in more detail in Chapter 5, "Adapters, Overview and Configuration" on page 147 and storage subsystems are described in Chapter 4, "Internal Storage Features" on page 91.

Additional Processor

A second, third, or fourth 332 MHz processor can be added.

Memory

Either 64 MB (two 32 MB) or 256 MB (two 128 MB) SDRAM DIMMs (200 pin, 10 ns) can be added. For using more than 16 DIMMs, a second memory expansion card must be added. Due to addressing restrictions, the maximum memory is limited to 3 GB.

Optional Disk Bays

The Model H50 can accommodate up to two hot-swappable 6-Packs. The hot swap capability is enabled through the backplane associated with each 6-Pack. The H50 provides a hot-swappable backplane as standard with the first 6-Pack.

Internal Disk Drives

The Model H50 disk offerings allow internal hot-swappable disk capacities ranging from 4.5 GB to 118.2 GB in either 4.5 GB or 9.1 GB increments. The design feature that provides the upgradeability is the 6-Pack backplane into which up to six disk drives can be installed. Individual drives are mounted in the 6-Packs through carriers.

The 9.1 GB (# 2911, # 3019, and # 3072) disk drives require two bays. This allows for up to three 9.1 GB drives maximum per 6-Pack. The 9.1 GB Ultra SCSI 16-bit 1-inch (25 mm) High Disk Drive (# 2913) and Select (# 2919) are also available. These disk drives each require only one bay, which allows for up to six 9.1 GB drives per 6-Pack. Hard disks located in the media bays are non-hot-swappable. They can be mirrored if desired by using the RAID mirroring function contained within AIX. The H50 has one available media bay. If the bay is occupied by a hard disk, then any other optional devices, such as a tape back-up unit, must be attached externally to the system. Hot swap capability is standard on all supported drives when installed in either of the two SCSI or SSA 6-Packs available. The hot swap capability is enabled through the backplanes located within the SCSI or SSA 6-Packs.

For more information, refer to 4.2.3, “Internal SSA Devices” on page 118 and to “Disk Drives” on page 98.

Internal Tape Drives

All four types (4 GB/8 GB 4 mm, 12 GB/24 GB 4mm, 5 GB/10 GB 8 mm, or 20 GB/40 GB 8 mm) of internal tape drives can be installed in an internal non-hot-swappable media bay.

Refer to Table 17 on page 100 for more information.

Internal CD-ROM Drives

The CD-ROM provides 20X performance with the convenience of loading CDs without a caddy. Refer to “CD-ROM Drives” on page 99 for details.

SCSI Adapters

The Model H50 disk subsystem incorporates dual integrated SCSI-2 Fast/Wide controllers. The standard configuration controls the 6-Pack number one and the removable media devices (for example, CD-ROM) with the first integrated SCSI-2 Fast/Wide controller.

Feature # 2444 provides the necessary cable to control 6-Pack number two from the second integrated SCSI-2 controller. If desired, 6-Pack number two can be added and controlled from an optional PCI SCSI-2 SE Fast/Wide adapter (# 6208 or equivalent) through the cable provided in feature # 2447.

The following five optional PCI SCSI adapters are supported:

- Ultra SCSI Adapter (# 6206)
- SCSI-2 Fast/Wide PCI-Bus Adapter (# 6208)
- Ultra SCSI Differential Adapter (# 6207)
- SCSI-2 Differential Fast/Wide Adapter (# 6209)
- SCSI-2 Fast/Wide RAID Adapter (# 2493)

For initial orders, a single SCSI 6-Pack connected to an integrated SCSI controller is provided as the default through a base SCSI cable. This base cable also connects the SCSI media devices. A single hot-swappable 16-bit 4.5 GB Ultra SCSI-enabled hard disk is standard.

SCSI RAID: If a SCSI RAID configuration is desired, the 16-bit 4.5 GB Ultra SCSI-enabled hard disk must be moved into a media bay for use as a boot drive by selecting feature # 6424.

Ultra SCSI: If an Ultra SCSI configuration is desired, the following discussion is applicable. A SCSI 6-Pack supports Ultra SCSI drives at Ultra speed. However, Ultra SCSI speed is not supported by the integrated dual SCSI Fast/Wide controllers.

Ultra SCSI RAID: If an Ultra SCSI RAID configuration is desired, the PCI Ultra SCSI Differential adapter (# 6207) can be used to support external Ultra SCSI RAID subsystems. Ultra SCSI RAID is supported in an under-the-covers configuration through the striping (RAID 0) and mirroring function (RAID 1) of AIX and using the PCI Ultra SCSI SE adapter (# 6206). Feature # 2493 does not support RAID under-the-covers at ultra speed.

SSA and SSA RAID

If an all SSA configuration is desired on the initial order, the 16-bit 4.5 GB ultra-enabled hard disk must be moved into a media bay for use as a boot drive by selecting feature # 6427. If the initial order includes a SCSI 6-Pack in the first position, SSA can be ordered later as an MES for the second 6-Pack only using feature # 6539.

The PCI SSA Multi-Initiator/RAID EL adapter (# 6215) is supported only in slots 1, 4, 7, and 9. When an SSA adapter (# 6215) is installed, no adjacent slot may be occupied by a Digital Trunk adapter (# 6309).

Graphics Adapter

The Model H50 provides support for the GXT120P entry-level (# 2838) 2D graphics adapter. This adapter is used in conjunction with graphics consoles as an alternative to ASCII if desired. If an ASCII console is desired, it would be supported through serial port S1 and would not require the use of this adapter. AIX systems can be ordered with or without a keyboard, mouse, or graphics adapter. For details, refer to 5.2, "Graphics Adapters" on page 158.

Communication Adapters

The following communication adapters are supported:

- Auto LANStreamer Token-Ring PCI Adapter (# 2979, until July 24, 1998)
- IBM PCI Token-Ring Adapter (# 2920)
- PCI Ethernet Adapters
 - PCI Ethernet BNC/RJ-45 Adapter (# 2985)
 - PCI Ethernet AUI/RJ-45 Adapter (# 2987)
 - 10/100 Mbps Ethernet (# 2968)
 - Fast Etherlink XL (# 2986)

- ATM Adapters
 - TURBOWAYS 155 PCI UTP ATM Adapter (# 2963)
 - TURBOWAYS 155 PCI MMF ATM Adapter (# 2988)
 - TURBOWAYS 25 PCI ATM Adapter (# 2998)
- ISA X.25 Adapter (# 2961)
- ISA 4-Port Communications Controller (# 2701)
- ISA 2-Port Multiprotocol PCI Adapter (# 2962)
- FDDI Adapters
 - SysKonnnect SK-NET FDDI-LP SAS (# 2741)
 - SysKonnnect SK-NET FDDI-LP DAS (# 2742)
 - SysKonnnect SK-NET FDDI-UP SAS (# 2743)
- Eicon ISDN DIVA PRO 2.0 PCI S/T Adapter (# 2708)

More detailed technical information about these PCI and ISA adapters can be found in section 5.1, “Communication Adapters” on page 147.

Model S00 Racks

Up to three Model H50 drawers, depending on their configurations, may be installed in the IBM 7014 RS/6000 Rack Model S00. The S00 rack is 1.6 meters high and meets the Electronics Industries Standards EIA-310C standard. An EIA is a unit measure of vertical mounting space that equals 44.45 mm (1.75 in). The Rack Model S00 has 32 EIA units of vertical mounting space for 19-inch-wide drawers. The standard rack comes with an AC power distribution unit that is mounted in an area separate from the drawer mounting area. AC and DC power distribution units are available.

The rack is black with a front door consistent with the new Enterprise Server image. Also available are features for bolting the rack to either a concrete or raised floor and for ruggedizing the rack for central office environments.

Figure 27 on page 88 shows the S00 rack with an H50 installed.

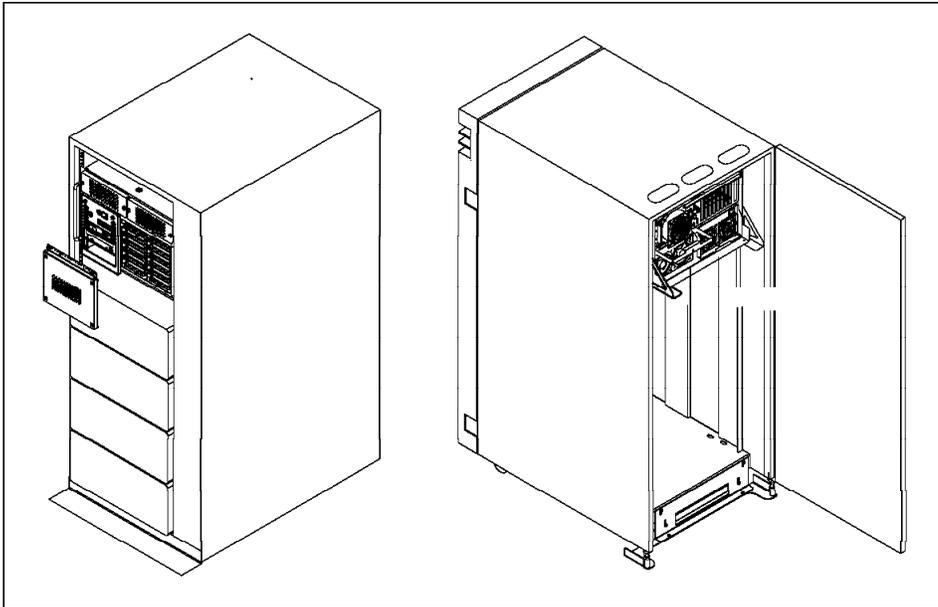


Figure 27. RS/6000 Model F50 S00 Rack, Front and Back Views

Redundant AC Power Supply

The Redundant AC Power Supply (# 6296) is a modular power supply which is identical to the power supply that is standard in each system unit. Each system has two power supply bays. This optional power supply allows the system to receive power from two independent sources and occupies the second power supply bay. This helps to improve reliability in the event of failure from one of the power sources.

RS/6000 High Available Cluster Server HA50

The HA50 is a two-node high availability cluster consisting of two RS/6000 Model H50s mounted in a 7014 Model S00 rack with AIX Version 4.2.1 or AIX Version 4.3 system software and HACMP 4.2.2 high-availability cluster software. One HA50 configuration includes the Model 7133-020 SSA Disk Storage Subsystem. Additional storage options can be configured with the HA50.

A sample HA50 server includes (based on early marketing information):

- One 332 MHz 604e PowerPC processor
- 256 MB of memory
- Two 4.5 GB internal SCSI IPL boot disks
- Redundant AC power supply

- Two PCI SSA adapters and SSA cabling
- Serial to Serial Port Cable for drawer/drawer connection
- Two 10/100 Mbps Ethernet Adapter
- One CD-ROM drive
- One AIX 4.2.1 or AIX Version 4.3 (one or two user) license
- One HACMP 4.2.2 license
- One 7133-020

The base HA50 offers configuration flexibility. Since the HA50 is comprised of RS/6000 Model H50, all H50 supported features, except SCSI and Ultra SCSI HACMP configurations, are supported by the HA50.

3.4.5 Configuration Notes

When planning or implementing an RS/6000 Model H50, give special consideration to the following limitations, which are also listed in the announcement letter for this product.

- 9.1 GB Disk drives (1.6 inch height) require two hot-swappable bays.
- There can be a maximum of two graphics adapters installed in the system.
- A SCSI boot disk must be configured in the media bay if all three 6-Packs use SCSI RAID or SSA disks.

Chapter 4. Internal Storage Features

In this chapter, the internal storage features of the Models E30, F40, F50, and H50 are described. Both storage standards, SCSI and SSA, are covered. The different technologies are discussed, and a short description of the storage devices and adapters follows. Additional information is provided in the cabling sections. The Hot Swap section describes the handling and management of the hot swap disk subsystems of the models F40, F50, and H50.

4.1 Small Computer System Interface

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: 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.

4.1.1 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 MBps 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.

4.1.2 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 MBps synchronous transfer rate for 8-bit transfers and 20 MBps for 16-bit transfers, and other enhancements 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.

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.

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.

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 multithreaded I/O on the SCSI bus. This feature is very important in multitasking 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.

4.1.3 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:
 - Fiber 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 MBps 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 expandability 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.

Overview of SCSI-III Standards

The breakdown of the SCSI-III standards is shown in Figure 28 on page 94 and described in Table 13 on page 94.

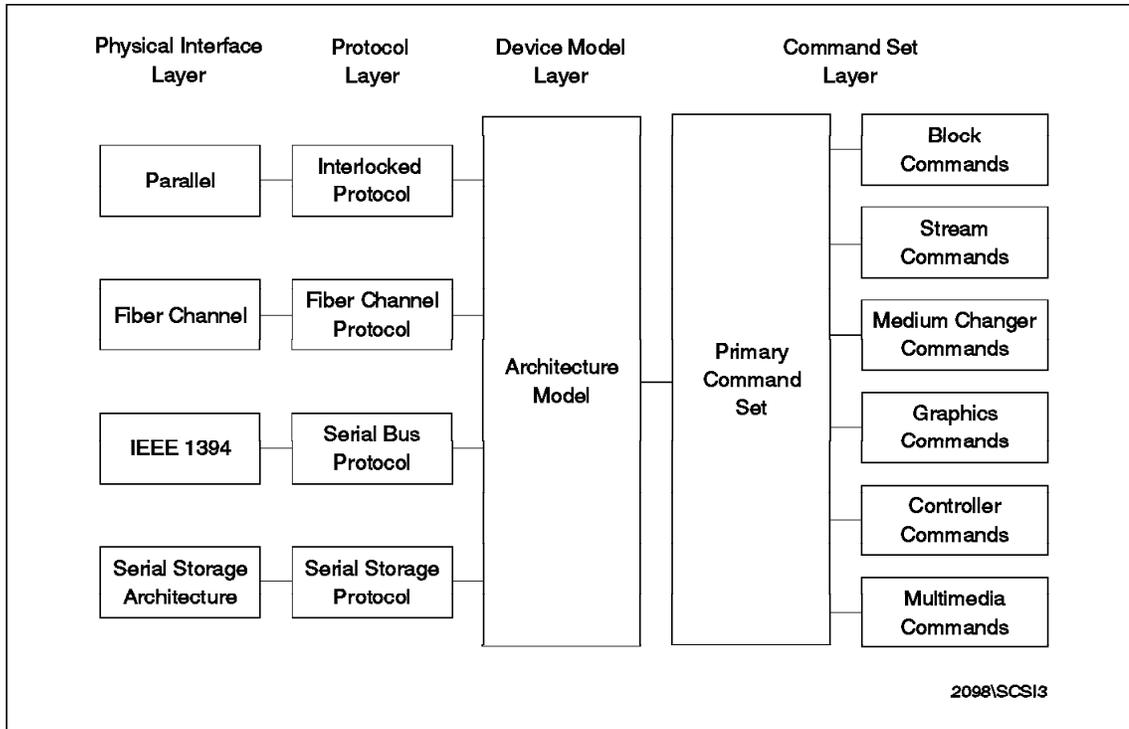


Figure 28. SCSI-III Standards Overview

Table 13 (Page 1 of 2). Overview of SCSI-III Standards		
Standard	Abbrev	Function
SCSI-III Interlocked Protocol	SIP	Describes the protocols used on the SPI bus.
Fiber 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.

Table 13 (Page 2 of 2). Overview of SCSI-III Standards

Standard	Abbrev	Function
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.

1. SCSI-III Parallel Interface (SPI)
2. Fiber Channel Physical and Signaling Interface (FC-PH)
3. IEEE 1394 High Performance Serial Bus
4. 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-II and SCSI-I. 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 MBps. Fast-20 Wide is a 16-bit bus with a maximum data transfer rate of 40 MBps. The “20” in Fast-20 refers to the clock speed of bus: 20 MHz, which is double the speed of SCSI-2 Fast.

The term *Ultra SCSI* is also 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).

4.1.4 SCSI Terminology

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

Fast Fast refers to the doubling of the data transfer rate from the SCSI 5 MBps to 10 MBps by doubling the clock rate. SCSI (that is, the original SCSI specification, or SCSI-I) is 5 MBps, which is produced by a clock speed of 5 MHz sending data down eight wires. SCSI-II Fast achieves 10 MBps 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 *burst* rate) of 20 MBps.

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 *Fast-20 Narrow*. *Fast-20 Wide* 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 MBps over an 8-bit data path. Wide Ultra SCSI, the 16-bit version of Ultra SCSI, can transmit a maximum of 40 MBps.

Ultra2 SCSI Ultra-2 is the latest extension to the SCSI-3 specification which allows transfers to take place at 80 MBps on a wide bus, or 40 MBps 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).

Wide SCSI

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).

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 acknowledgements 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.

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.

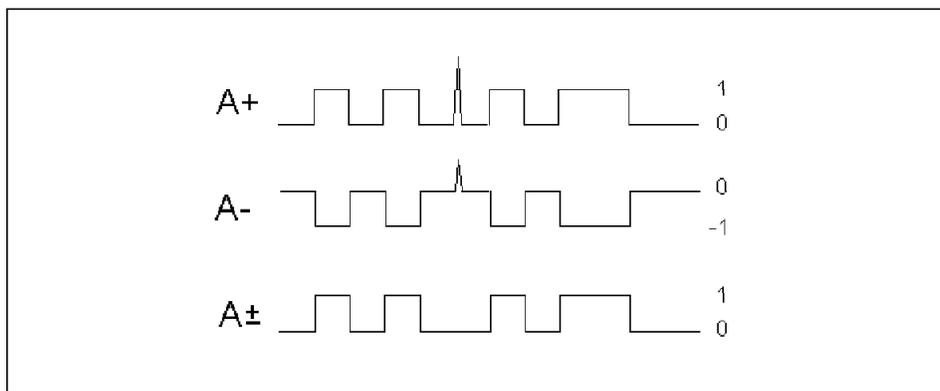


Figure 29. SCSI Differential Example

Consider Figure 29 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*, the result being A± as shown in the diagram. The spike is canceled out by the subtraction.

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 4.1.5, “SCSI Repeaters” on page 98, are more common.

4.1.5 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 14, the maximum length of the SCSI bus for a SCSI-2 F/W configuration is 3.0 meters. With a repeater, it is possible to extend this distance.

4.1.6 Summary of SCSI Specifications

SCSI defines many different modes of operation, including several different data transfer modes. The most common are shown in Table 14.

<i>Table 14. Different SCSI Modes</i>				
Mode	Bandwidth	Frequency	Max Data Transfer Rate	Max Cable Length Single Ended*
SCSI-I	8-bit	5 MHz	5 MBps	6 meters
Fast SCSI	8-bit	10 MHz	10 MBps	3 meters
Fast/Wide SCSI	16-bit	10 MHz	20 MBps	3 meters
Ultra SCSI (or Fast-20 SCSI)	8-bit	20 MHz	20 MBps	1.5 meters
Wide Ultra SCSI (or Fast-20 Wide SCSI)	16-bit	20 MHz	40 MBps	1.5 meters
Note: *The single-ended length refers to the length of the SCSI cable without using differential technology and without the use of a SCSI repeater.				

4.1.7 Internal SCSI Devices

The following section gives an overview of the internal SCSI devices and their main characteristics.

Disk Drives

Table 15 on page 99 shows the main characteristics of the optional disk drives that are available for the models E30, F40, F50, and H50.

<i>Table 15. Optional Disk Drives on Model E30, F40, F50, and H50</i>					
Formatted Capacity	9.1 GB	9.1 GB	4.5 GB	2.2 GB	2.1 GB
Feature Number	2908 / 2913	2911	2900 / 2901	3033 / 3083 (E30, F40 only)	3095 (E30 only)
SCSI Bus Rate MBps	40	40	40	20	40
Average Seek Time ms	6.7	8.9	7.9	8.5	8.5
RPM	7200	7200	7200	7200	5400
Average Latency ms	4.17	4.17	4.17	4.17	5.6
Interface	Wide Ultra SCSI	Wide Ultra SCSI	Wide Ultra SCSI	SCSI-2 F/W	Wide Ultra SCSI
Data Width	16 bit	16 bit	16 bit	16 bit	16 bit

CD-ROM Drives

Table 16 gives an overview of the two supported CD-ROM drives of the models E30, F40, F50, and H50.

<i>Table 16. Optional CD-ROM Drives on Model E30, F40, F50, and H50</i>		
CD-ROM Speed	8x (limited availability)	20x
Feature Number	# 2618	# 2619
Interface	SCSI-2 (8 bit)	SCSI-2 (8 bit)
Interface Speed (average)	1200 KB/sec	3000 KB/sec
Interface Speed (burst)	5.0 MBps	5.0 MBps
Bays	1 half-height	1 half-height
Note: Maximum CD-ROM Drives: 2		

Internal Tape Drives

Table 17 on page 100 shows the available internal tape drives for the models E30, F40, F50, and H50.

Tape Capacity	Feature number	Cartridge	Max Effective Transfer Rate	Interface
4 GB / 8 GB*	# 6142	IBM 4 mm DDS	400 KB/sec / 800 KB/sec*	SCSI-2 (8 bit)
12 GB / 24 GB*	# 6159	IBM 4 mm DDS	1.1 MB/sec / 2.2 MB/sec*	SCSI-2 (8 bit)
5 GB / 10 GB*	# 6147 (not for E30)	IBM 8 mm DDS	500 KB/sec / 1000 KB/sec*	SCSI-2 (8 bit)
20 GB / 40 GB*	# 6154 (not for E30)	IBM 8 mm DDS	3000 KB/sec / 6000 KB/sec*	SCSI-2 (8 bit)
Note: *Without / With Hardware Compression				

4.1.8 SCSI Adapters

In the E30, F40, F50, and H50, the following five PCI SCSI adapters are supported:

- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI Differential Ultra SCSI Adapter (# 6207)
- PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)
- PCI SCSI-2 Fast/Wide Differential Adapter (# 6209)
- PCI SCSI-2 Fast/Wide RAID Adapter (# 2493)

PCI Single-Ended Ultra SCSI Adapter

The PCI Single-Ended Ultra SCSI Adapter (# 6206) is an ideal solution for applications requiring large block data transfers (> 64 KB block size) in a multiple disk drive environment utilizing the SCSI-2 protocol. This adapter has a maximum data transfer rate of 40 MBps, which is twice the maximum data transfer rate of SCSI-2 Fast/Wide adapters. The PCI Single-Ended Ultra SCSI Adapter conforms to the SCSI-2 standard and Fast-20 (Ultra) documentation. Industry standard 68-pin connectors are incorporated. It provides 15 internal or external SCSI-2 addresses and occupies one PCI slot. Only AIX 4.2.1 or later supports this adapter.

PCI Differential Ultra SCSI Adapter

The PCI Differential Ultra SCSI Adapter (# 6207) is the next generation of SCSI-2 performance with maximum data transfer rates of 40 MBps. This adapter allows connection to external differential SCSI-2 F/W or 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 models and certain 7204 external disk drives.

The adapter conforms to SCSI-2 standard and the Fast-20 (Ultra) documentation. The industry standard 68-pin connectors are incorporated. It provides 15 external SCSI-2 addresses and occupies one PCI slot. AIX 4.2.1 and AIX 4.3 support this adapter.

PCI SCSI-2 Fast/Wide Single-Ended Adapter

The PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208) provides a single-ended SCSI-2 Fast/Wide interface that can burst data to devices on the SCSI bus at 20 MBps. This feature can provide both internal and external device support, connected to the same SCSI bus.

Both internal and external single-ended 8-bit or 16-bit SCSI devices can be attached. It occupies one PCI slot. AIX 4.2.1 or later is required to support this adapter.

PCI SCSI-2 Fast/Wide Differential Adapter

The PCI SCSI-2 Fast/Wide Differential Adapter (# 6209) is a PCI adapter that provides a differential SCSI-2 Fast/Wide interface. It can burst data to devices on the SCSI bus at 20 MBps. External differential 8-bit or 16-bit SCSI devices can be attached. One PCI adapter slot is required.

SCSI-2 Fast/Wide PCI RAID Adapter

The PCI RAID Adapter implements RAID level 0, 1, and 5 support for SCSI-2 attached disks. Configuration management, RAID algorithms, and error recovery are handled by an onboard 403 PowerPC. Three internal 16-bit SCSI-2 single-ended connectors are provided, one of which is shared with an external connector. The maximum number of addressable device IDs enabled by the adapter is 15 per bus or 45 devices per adapter. Cabling and housing restrictions may limit the actual total number of supported devices attached to the internal or external connector.

The main characteristics are as follows:

- Supports RAID 0, 1, and 5
- Controller conforms to ANSI Doc X3T9.2/86-109 revision 10k (SCSI-2) and the PCI Local Bus Specification Revision 2.1
- Accepts multiple commands per logical device from the system
- Configuration utilities provided through command line, SMITTY, SMIT, visual SMIT and web-based system management

- SCSI-2 data rate of up to 20 MBps (synchronous protocol)
- PCI bus interface; 132 MBps burst rate on the PCI bus
- 4 byte (32-bit) Bus Master
- Address and Data Parity Support
- Occupies one PCI bus slot
- Supports Command Tagged Queueing (as SCSI initiator)
- Adapter connector — A shielded 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1.27 mm (0.05 inches) apart
- Cable or mating connector — A shielded 68-conductor connector consisting of two rows of 34 female contacts with adjacent contacts 1.27 mm (0.05 inches) apart
- Attachment of internal/external single-ended 8-bit or 16-bit SCSI devices
- AIX 4.2.1 or later is required

4.1.9 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 face. Lots of questions always arise: 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 written above.

Described are the SCSI cabling of the five PCI SCSI adapters supported by the PCI-based RS/6000 servers:

- PCI Single-Ended Ultra SCSI Adapter (# 6206)
- PCI Differential Ultra SCSI Adapter (# 6207)
- PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)
- PCI SCSI-2 Fast/Wide Differential Adapter (# 6209)
- SCSI-2 Fast/Wide PCI RAID Adapter (# 2493)

Note

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

General SCSI Considerations

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

- Maximum SCSI bus length
- SCSI terminators
- SCSI device addresses
- SCSI bus width

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-2 Single Ended: 6 meters (20 feet)
- SCSI-2 Fast/Wide Single Ended: 3 meters (10 feet)
- SCSI-2 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 which have two SCSI connectors have internal cabling which 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 30 on page 104.

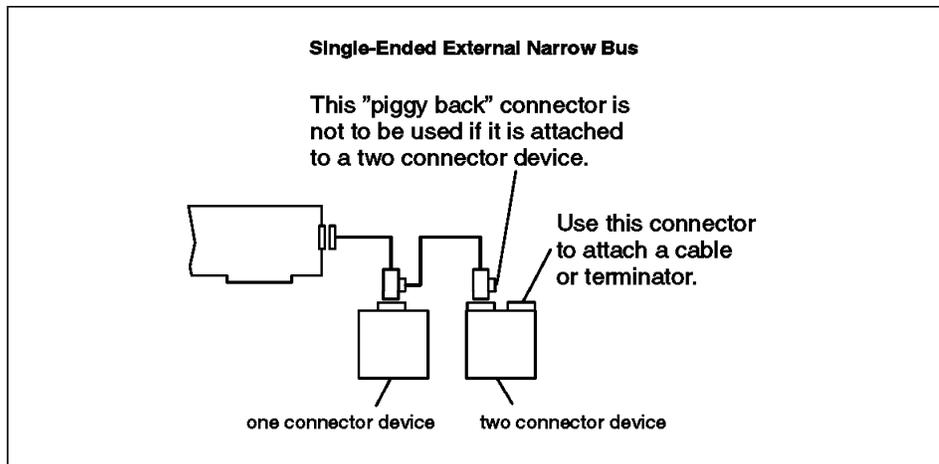


Figure 30. Piggy Back Connectors

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

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.

SCSI Device Addresses: The SCSI-2 Fast/Wide adapters support a maximum of 16 SCSI addresses, including devices and adapters. The default address for an adapter is 7, but it can be changed easily 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:
 - 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.

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.

Cabling the SCSI-2 Fast/Wide Single-Ended PCI Adapter (# 6208)

SCSI Fast/Wide devices support data rates of up to 10 MBps for 8-bit or 20 MBps for 16-bit transfers. If a configuration includes SCSI-2 Fast/Wide devices, the maximum cable length supported is three meters (approximately 10 feet). The maximum length includes the internal cabling of any device that has two SCSI connectors.

When attaching external devices, a maximum of *four independent physical enclosures* is allowed, provided each physical enclosure presents one device load to the SCSI bus. For example, a 7131-105 Multi Storage Tower cannot be attached in combination with any other internal or external load, but up to four different boxes, such as the 7204-112 External Disk Drives, can be attached.

If four external devices are attached and any one device is capable of SCSI-2 Fast transfers, then the total bus length is limited to three meters. In this case, the maximum bus length has been reached, and therefore no additional external or internal device can be attached to the SCSI bus.

A single enclosure containing any amount of multiple SCSI device loads attached externally to this adapter is supported, but subject to the following restrictions:

- Maximum combined internal and external cable length is three meters.
- Cable length between devices (loads on the cable) must be 0.1 meters apart at a minimum.
- No mixing of device widths (8-bit and 16-bit) unless the 68-pin-to-50-pin interposer (P/N 92F2565) is used.
- No internal devices are allowed to be attached to the adapter.
- It is recommended that dedicated adapters be used for external SCSI enclosures that contain more than two devices.

Section “Cables and Terminators for Single-Ended SCSI Adapters” on page 108 describes the cables, terminator features, part numbers, and lengths required.

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

The maximum supported cable length for this adapter depends on what type of devices are attached (SCSI-1, SCSI-2, to name two) and where they are attached (to the internal or external connector). Device types are classified as:

- SCSI-1 — maximum transfer rate of 5 MBps (one byte transfers)
- SCSI-2 Fast — maximum transfer rate of 10 MBps (one byte transfers)

- SCSI-2 Fast/Wide — maximum transfer rate of 20 MBps (two byte transfers)
- Ultra SCSI — maximum transfer rate of 20 MBps (one byte transfers)
- Ultra SCSI Wide — maximum transfer rate of 40 MBps (two byte transfers)

For this adapter:

- The maximum supported cable length for configurations without any SCSI-2 Fast or Ultra SCSI devices is 6 meters (approximately 20 feet).
- The maximum supported cable length for configurations that include SCSI-2 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. This adapter has circuitry that can detect the presence of a cable on the external connector, and the default configuration limits the SCSI bus speed to SCSI-2 Fast and Wide (20 MBps).
- 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-2 Fast devices running at SCSI-2 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-2 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). 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 tables in “Cables and Terminators for Single-Ended SCSI Adapters” 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)
- PCI SCSI-2 Fast/Wide Single-Ended Adapter (# 6208)

Adapter-to-First Device Cables: Table 18 on page 109 describes the cables for the SCSI-2 Fast/Wide Single-Ended Adapter and Single-Ended Ultra SCSI Adapter for connection to the first device.

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	2113	52G0174	1.5	Adapter-to-first device (where first device has one connector), 8-bit narrow bus
Host System	2115	06H6036	1.0	Adapter-to-first device (where first device has two connectors), 16-bit wide bus
Host System (# 6208 only)	2117	52G0173	0.94	16-bit Y-cable

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.
- For these adapters, the same cable can be used for either single-ended or differential attachments. The difference in Feature Code orders is the terminator type. Refer to Table 20 on page 111.
- The external connector on these adapters are the SCSI-3 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 19 on page 110 describes the device-to-device cables for single-ended applications.

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

Note: Most feature codes for cables are only orderable against the attachment device (7204, 7206, 7208, etc). 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.

Terminators for Use with These Adapters: Table 20 on page 111 describes the 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 Internal Cabling	part of cable FC	92F0322 (92G2566)	68-pin high density	16-bit external bus terminator

Cabling the Differential Adapters

The following tables describe the cables, terminator features, part numbers, and lengths for the following SCSI differential adapters:

- PCI Differential Ultra SCSI Adapter (# 6207)
- PCI SCSI-2 Fast/Wide Differential Adapter (# 6209)

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

Adapter-to-First Device Cables: Table 21 contains the adapter to first device cables.

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

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 Table 23 on page 114.

- For these adapters, 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 these adapters is the SCSI-3 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 22 contains the device to device cables.

<i>Table 22 (Page 1 of 2). Device-to-Device Cables for Differential Installations</i>				
Machine Type	FC	Cable Description	Part Number	Cable Length (meters)
DE External Device	2848/9134	Device-to-device (where second device has two connectors), 8-bit narrow bus	74G8511	0.6
DE External Device	2860/9139	Device-to-device (where second device has two connectors), 16-bit wide bus	52G9921	0.3
DE External Device	2884/9160	Device-to-device (where second device has two connectors), 16-bit wide bus	52G4291	0.6
DE External Device	2846/9132	Device-to-device (where second device has two connectors), 16-bit wide bus	52G4233	2.5
7027 HSD	2425			
DE External Device	2885/9161	Device-to-device (where second device has two connectors), 16-bit wide bus	88G5749	4.5
7027 HSD	3132	Device-to-device (where second device has two connectors), 16-bit wide bus	40H7351	6.0

<i>Table 22 (Page 2 of 2). Device-to-Device Cables for Differential Installations</i>				
Machine Type	FC	Cable Description	Part Number	Cable Length (meters)
DE External Device	2870/ 9146	Device-to-device (where second device has two connectors), 16-bit wide bus	88G5747	12.0
7027 HSD	3135			
DE External Device	2869/ 9145	Device-to-device (where second device has two connectors), 16-bit wide bus	88G5748	14.0
DE External Device	2868/ 9144	Device-to-device (where second device has two connectors), 16-bit wide bus	88G5746	18.0
7027 HSD	3136			

Note:

- Most feature codes for cables are only orderable against the attachment device (7204, 7206, 7208, etc). 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 codes are used.
- 9xxx feature codes are used for new build orders; 2xxx feature codes are used for MES orders.

Terminators for Use with These Adapters: These adapters have 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 adapter's 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 differential 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 time of removal. To maintain data integrity, the SCSI

bus should be inactive during the removal of adapters, cables, or terminators. Table 23 contains the terminators for differential installations.

<i>Table 23. Terminators for Differential Installations</i>				
Machine Type	FC	Terminator Description	Part Number	Connector
DE External Devices	part of cable FC	8-bit external bus terminator	87G1356	50-pin low density
DE External Devices	part of cable FC	16-bit external bus terminator	61G8324	68-pin high density

Cabling the PCI SCSI-2 F/W RAID Adapter (# 2493)

This section provides specific cabling information for the PCI SCSI-2 F/W RAID Adapter.

Note: This adapter cannot be used as a boot adapter, meaning you cannot put the rootvg volume group on any disk drive connected to this adapter.

SCSI-2 Single-Ended Cable Lengths for This Adapter: The maximum cable length supported for configurations without any SCSI-2 fast devices is 6 meters (approximately 20 feet). If a configuration includes SCSI-2 fast devices, the maximum cable length supported is 3 meters (approximately 10 feet). The maximum length includes the internal cabling of any device that has two SCSI connectors.

When you connect external devices, you can attach only one independent physical enclosure per SCSI bus (also called a *channel*). For example, you cannot attach a 7131-105 in combination with another external load. If you attach any external devices and any one device is capable of SCSI-2 Fast transfers, total bus length is limited to 3 meters.

You can externally attach a single enclosure that contains any amount of multiple SCSI device loads to the PCI SCSI-2 F/W RAID Adapter, but the following restrictions apply.

- The maximum combined internal and external cable length is 3 meters.
- Loads on the cable (cable length between devices) must be a minimum of 0.1 meters apart.
- You can mix device widths (8-bit and 16-bit) only if you use the 68-pin-to-50-pin interposer (FRU 92F0324, ASM 92F2565).

- You cannot attach any internal devices to the internal port adapter (J10) when the external port (J11) is in use.

No other adapters can be connected to any of the adapter SCSI ports. The adapter must have exclusive use of all attached drives. The adapter has on-board terminators that are always enabled. No terminator is needed at the adapter end of a SCSI cable.

An extender cable, Feature # 3131, allows internal adapter SCSI ports to be made available for external connection through an unused card slot.

4.2 Serial Storage Architecture

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 has been designed to overcome the bottlenecks of existing SCSI bus architectures.

Figure 31 shows SSA in a loop topology.

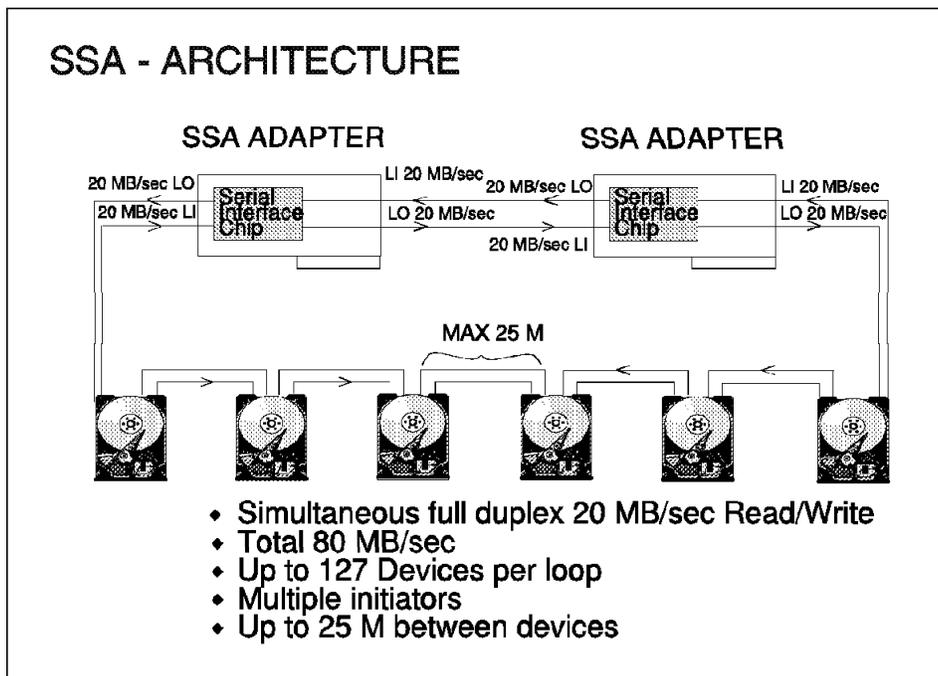


Figure 31. SSA Architecture

SSA uses a port design that enables inbound and outbound conversations simultaneously. An SSA Serial Interface Chip (SIC) on an adapter that is configured into a loop has a bandwidth of 80 MBps. There might be several hosts on a loop with the possibility of multiple hosts in a single server, thereby enabling a server to connect to several loops. SSA connects to its transport medium by means of 9-way micro-miniature D connectors.

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.

4.2.1 SSA Functions

Because SSA allows SCSI-2 mapping, all functions associated with initiators, targets, and logical units are translatable. Therefore, SSA can use the same command descriptor blocks, status codes, command queuing, and all other aspects of current SCSI systems.

Unlike SCSI and Fiber Channel Arbitrated Loop (FC-AL), SSA does not use bus phases. SSA does not use arbitration to contend for the bus. Instead, it makes use of frame multiplexing. All data is transported as frames.

SSA can be set up in three different topologies:

- String topology

A string is a simple linear network of two or more nodes. The ports at either end can be single-port nodes, and the others are dual port nodes.

- Loop topology

Loop topology is by far the most popular and most highly recommended topology. A loop is a cyclic network that contains only dual port nodes. Loops provide two data paths between nodes. The resulting benefit is higher bandwidth and higher reliability. Single node failures and single link breaks do not halt communication.

- Switch topology

Switch topology requires additional hardware and is the most complicated to set up and maintain. Its benefit is that it is the most fault tolerant topology of the three and allows a virtually unlimited network configuration.

4.2.2 SSA Advantages

In this section, each of the following advantages are discussed.

- High performance
- High availability
- Fault tolerance
- Hot plugging
- Connectivity
- Ease of use
- Reliability

High Performance

Only the nodes on either side of a virtual link are involved in transferring frames. A port can send a frame on its line-out while it receives frames on its line-in, thus enabling other nodes to initiate or process I/O requests simultaneously. This concept is known as *spatial reuse* and is very different from bus architectures which require the dedication of the entire bus for the duration of each data transfer. Fairness algorithms are used to divide bandwidth equally when there is contention for the same virtual loop, preventing a single device from saturating a virtual link. Spindle synchronization is integrated into SSA without the requirement of additional cables. SSA bandwidth is four times that of SCSI-2 and twice that of SCSI-3. Because multiple devices can process transfers, SSA uses its bandwidth much more efficiently than SCSI. SSA overhead is approximately 6 percent, effectively allowing SSA to run at a 94 percent efficiency rate (94 percent of the data that SSA handles is customer data).

High Availability, Fault Tolerance, and Hot Plugging

The loop and switch topologies do not have a single point of failure. As displayed in Figure 31 on page 115, the loop topology has alternate paths to each node and eliminates a single point of failure. It can connect 2 to 128 dual port nodes. A single break in the loop does not bring down the network. Frames are automatically rerouted across the available link. SSA devices are hot-pluggable by design; so nodes can be inserted into the loop without disruption. Redundant copies of data stored on SSA can be placed up to 2400 m away. SSA used in RAID systems makes up for overhead with its high speed and user data efficiency.

Greater Connectivity and Ease of Use

SSA allows the connection of 2 to 129 nodes in a string, 2 to 128 devices in a loop, and virtually unlimited nodes on a switch. Compare this with the 16 devices for SCSI. SSA supports multiple initiators. Manual ID selection and

termination is not required because this is an automated function of SSA and minimizes the possibility of configuration errors.

Reliability

SSA devices feature advanced diagnostics and monitoring, including a power-on self-test (POST), which is performed by all SSA devices. Each port remains in a disabled state until it has successfully completed a beginning-communication process.

4.2.3 Internal SSA Devices

The following section gives an overview of the internal SSA devices and their main characteristics.

- 4.5 GB SSA Hot Swap Disk Drive (# 3071)
- 9.1 GB SSA Hot Swap Disk Drive (# 3072)
- SSA Blank Disk Drive Module (# 2452)

4.5 and 9.1 GB SSA Hot Swap Disk Drive

The 4.5 GB SSA Hot Swap Disk Drive (# 3071) and the 9.1 GB SSA Hot Swap Disk Drive (# 3072) supplies one 4.5 GB and 9.1 GB SSA disk drive mounted in an auto-docking carrier. The auto-docking carrier includes a lever that facilitates safe insertion or removal of the disk drive module. The carrier also has Power, Ready, and Check light emitting diodes (LEDs) that show the status of the drive.

These are the main characteristics of both disks:

- Formatted Capacity: 4.5 GB (9.1 GB)
- Form Factor: 3.5 inch
- Average Seek Time: 8.0 ms
- Latency: 4.17 ms
- Rotational Speed: 7200 RPM
- SSA Data Transfer Rate: 20 MBps Full Duplex
- The 4.5 GB disk drive requires one disk bay, and the 9.1 GB disk drive requires two disk bays, with an SSA backplane and one SSA adapter in the system.

SSA Blank Disk Drive Module

The SSA blank disk drive module (# 2452) is needed to provide SSA loop signal continuity when an SSA disk drive is not used in an SSA disk drive slot. It is necessary that each SSA 6-pack backplane disk drive position be

populated with either an SSA disk drive (of any capacity) or the SSA blank disk drive module.

Each SSA 6-Pack backplane feature comes populated with the necessary SSA blank disk drive modules based on the order configuration. This is being offered as an MES only in the event that reconfiguration of the SSA 6-pack backplanes is required.

4.2.4 SSA Adapter

The PCI SSA Multi-Initiator/RAID EL Adapter (# 6215) is the only SSA adapter supported in Models E30, F40, F50, and H50.

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 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 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 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 plugged into SSA Multi-Initiator/RAID EL Adapter will operate in either non-RAID or RAID 5 mode, in a one-initiator configuration.

4.2.5 SSA Cabling

This section provides the configuration rules for the SSA adapter, a general introduction to SSA cabling, and details of the SSA cables.

The installation and service manuals for each SSA subsystem unit have more information that relates to connecting that unit; those manuals contain cabling details for sample configurations that meet the simplest requirements, where performance or availability are particularly important.

Rules for SSA Loops

For SSA loops that include the PCI SSA Multi-Initiator/RAID EL Adapter (# 6215, type 4-N), the following rules apply:

- Each SSA loop must be connected to a valid pair of connectors on the SSA adapter (that is, either connectors A1 and A2, or connectors B1 and B2).
- A maximum of 48 devices can be connected in a particular SSA loop.
- If the fast-write cache or RAID functions of the adapter are used, no other adapter can be connected in an SSA loop.
- If the fast-write cache or RAID functions of the adapter are not used, a second PCI SSA Multi-Initiator/RAID EL adapter (or a Micro Channel SSA Multi-Initiator/RAID EL Adapter) can be connected in the loop.
- Ports B1 and B2 have two sets of connectors. There is an internal B1, B2 pair and an external B1, B2 pair. Only one pair of these loop connectors can be used at a time.

SSA Loops, Links, and Data Paths

In the simplest SSA configuration, SSA devices are connected through two or more SSA links to an SSA adapter that is located in a using system. The devices, SSA links, and SSA adapter are configured in loops. Each loop provides a data path that starts at one connector of the SSA adapter and passes through a link (SSA cable) to the devices. The loop continues through the devices and returns through another link to a second connector on the SSA adapter.

The maximum permitted length for an external copper cable that connects two SSA nodes (disk drives or adapters) is 25 meters (82 feet).

Nodes that have the fiber-optic extender feature can be connected by fiber optic cable over a maximum permitted length of 2.4 kilometers (7874 feet).

When using the fiber optic extender (# 5500), you should consider the following:

- The feature supplies a pair of optical converters to perform the function of a single copper cable. To complete a loop you need four optical extenders or the order of two feature codes.
- The performance degradation depends on the optical fiber length. For a length up to 200 meters, the degradation is about 10 percent. For a length up to 2.4 kilometers, the degradation goes up to 75 percent.
- Overall, the maximum allowed loss for the complete path length is 3 dB.

For the following discussion, refer to Figure 32.

All devices that are attached to an SSA adapter card **1** are connected through SSA links **2**. Data and commands to a particular device pass through all other devices in the loop between the adapter and the target device.

Data can travel in either direction around a loop. The adapter can, therefore, access the devices **3** through two data paths. The using system cannot detect which data path is being used.

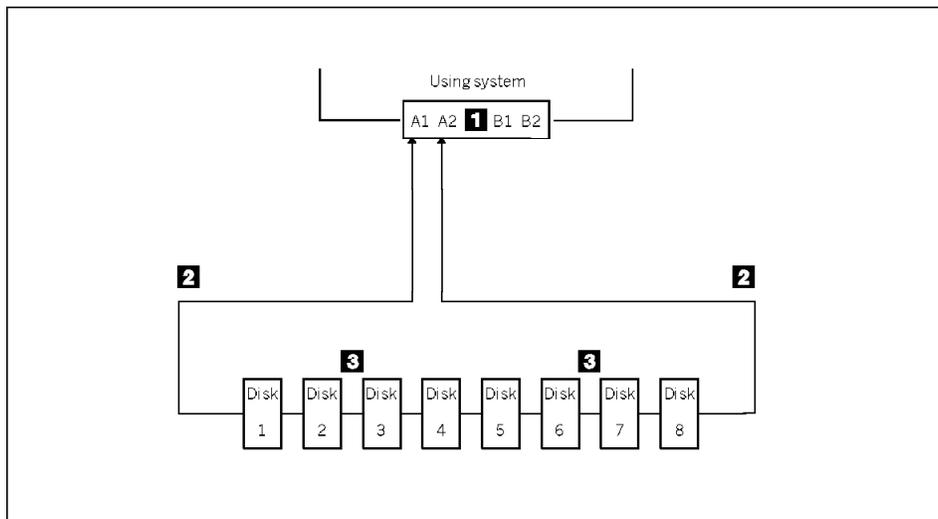


Figure 32. SSA Loops, Links, and Data Paths

If a disk drive fails or is turned off, the loop is broken. At this moment, the data paths to each disk drive is are longer available. The other disk drives continue to work through the remaining data path, but an error is reported to the system.

SSA Cables

Table 24 lists the available SSA cables.

Part Number	Feature Code	Length	
		m	ft
07H9163	5002 (Model 010 and 500 only)	0.18	0.6
31H7960	5006 (Model 010 and 500 only)	0.6	1.9
07H8985	5010	1.0	3.3
32H1465	5025	2.5	8.2
88G6404	5050	5.0	16.4
32H1466	5100	10	32.8
88G6406	5250	25	82.0

7133 SSA Subsystem Cabling and Configurations

This section provides information on cabling for all the 7133 models to the PCI SSA Multi-Initiator/RAID EL Adapter.

If you order one or more 7133 Model 010 or 020 units as part of a new system, some cables are supplied with each 7133 unit. These cables might be connected to the unit, the adapter, or both. The cables supplied depends on the number of disk drives that are installed in the 7133 unit. Table 25 contains the SSA cables that are provided with a 7133 Model 010 or 020.

Number of Disk Drives	Number of Cables	Part Number	Feature Code	Length	
				m	ft
1 through 4	2	88G6404	5050	5.0	16.4
5 through 8	2	88G6404	5050	5.0	16.4
	1	07H8985	5010	1.0	3.3

Number of Disk Drives	Number of Cables	Part Number	Feature Code	Length	
				m	ft
9 through 12	2	88G6404	5050	5.0	16.4
	1	07H8985	5010	1.0	3.3
	1 (Model 010 only)	07H9163	5002	0.18	0.6
13 through 16	2	88G6404	5050	5.0	16.4
	2	07H8985	5010	1.0	3.3
	1 (Model 010 only)	07H9163	5002	0.18	0.6

Some of these cables will become surplus after an installation.

7131 SSA Subsystem Cabling and Configurations

This section provides information on cabling for 7131 (model 405) to the PCI SSA Multi-Initiator/RAID EL Adapter.

SSA Subsystem Cabling Considerations: “Base Configuration” on page 124 and “Loop Configuration Using Four 7131 Units” on page 124 show two configurations for SSA subsystems using 7131 units.

SSA Cables: This section gives details of the adapter and cables shown in the diagrams in the section.

Model 405

Note: All of the feature codes listed for the SSA Copper Cables (item **2** in Table 26) are feature codes of machine type 7131.

Item	Description	Part Number	Feature Code	Length	
				m	ft
1	PCI SSA Multi-Initiator/RAID RAID EL Adapter	96H9938	6215 feature of system unit	N/A	N/A
2	SSA Cable, 7131 to Adapter	07H8985	2895	1.0	3.3
		32H1465	2896	2.5	8.2
		88G6404	2897	5.0	16.4
		32H1466	2898	10	32.8
		88G6406	2899	25	82.0

Base Configuration: A 7131 unit base configuration includes two disk drive carrier assemblies. These carrier assemblies are installed in the lower two slots of the 7131 unit. The other three disk drive positions contain dummy disk drive carrier assemblies.

In the simplest configuration, the SSA loop is completed by connecting the two SSA connectors on the 7131 unit to one of the two pairs of connectors on an SSA adapter card in the using system, as shown in Figure 33.

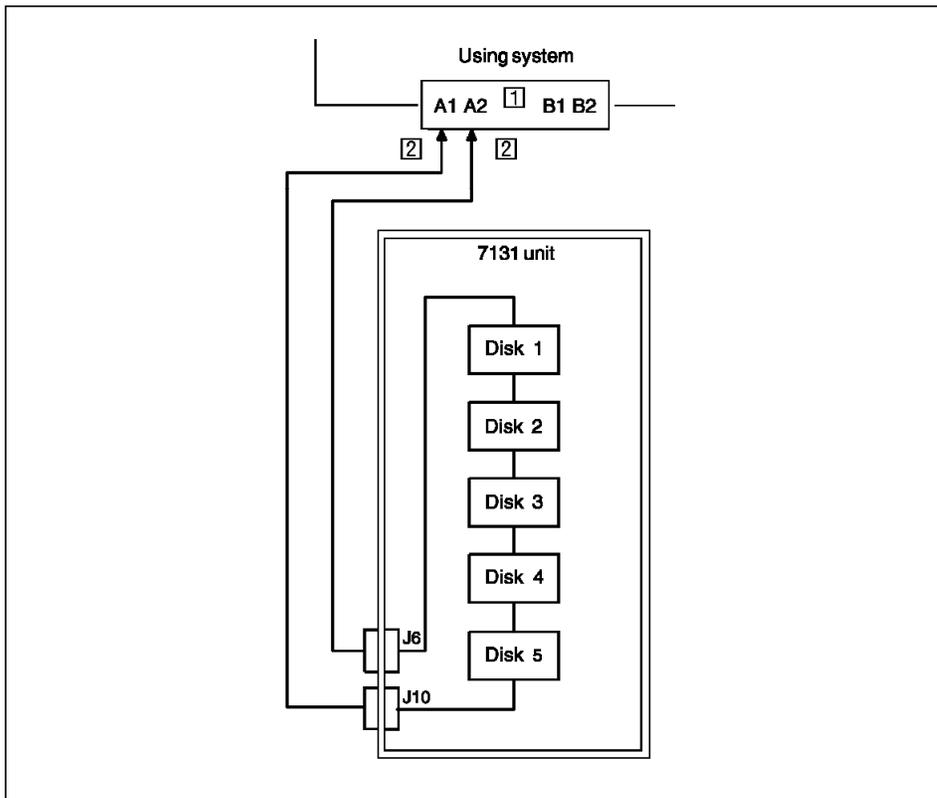


Figure 33. 7131 Base Configuration

Note: The cables shown in Figure 33 are described on page 123.

Loop Configuration Using Four 7131 Units: All twenty disk drive carrier assemblies in the four full 7131 units are connected in a single loop, as shown in Figure 34 on page 125.

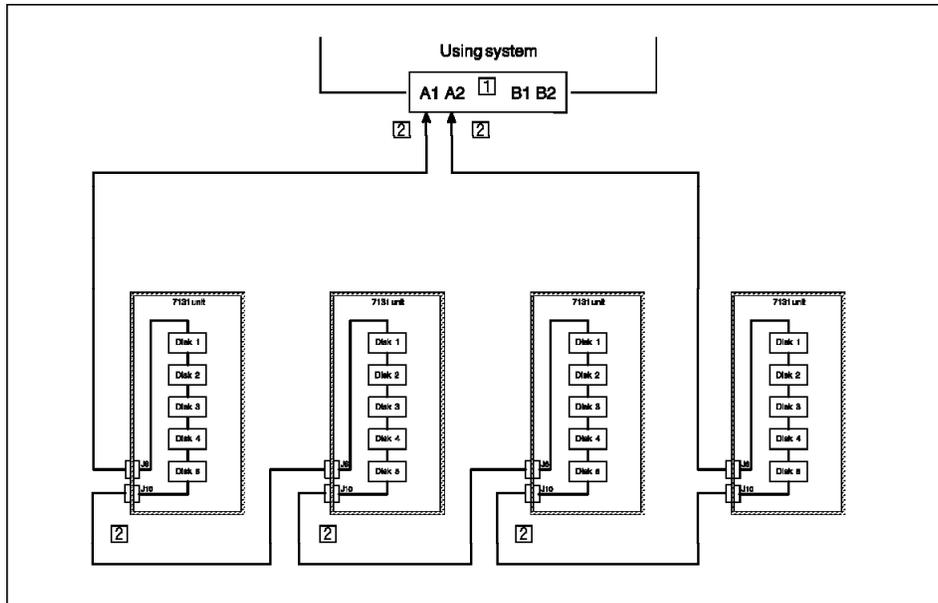


Figure 34. Loop Configuration with Four 7131

For a higher-availability configuration, a second SSA adapter can be included in the loops.

Note: The cables shown in Figure 34 are described on page 123.

Cabling SSA Adapters Using External Cables to Internal Devices

Review this section to connect the PCI SSA Multi-Initiator/RAID EL adapter's external ports to internal SSA disk drives.

Some system units have a pair of SSA bulkhead connectors at the rear of the system unit. The bulkhead connectors allow a pair of the SSA adapters external ports to be used to support internal SSA devices. This configuration is used when it is necessary to support both internal and external SSA disk drives on a single pair of SSA ports.

A pair of SSA ports, A1 and A2 or B1 and B2, can support a maximum of 48 SSA disk drives.

Internal SSA Backplanes: As shown in Figure 35 on page 126, the internal SSA backplanes can support up to 18 one-inch SSA drives. The backplanes are shown with six one-inch wide bays or SSA drives. When 1.6-inch SSA drives are used, only three drives can be installed (each drive takes two bays), and a dummy is required in each of the three unused bays to complete the loop.

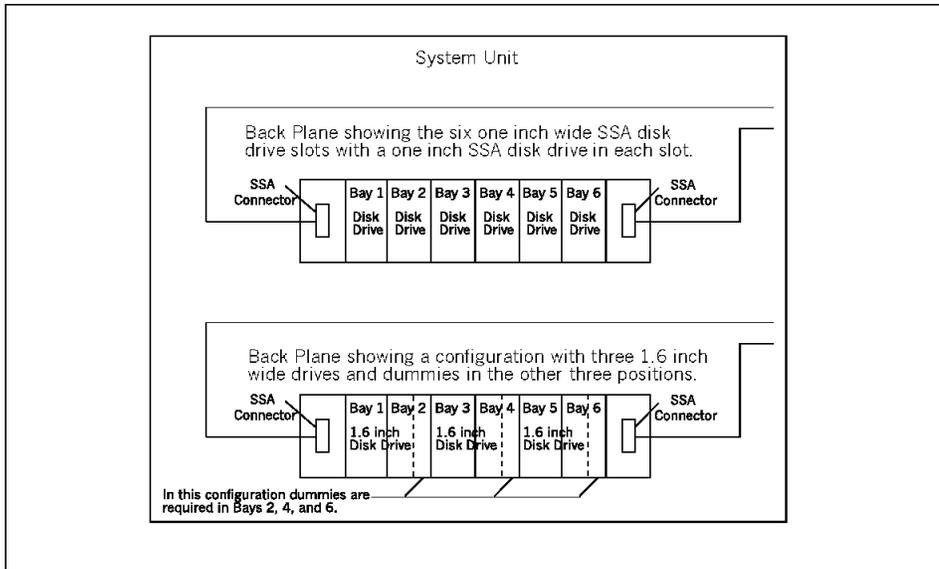


Figure 35. Internal SSA Backplanes

Simple SSA Loop to Internal SSA Drives: Figure 36 shows a simple loop configuration using two short (0.6 m or 2 ft. P/N 93H4340) external cables from the SSA adapters external ports to the bulkhead connectors.

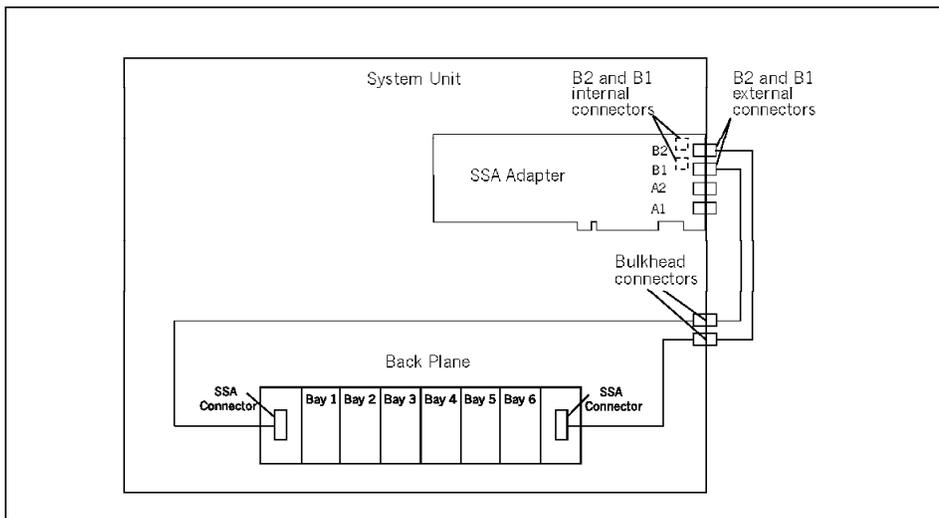


Figure 36. SSA Loop to Internal SSA Drives

Note:

- The cables required to cable up SSA subsystems are described on page 123.
- When cabling to external cable pair B1 and B2, insure that the internal cable pair B1 and B2 are not being used. These port pairs are mutually exclusive. They cannot be used at the same time.

SSA Loop to Internal and External SSA Drives: Figure 37 shows a configuration using both internal and external SSA disk drives to form the SSA loop. The external SSA devices can be installed in any combination of 7131 model 405 or 7133 external SSA subsystem storage units.

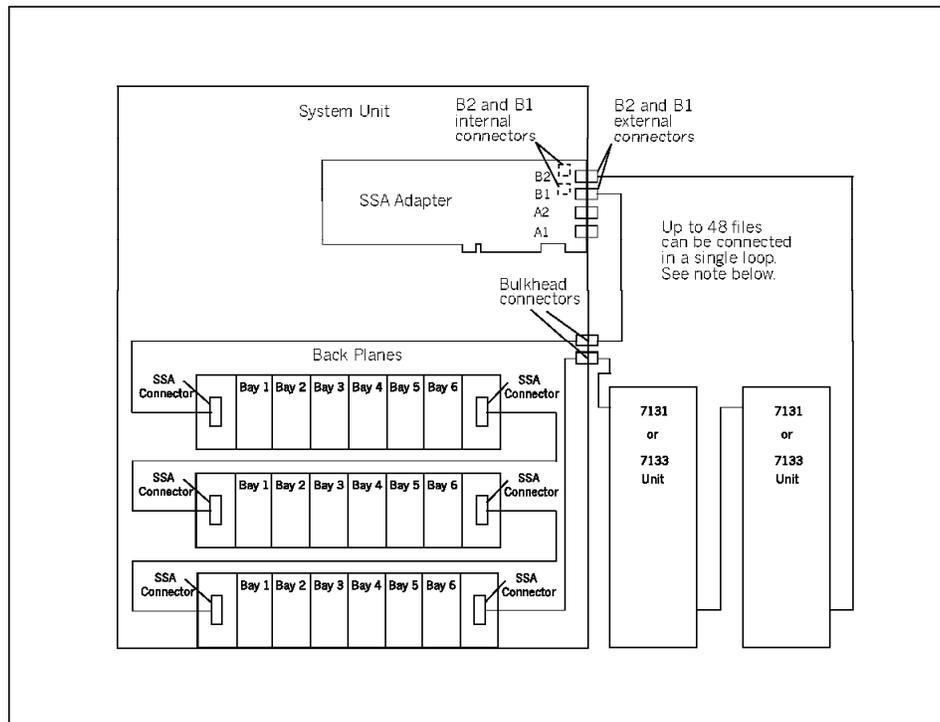


Figure 37. Loop to Internal and External SSA Drives

Note:

- A pair of SSA ports can support up to 48 disk drives. Two to 18 can be internal and the remainder can be external. Two 7131/7133 disk drive units are shown in the illustration above. Several can be cabled together until a total of 48 disk drives are in a single SSA loop.

- The cables required to cable up SSA subsystems are described on page 123.
- When cabling to external cable pair B1 and B2, insure that the internal cable pair B1 and B2 are not being used. These port pairs are mutually exclusive; that is, they cannot be used at the same time.
- For information on cabling the internal SSA backplanes to the internal connectors on the SSA adapter, see your System Unit User's Guide.

4.3 Hot Swap Subsystem

The hot swap subsystem allows the addition or removal of disk drives without powering-down the system. The hot swap subsystem also allows applications that do not use the disk that must be replaced or removed to keep running.

The hot swap subsystem uses special, hot-swappable disk carriers that allow you to easily plug and unplug disks in the system. The connectors, however, have a limited life of insertions and removals. If the user expects to insert and remove a drive daily, it may begin to fail intermittently resulting in a hot swap carrier, backplane, or drive having to be replaced. This subsystem, therefore, is not intended to be used as a removable disk subsystem.

This section describes the F40, F50, and H50 hot swap subsystem and the AIX features which can take advantage of this capability.

4.3.1 Components

The hot swap subsystem hardware is divided into three major components; the disk, carrier, and backplane. These components are shown in Figure 38 on page 129 and are described in the text that follows.

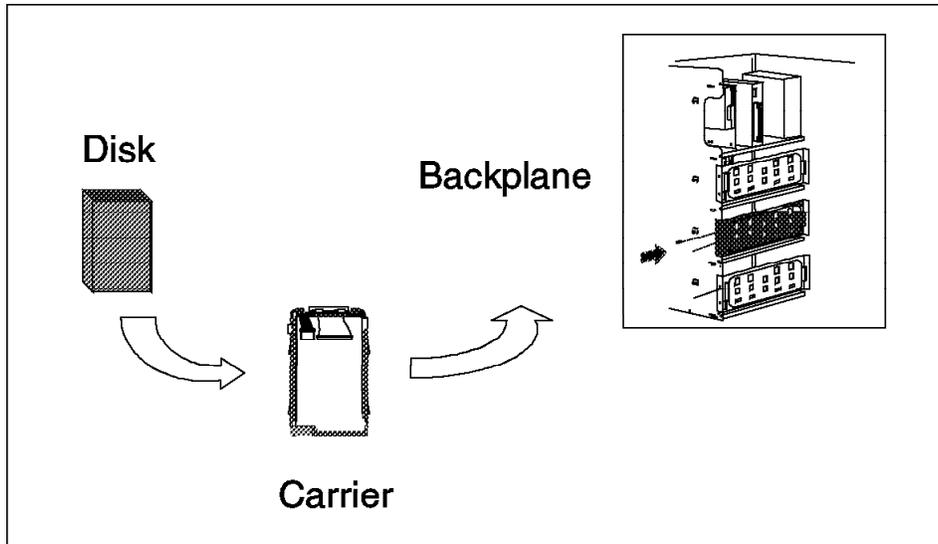


Figure 38. Hot Swap Components

- Backplane

The backplane card provides connectors and electronics support for up to six half-height or three full-height 3.5 inch 8-bit or 16-bit disk drives. The backplane card includes power management and hot-plug sense logic. The backplanes with all the necessary hardware are called *6-Packs*.

- Hot Swap Disk Drive Carrier

The hot swap SCSI disk drive carrier is a tray with built-in SCSI connectors that provide the interface between an industry standard SCSI 16-bit or 8-bit disk drive and the 80-pin backplane connector.

The carrier includes two LEDs. One of the LEDs is used to indicate the status of the disk drive, and the other one is the power LED indicator for the carrier.

Carriers are available for two drive types: narrow (half-height) and wide (full-height). The narrow carrier has a 50-pin drive connector. The wide carrier has a 68-pin drive connector. Both narrow and wide carriers have the same connector at the rear of the hot-swappable carrier. The different carriers, therefore, can be installed into the same backplane card.

The SSA Hot Swap Disk Drive Carrier is also available either for narrow and wide disks.

- Disk Drive

The following IBM disk drives are supported in this environment:

- 2.2 GB SCSI-2 Fast/Wide disk drive (F40 only)
- 4.5 GB Ultra SCSI Hot Swap disk drive
- 9.1 GB Ultra SCSI Hot Swap disk drive
- 9.1 GB 1" Ultra SCSI Hot Swap disk drive
- 4.5 GB SSA Hot Swap disk drive
- 9.1 GB SSA Hot Swap disk drive

The hot swap subsystem additionally requires a power source and a SCSI-2 Fast/Wide SE or SSA adapter. In the standard configuration, the first backplane is connected to the first power supply and to the integrated SCSI-2 Fast/Wide adapter. Additional backplanes will be connected to the second SCSI-2 Fast/Wide PCI adapter and obtain their power from the second power supply.

Instead of the standard SCSI-2 Fast/Wide adapter, an optional SSA, SCSI RAID, or Ultra SCSI adapter can be chosen. For that, the appropriate adapters and cables must be configured and a SCSI boot disk must be available either in the first 6-Pack or in the media bay. For more information, refer to the specific sections of the different models in Chapter 3, "Detailed Product Descriptions" on page 41.

Models F40 and F50 6-Packs

The Models F40 and F50 can accommodate up to three 6-Packs, each of which can accommodate a maximum of six hot-swappable disks. The first SCSI hot swap 6-Pack is installed in bank C (as shown in Figure 39 on page 131) and is supplied with the F50 as standard and as an optional feature on the F40 (F40 comes standard with a non-hot swap 6-Pack).

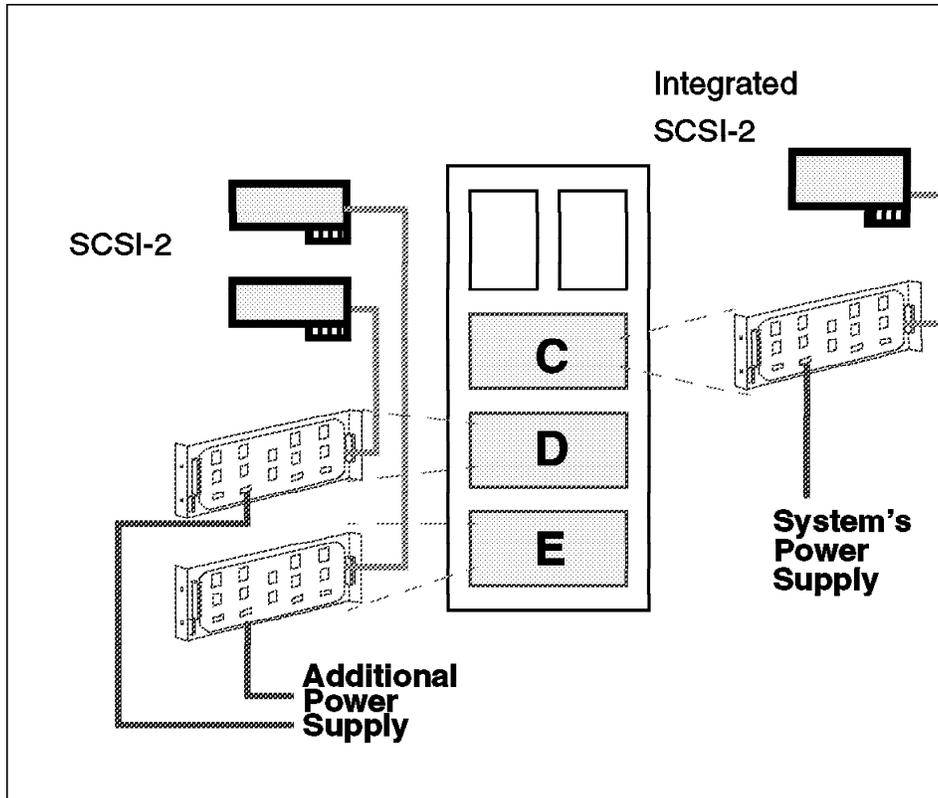


Figure 39. Installing Hot Swap

The second and the third 6-Packs are installed in banks D and E, respectively. The 6-Packs contain all of the necessary hardware to completely enable up to six hot swap disks installed in the machine: the backplane, the power cable, the SCSI cable, and the mounting hardware.

An optional feature is the use of SSA disks in SSA 6-Packs. At least one optional SSA adapter with connection cables must be ordered. At least two SSA disks must be configured for each SSA 6-Pack. Since booting only from SCSI disks is supported, the first 6-Pack must be SCSI, or a SCSI disk in media bay must be available.

The Models F40 and F50 default configuration is to control bank C's hot swap disks and the removable media devices (CD-ROM and so forth) with the first integrated SCSI-2 Fast/Wide controller. For improved performance, the second SCSI-2 SE Fast/Wide adapter can be connected to the first 6-Pack to provide an independent SCSI bus for Bank C's disk drives. Feature # 2443 provides the necessary cable to control bank C's disks from the second SCSI adapter.

Model H50 SCSI 6-Packs

As shown in Figure 40, the H50 features only two SCSI or SSA 6-Packs. The lack of a third 6-Pack is the only difference between the H50 and the F50 regarding the storage subsystem.

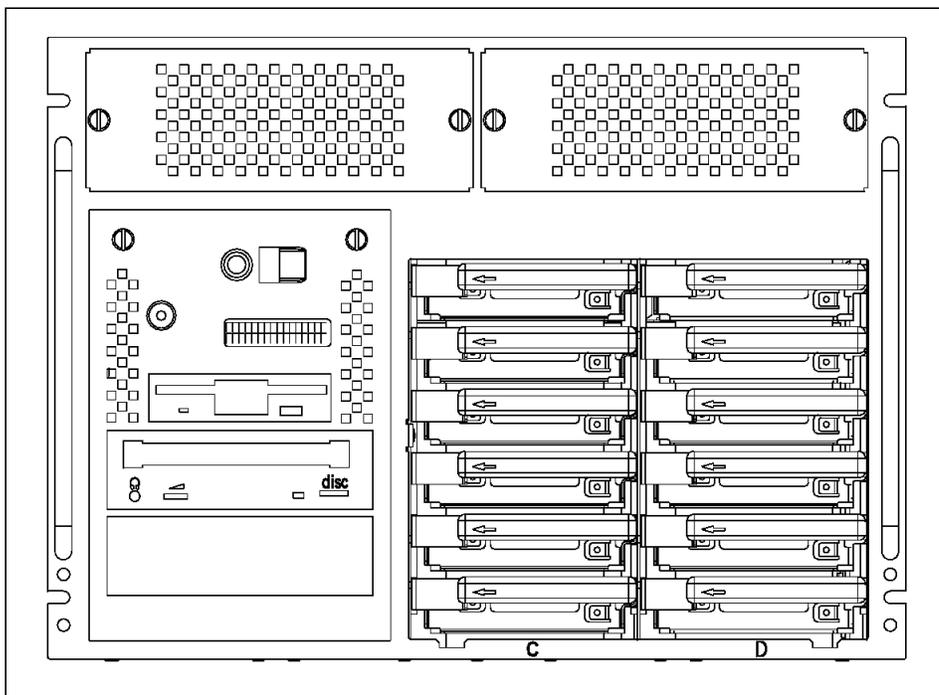


Figure 40. RS/6000 Model H50 Bays

Hot Swap Bays and 6-Packs

For a disk drive to be installed in one of the hot-swappable bays, the disk must use a hot-swappable carrier, a 6-Pack.

Each 6-Pack contains six bays and can accommodate three full-height disk drives/carriers or six half-height disk drives/carriers. A combination of half- and full-height disk drives/carriers is possible.

SCSI hot swap bays are numbered from 1 to 6 from left to right. SCSI IDs ranging from 8 to 13 are assigned automatically from left to right. When disks are installed in the carriers, their SCSI ID jumpers have to be removed. The hot swap bay SCSI ID numbering is shown in Figure 41 on page 133.

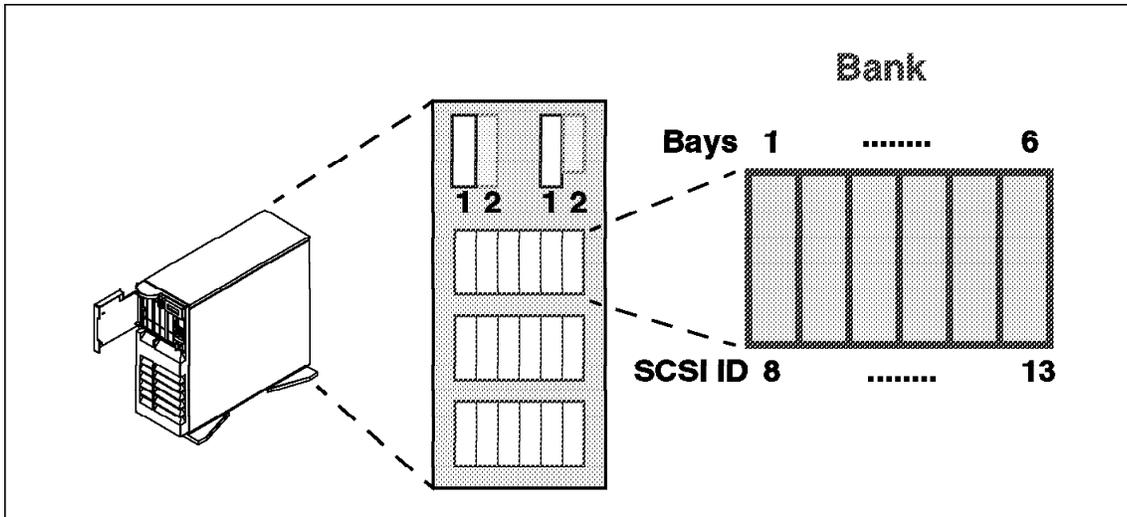


Figure 41. Front Bays

4.3.2 Models F40 and F50 Hot Swap Disk Installation

The requirements for installing hot swap disk drives are not the same for all three banks. A general requirement is that each hot swap bank requires a separate SCSI-2 adapter. In the case of the first and second bank, bank C and D in Figure 39 on page 131, they use the two integrated SCSI-2 controllers. The use of the third bank requires an additional SCSI-2 Fast/Wide Adapter. Bank D and E require the installation of one additional backplane each. The F40 also requires one additional power supply when the second and third backplanes are used.

Accessing the Models F40's and F50's Hot Swap Banks

To access the hot swap banks, you must first remove the front cover of your machine as shown in Figure 42 on page 134. To do this, refer to the "Installing Options" Chapter in the F40 or F50 User's Guide; then look at the section titled "Removing the Front Covers Only." In the same order as this section describes, you should:

1. Open and remove the door.
2. Release the front cover latch.
3. Unlock the keylock.
4. Remove the front cover.

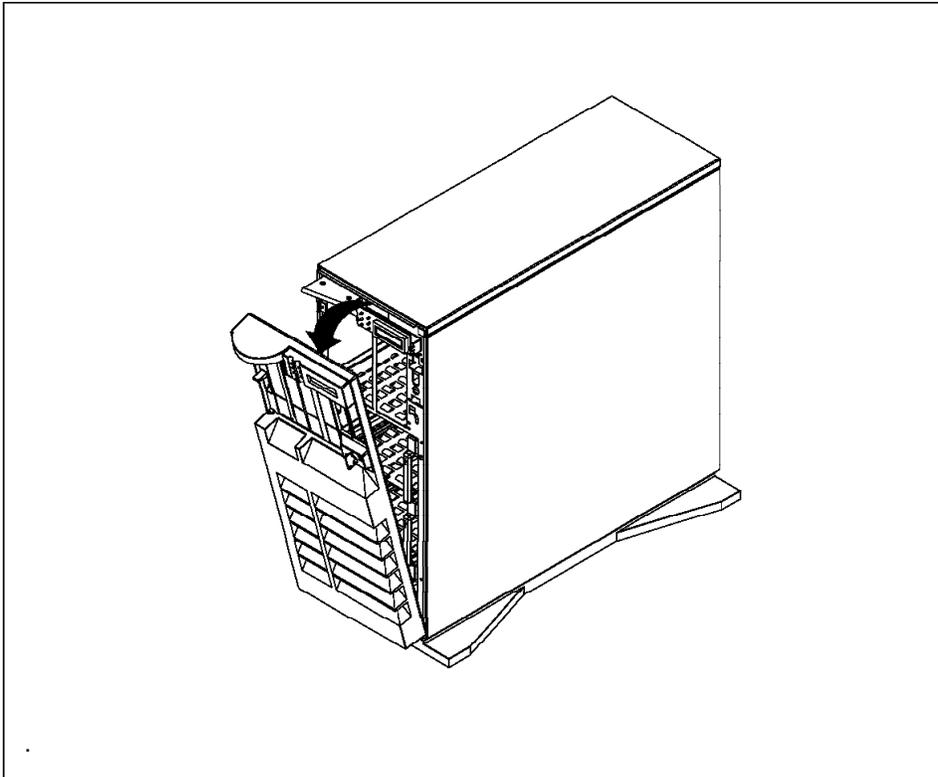


Figure 42. Removing the Front Panel

This will provide you access to the hot swap banks. If you have only one backplane installed, you will be able to add or remove up to six disks.

Installing SCSI Drives in Bank C

The backplane for bank C comes installed in the system. The SCSI IDs are automatically assigned by the backplane for each bank. Therefore, no jumpers are required on the drives. If your disk drive has SCSI ID jumpers set, you must remove them. Refer to the disk drive booklet provided with the system or with your disk drive. Attaching SCSI address jumpers on the drive rather than using the automatic ID supplied by the backplane is not advised.

Each drive you plan to install in bank C must have the hot swap carrier attached. It is recommended that you install the hard disk in the next unused position.

SCSI disk drives and SSA disk drives are not interchangeable.

Installing SCSI Drives in Bank D

To install drives in bank D, you must first install a backplane in the back of the bank. The installation procedure is as follows:

1. Install a backplane in the bank.
2. Connect the SCSI adapter to the backplane.
3. Connect the (additional, in case of F40) power supply to the backplane.
4. Install the hot swap drive as you would for bank C.

Installing SCSI Drives in Bank E

One additional SCSI-2 adapter must be installed in a PCI slot, then the installation of drives in bank E is exactly the same as the procedure used for installing drives in bank D.

Installing SSA Drives

SSA disks can be used in two configurations:

- SSA disks in bank D and E, bank C remains SCSI
- SSA disks in all three banks.

The second option requires a SCSI boot disk in the media bay.

An SSA adapter must be installed first; then the first SSA backplane can be installed and connected to the SSA adapter either with an internal cable or with an external bulkhead cable. Use a daisy chain cable for connecting the second and third SSA backplanes. Figure 43 on page 136 shows how to cable the SSA loop.

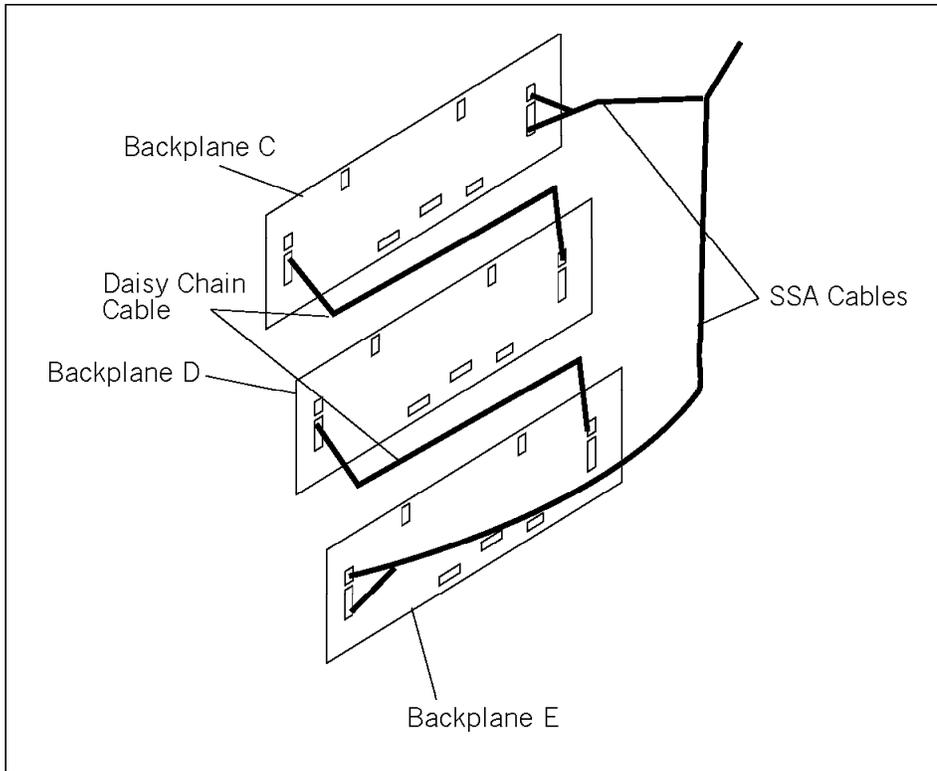


Figure 43. SSA Cabling

A minimum of two SSA disk drives are required when populating individual banks C, D or E. They must be located in bays 1 and 5 prior to installing additional disk drives in the next unused position of banks C, D or E.

SSA backplanes must be fully populated. Each backplane connector must have a disk drive or a blank disk drive module installed at all times. If you are installing a disk drive in a bay that was not previously utilized, you must remove the blank carrier prior to installing the disk drive.

Each SSA 6-Pack backplane feature will come populated with the necessary SSA blank disk drive modules based on the order configuration.

4.3.3 Model H50 Hot Swap Disk Installation

The H50 provides space for two hot swap banks. The first, bank C, comes standard as a SCSI hot swap bank; bank D is optional. The two integrated SCSI-2 Fast/Wide adapters are sufficient to run the two banks. A third SCSI adapter can be added for increased performance. Internal SCSI RAID and SSA are also available as optional features.

Accessing the Model H50's Hot Swap Banks

To access the hot swap banks, you must first open the front door of the rack unit and then loosen the screws of the disk drive covers as shown in Figure 44. To do this, refer to the "Installing Options" Chapter in the H50 User's Guide.

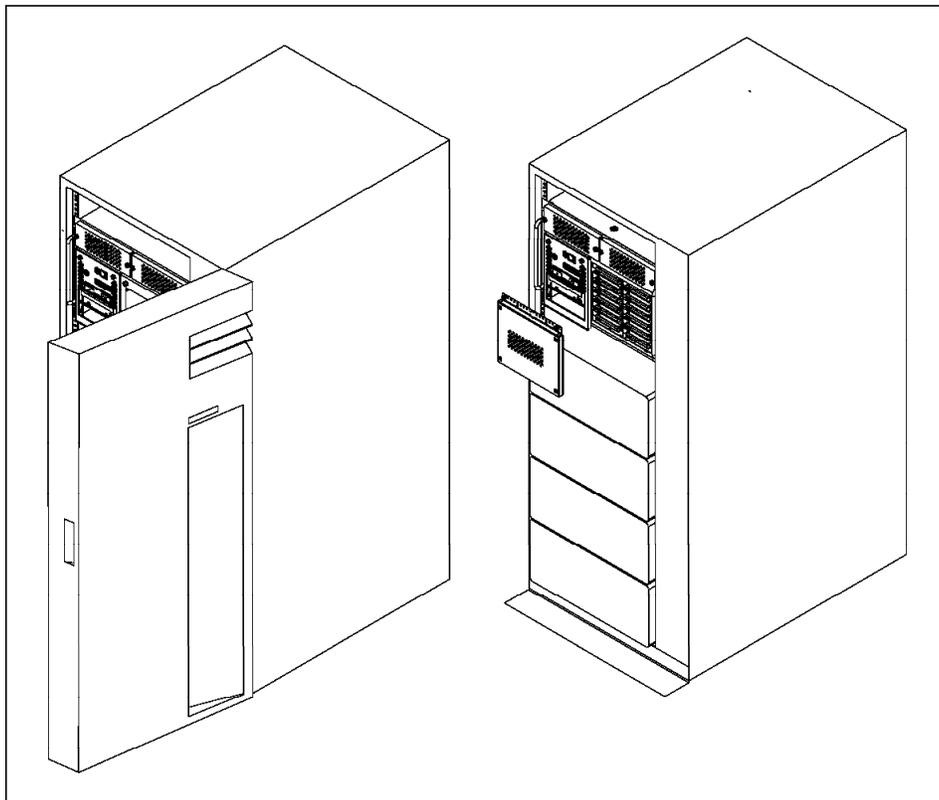


Figure 44. Accessing the Hot Swap Bays

This will provide you access to the hot swap banks. If you only have one backplane installed, you will be able to add or remove up to six disks.

Installing SCSI Disk Drives in Bank C and D

The Model H50 uses the same banks as the F40 and F50. Refer to "Installing SCSI Drives in Bank C" on page 134 for more information.

Installing SSA Drives

The Model H50 uses the same SSA configuration as the F40 and F50, but only two banks are provided.

Refer to “Installing SSA Drives” for more information. Figure 45 shows how to remove a hot swap disk drive.

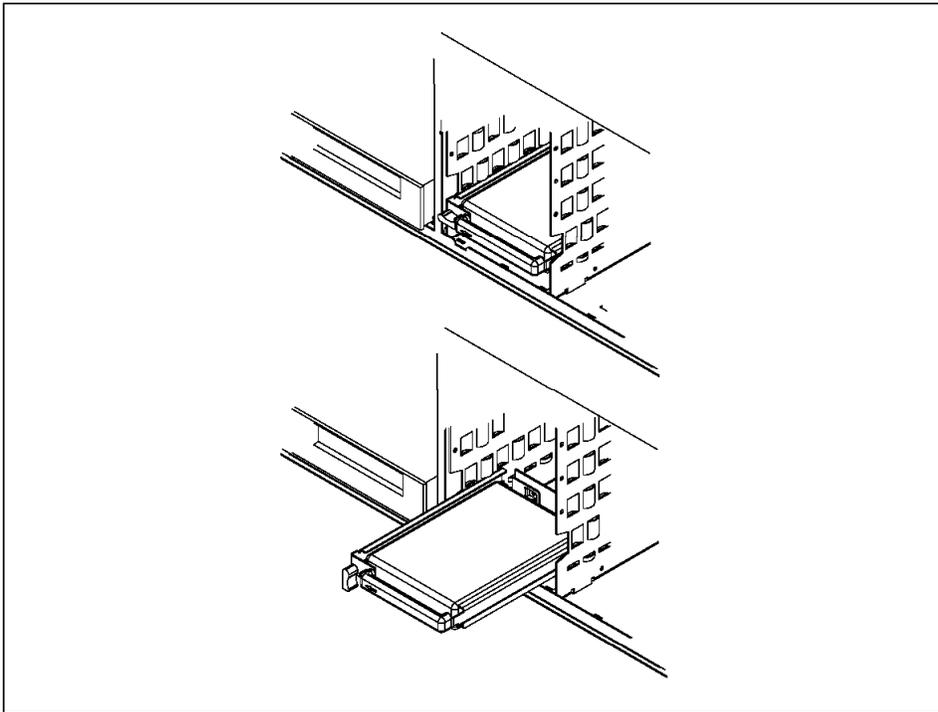


Figure 45. Hot Swap Disk Drives

4.3.4 SCSI Disk Carriers

The SCSI disk carriers have two lights which enable you to determine the status of the drive. Figure 46 on page 139 shows a diagram of the lights and the spin down button, and Table 27 on page 139 explains their meanings.

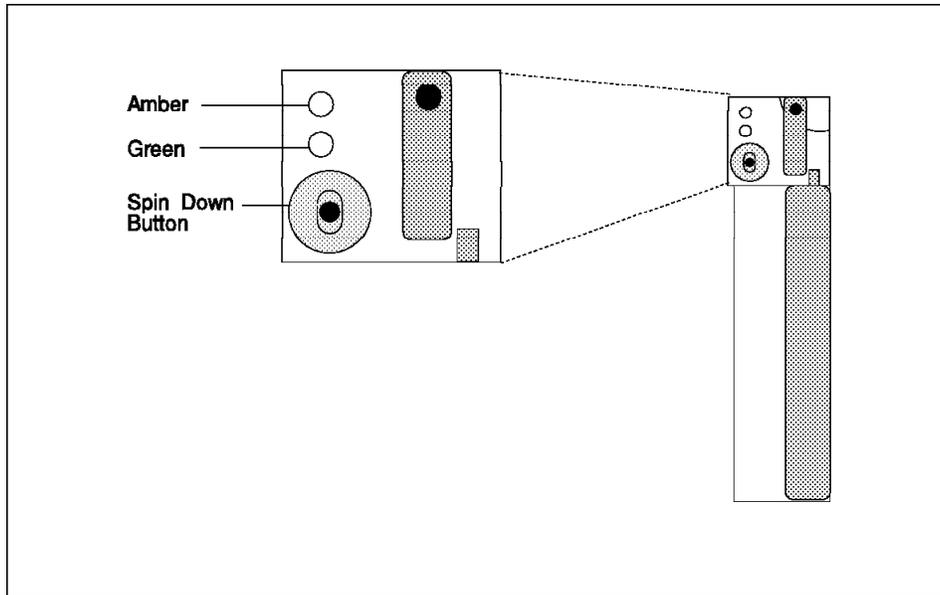


Figure 46. Disk Drive Light

LED or Button	Status	Definition
Amber	On	Drive spinning
	Off	Drive not spinning
Green	On	Power On
	Off	Power Off
	Blinking	Power Off/Drive identify
Spin down	Depressed	Spin down drive and remove power

During normal operation of AIX, both of the lights on the disk carrier will be on. If a drive is removed from the ODM, for example, by using:

```
rmdev -d -l hdiskx
```

the drive will spin down, and the amber light will go out.

4.3.5 SSA Disk Carriers

The SSA disk carriers have three LEDs showing the status of the drive. Figure 47 on page 140 shows the location of the three lights, and Table 28 on page 140 shows their meanings.

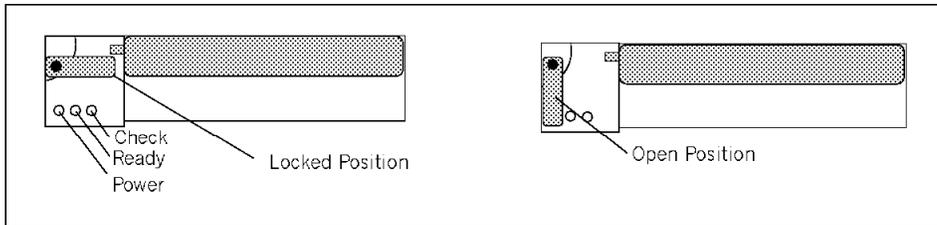


Figure 47. SSA Hot Swap Disks

LED	Status	Definition
Power	On	Power On
	Off	Power Off
Ready	On	Both SSA connections are good and drive is ready
	Blinking	Only one SSA connection is good
	Flickering	Drive is executing a command
Check	On	Disk drive failure
		Self-test running
		Drive in service mode / identify mode
	Blinking	Drive in service mode

4.3.6 Using the Hot Swap Subsystem

Having the hot swap subsystem is very useful. As previously mentioned, it allows you to add, remove, or replace disk drives without shutting down and powering-off your system. There are, however, three ways that you can use the hot swap subsystem.

The first way you can use the hot swap subsystem is to add and remove disk drives from the system. You can buy and add new drives with ease, and these can be quickly configured into your system without affecting the users.

The second way is for replacement usage. You may only have one system with the hot swap capability and will only remove drives if they become faulty. In this scenario, you can replace the drive without powering-off and shutting down the system. To enable the removal of the drive, tell AIX that the disk is to be removed from the system. If the data on the disk can still be accessed, then migrate the contents of the disk to another one or back up the data on the disk. The disk can then be replaced and the data restored.

The third usage of the hot swap subsystem is mirroring, effectively giving the subsystem RAID 0. The data is highly available since it is duplicated on two drives.

Adding a New Drive to a Live System

When you install a new drive in the hot swap subsystem, the amber light in the carrier will flash and then go out. This indicates that the drive has been identified and is not spinning. At this point, AIX does not know that the drive is present in the system. To tell AIX that you have added a new drive, you can use SMIT or type:

```
cfgmgr
```

from the command line. Running this command will cause AIX to find the drive and spin it up. You will see the amber light come on and hear the sound of the drive spinning up. Using the `lspv` command, you will be able to see your new drive added. If the drive contains a volume group, you access it from your system by running:

```
importvg -y VGname hdiskx
```

Otherwise, you can add the disk to an existing volume group by using `smit extendvg` or create a new volume group by using `smit mkvg`.

Removing a Drive from a Live System

To be able to remove a hot swap drive from the system without causing problems, you will have to tell AIX that you are removing the drive.

Removing a Disk from an Existing Volume Group: If you want to physically remove the disk and it belongs to an existing volume group, you would either remove the logical volumes which are present on the disk (you can determine which logical volumes are present on a disk by using the `lspv -l hdiskx` command), or migrate the physical partitions from the disk to other disks in the same volume group. To remove a disk from an existing volume group, you can use the following procedure:

1. Unmount all the file systems on the disk if you are removing the logical volumes on the disk.
2. Remove all data from the drive by either removing the logical volumes or by migrating the partitions on the disk to another disk in the same volume group (`smit migratepv`). If you are removing the logical volumes, you may wish to back up the data prior to removal.
3. Remove the drive from the volume group:

```
reducevg VGname hdiskx
```

4. Remove the device from the ODM:

```
rmdev -d -l hdiskx
```

When you run the `rmdev` command, the amber light on the drive will switch off. If you run `lspv`, you will see that the disk is no longer defined on the system.

5. Physically remove the drive from the system.

Removing a Drive with Its Own Volume Group: To remove a drive which has its own volume group, you can use the following procedure:

1. Back up any data that you require from the volume group.
2. Unmount all the file systems on the disk.
3. Varyoff the volume group by issuing:

```
varyoffvg VGname
```

4. Export the volume group by issuing:
5. Remove the device from the ODM by issuing:

```
rmdev -d -l hdiskx
```

6. Physically remove the device from the system.

Replacing a Previously Defined Drive

If you add a disk drive which was already configured to the system and was removed using the procedures described above, then you can simply add the new drive as described in “Adding a New Drive to a Live System” on page 141.

Replacing a Previously Defined Drive into the Same Bay: If a drive was physically removed without first being logically removed from the operating system, then AIX may have problems. If there were no writes to the disk after the removal of the disk, then further action is not required. If a write occurred after or during the removal of the drive and the drive has been re-added, then you should perform the following:

1. Unmount all file systems on the disk.
2. `fsck -y` file systems on the disk.
3. Remount the file systems on the disk.

Replacing a Previously Defined Drive into a Different Bay: If you place a drive into a bay different than the one from which it was removed (the one configured to the system by running the `cfgmgr` command) and you did not remove the device from the ODM before physically removing the device, you

will have to clear up the ODM because there will be a duplicate entry for the drive.

The following shows a scenario where there were three disks in the system and two volume groups on two separate disks. The disk belonging to testvg was physically removed without telling AIX.

Running the lspv command shows the three disks:

```
lspv
hdisk0      00000000a641877c   rootvg
hdisk1      00000000b0a645b0   testvg
hdisk2      00000000a6274604   None
```

hdisk1 was removed from the system before telling AIX. The disk was replaced into a different bay than it was removed from. The cfgmgr command was run to configure the disk back into the system, and now lspv shows:

```
lspv
hdisk0      00000000a641877c   rootvg
hdisk1      00000000b0a645b0   testvg
hdisk2      00000000a6274604   None
hdisk3      00000000b0a645b0   testvg
```

There are now two entries for the testvg volume group, both with the same physical volume identifier. This is incorrect and can be cleared up by performing the following:

1. Unmount all file systems in the testvg volume group.
2. varyoffvg testvg
3. exportvg testvg

At this point, you may get the following error which you can ignore:

```
0516-024 /usr/sbin/lqueryvg: Unable to open physical volume.
      Either PV was not configured or could not be opened. Run
      diagnostics.
```

Running lspv now shows:

```
lspv
hdisk0      00000000a641877c   rootvg
hdisk1      00000000b0a645b0   None
hdisk2      00000000a6274604   None
hdisk3      00000000b0a645b0   None
```

4. rmdev -d -l hdisk1
5. rmdev -d -l hdisk3

6. `cfgmgr`

Running `lspv` now shows the correct disks:

```
lspv
hdisk0      00000000a641877c   rootvg
hdisk1      00000000b0a645b0   None
hdisk2      00000000a6274604   None
```

7. `importvg -y testvg hdisk1`

8. Mount all the file systems. You may have to run `fsck -y` on the file systems first if they were written to while the disk was removed.

Replacing Mirrored Disks

In this section an outline of mirroring is given. Mirroring takes maximum advantage of the hot swap subsystem.

The hot swap subsystem means that AIX has to be explicitly told about the removal and addition of disks. In a normal AIX environment the system would have been shut down and powered-off. The system would then recognize the removal or addition of disks.

Removing and Adding a Mirrored Disk: If you want to remove a disk which is a mirror of other disks in its volume group, you can either remove the logical volume's copies which are on the disk and follow the procedure outlined in "Removing a Disk from an Existing Volume Group" on page 141, or you can remove the disk without telling AIX. If you choose to remove the disk without telling AIX, the volume group will stay on-line providing that quorum has been maintained (more than 50% of the disks in the volume group are still accessible after removing the disk) or if quorum checking has been turned off. When you re-add the disk, perform the following procedure:

1. Unmount all the file systems which are mirrored on the disk.
2. Change the state of the disk in the volume group to active:

```
chpv -v -r hdiskx
chpv -v -a hdiskx
```
3. Synchronize all the partitions on the disk from their mirrors:

```
syncvg -v VGname   or   syncvg -p hdiskx
```
4. Remount all the file systems which are mirrored on the disk.

This will ensure that AIX correctly knows about the disk being re-added and that all the partitions are correctly synchronized.

If the disk is part of the root volume group and the file systems which are on the disk cannot be unmounted, then you can either wait for a reboot of the

system or remove the disk from the volume group, and then re-add the disk by performing the following:

1. Remove all the mirrored copies from the mirrored disk (for example, if we had a logical volume called lv00, there were a total of two copies of the logical volume, and the disk which was removed was hdiskx):

```
mlvcopy lv00 1 hdiskx
```

2. Reduce the rootvg to remove the disk:

```
reducevg rootvg hdiskx
```

3. Remove the device from the ODM:

```
rmdev -d -l hdiskx
```

4. Insert the new drive.

5. Configure it to the system:

```
cfgmgr
```

6. Extend the root volume group to this disk:

```
extendvg rootvg hdiskx
```

7. Re-create the mirrors on the disk, for example:

```
mlvcopy lv00 2 hdiskx
```

Chapter 5. Adapters, Overview and Configuration

In this chapter, the major communication and graphics adapters are introduced. Following the introduction, a discussion of adapter configuration for ISA and PCI bus cards is given.

5.1 Communication Adapters

In this section, an introduction to the various communications adapters is given, including:

- Token ring
- Ethernet
- X.25
- ISDN
- FDDI
- ATM
- Multiple port
- Miscellaneous

5.1.1 Token-Ring Adapters

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

PCI Auto LANStreamer Token-Ring Adapter (# 2979)

The Auto LANStreamer Token-Ring PCI Adapter is designed to allow a PCI-based RS/6000 server to attach to 4 Mbps or 16 Mbps token-ring local area networks (LANs). The adapter automatically selects the correct token-ring speed (4-16 Mbps). It is compatible with all IBM PS/2 token-ring adapters; no new cables or network components are required.

The adapter has one connector, an RJ-45. The RJ-45 connector is used to attach to unshielded twisted pair (UTP) cabling. A 10-inch conversion cable is included with the adapter to attach to shielded twisted pair (STP) cabling.

The following are the main characteristics of this adapter:

- 4 Mbps or 16 Mbps speed (automatically set)
- Compatible with IEEE 802.5 specifications
- Full duplex enabled
- Occupies one PCI slot
- Maximum number of adapters: 6

This adapter will be replaced by feature 2920 after July 24, 1998. Check for availability on the Model E30.

IBM PCI Token-Ring Adapter (# 2920)

The IBM PCI Token-Ring Adapter is designed with high performance in mind. It is a 32-bit bus master adapter designed to operate on selected RS/6000 models that incorporate the PCI bus interface. Check for compatibility on the Model E30. 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 15 Mbps or 4 Mbps over UTP or STP
- Automatic ring speed detection
- External status LEDs providing visual indication of adapter and ring status
- Full duplex LAN operation
- Supports system book over token ring

5.1.2 Ethernet Adapters

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

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.

PCI 3Com 10/100 Mbps Fast Ethernet XL (# 2986)

The 3Com 10/100 Mbps Fast Ethernet XL adapter provides 10 Mbps or 100 Mbps connections to 10BaseT or 100BaseTX Ethernet LANs, including the ability to sense the hub port speed automatically. This adapter supports full duplex (FDX) operation at either speed.

Features include:

- Full 802.3u Fast Ethernet standard compliance
- Configurable transmit and receive buffering up to 8 KB
- Full duplex capability

- AIX SMIT and diagnostic routine support
- Remote boot capability on all RS/6000 models.

PCI IBM 10 Mbps Ethernet Adapters

Two types of PCI 10 Mbps Ethernet adapters are supported depending on the different kind of connections you need for your Local Area Network.

- PCI Ethernet BNC/RJ-45 Adapter (# 2985)
It supports connections to 10Base2 networks using BNC connectors and to 10BaseT networks using RJ-45 connectors.
- PCI Ethernet AUI/RJ-45 Adapter (# 2987)
It contains a 15-pin DIX Ethernet connector, and it supports 10BaseT networks using RJ-45 connectors.

Both adapters provide 10 Mbps Ethernet connectivity. They adhere to the PCI Revision 2.0 and IEEE 802.3 standards. They feature low-cost, high performance 32-bit cards. Each adapter includes three LEDs to provide the status of the card's operation and each includes auto-sense of the media connection. The controller chip on the card contains integrated DMA buffer management for CPU and bus utilization and look-ahead packet processing which allows protocol analysis to begin before the end of the frame is received.

5.1.3 X.25 Adapters

In this section, information on the available X.25 adapters is given.

ISA X.25 Adapter (# 2961)

The X.25 Interface Co-Processor Adapter ISA bus provides for the attachment to synchronous data networks at speeds up to 64 Kbps.

It is important to remember that the IBM AIXLINK/X.25 LPP for AIX is a software requirement to provide connection to X.25 packet-switched networks. This adapter provides a single port connection to X.25 networks that will accommodate one of the three selectable features:

- EIA-232D/V.24 interface up to 19.2 Kbps
- V.35 interface up to 56/64 Kbps
- X.21 (not switched) up to 64 Kbps

Each feature is selected by the cable, which can be ordered in 3 m (9 foot) or 6 m (18 foot) length. These are the main characteristics for this adapter:

- Code resident on the adapter off-loads low-level X.25 tasks from the system processor

- 512 KB RAM for buffering
- Synchronous protocols can be supported
- Automatic recognition of the selected interface cable
- Capacity of concurrent support up to 250 virtual channels
- Requires one ISA slot
- Maximum of three adapters of this type can be installed

Installation and configuration of this adapter is not automatically made by AIX; proper operations must be followed by the system manager. The DIP switches on the adapter must be set correctly before configuring the adapter to the AIX operating system. 5.3.7, "X.25 Interface Co-Processor ISA Adapter (# 2961)" on page 197 explains how to install and configure this adapter.

5.1.4 ISDN Adapters

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

PCI Eicon ISDN DIVA PRO 2.0 S/T (# 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 PowerPC Systems provides basic rate 2B + D with 64 Kbps on each B channel. TCP/IP applications are supported.

5.1.5 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, # 2986, and # 2988 (SysKconnect FDDI, 3Com Ethernet, 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.

PCI SysKonnnect SK-Net FDDI-LP SAS (# 2741)

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

Features include:

- Occupy 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

Note

This adapter is supported on all PCI-capable RS/6000 systems, but does not support remote boot over the FDDI network at this time.

The AIX Device Driver diskette is shipped with the adapter. Limitation: A configuration using more than one SysKonnnect FDDI adapter (# 2741, # 2742, or # 2743), where any one of them is feature # 2743 (SysKonnnect SK-NET FDDI-UP SAS PCI), constitutes a Class A system.

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:

- Occupy 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 connection requirements
- Connect to fiber FDDI concentrator ports using fiber optic cable with an FDDI SC (low cost fiber) connector

Note

This adapter is supported on all PCI-capable RS/6000 systems, but does not support remote boot over the FDDI network at this time.

The AIX Device Driver diskette is shipped with the adapter.

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

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).

Note

This adapter is supported on all PCI-capable RS/6000 systems, but does not support remote boot over the FDDI network at this time.

The AIX Device Driver diskette is shipped with the adapter. It also offers high speed, low latency, low CPU and bus utilization, and occupies a single PCI bus slot.

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

Features include:

- Occupy 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 with 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 Connection requirements:
- Connect to Category-5 unshielded twisted pair FDDI concentrator ports using shielded RJ45 connector

5.1.6 ATM Adapters

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

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.

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)
- 3Com 10/100 Mbps PCI Fast EtherLink XL for PowerPC Systems (# 2986)
- SysKonnect SK-NET FDDI-LP SAS PCI for PowerPC Systems (# 2741)
- SysKonnect SK-NET FDDI-LP DAS PCI for PowerPC Systems (# 2742)
- SysKonnect SK-NET FDDI-UP SAS PCI for PowerPC Systems (# 2743)

PCI TURBOWAYS 25 ATM Adapter (# 2998)

The TURBOWAYS 25 ATM PCI Adapter is for direct access to ATM networks, the TURBOWAYS 25 ATM PCI Adapter provides dedicated 25 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

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

5.1.7 Multiple Port Adapters

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

ISA 4-Port Multiprotocol Communication Controller (# 2701)

This adapter supports the attachment to wide area networks at speeds up to 64 Kbps. This adapter consists of a base card and daughter card. The two cards are physically connected and require a single ISA bus card slot.

This adapter can support four different physical interfaces by using an IBM 4-Port Multiprotocol Interface Cable (# 2705). This cable is actually a breakout box featuring four ports that support EIA-232, EIA-422A, V.35, and X.21 physical interfaces.

The main characteristics for this adapter are:

- 80C186 @ 10 MHz processor
- 1 MB RAM for data buffering
- Synchronous protocols
- EIA232, EIA422A, V.35, and X.21 physical interfaces supported
- Requires one ISA slot
- Maximum of three adapters of this type can be installed

Installation and configuration of this adapter is not automatically made by AIX; proper operations must be followed by the system manager. DIP switches on the adapter must be set correctly before configuring the adapter to the AIX operating system. 5.3.6, "4-Port Multi-Protocol ISA Adapter (# 2701)" on page 191 explains how to install and configure this adapter.

ISA 8-Port Asynchronous Adapter (# 2931)

This adapter is used to connect up to eight EIA-232 asynchronous serial devices (terminals, modem, printers, and so on). This feature provides both the adapter card and an 8-port DB25 connector box, composed of eight asynchronous ports from a single I/O card slot.

The main characteristics for this adapter are:

- Data rates up to 115.2 Kbps per port
- Fully buffered transmit and receive data
- Single 78-pin connector for connection to 8-port DB25 connector box
- Supports the following interface signals: TxD, RxD, RTS, CTS, DSR, DCD, DTR and RI
- Maximum EIA-232 distance supported is 62 meters (200 feet)
- Requires one ISA slot
- Maximum of three adapters of this type can be installed

Installation and configuration of this adapter is not automatically made by AIX; proper operations must be followed by the system manager. DIP switches on the adapter must be set correctly before configuring the adapter to the AIX operating system. 5.3.4, "8-Port EIA-232 Asynchronous ISA Adapter (# 2931)" on page 183 explains how to install and configure this adapter.

ISA 8-Port Asynchronous Adapter (# 2932)

This product is similar to # 2931, but also includes RS422A attachment.

ISA 128-Port Asynchronous Adapter (# 2933)

This adapter provides for a high concentration of asynchronous lines from a single ISA bus slot. Two synchronous channels link the adapter to a maximum of eight 16-Port Remote Asynchronous Nodes (RAN) (feature # 8130 in USA and Canada only; feature # 8134 World Trade). Each channel can support up to four 16-Port RANs.

One 4.5 meter cable (# 8131) is available to connect the adapter to the first RAN on each channel. Additional RAN modules may be connected to each channel in a *daisy chain* configuration, using either the 4.5 meter (15 foot) cable (# 8131) or the 23 cm (9 inch) cable (# 8132). Customer-supplied cables may also be substituted.

These are the main characteristics for this adapter:

- Data rates up to 57.6 Kbps per port
- Support to 128 ports per adapter
- Supports the following interface signals: TxD, RxD, RTS, CTS, DSR, DCD, DTR, and RI
- Maximum EIA-232 distance supported is 62 meters (200 feet)
- Requires one ISA slot
- Maximum of three adapters of this type can be installed

Installation and configuration of this adapter is not automatically made by AIX; proper operations must be followed by the system manager. DIP switches on the adapter must be set correctly before configuring the adapter to the AIX operating system. 5.3.5, "128-Port EIA-232 Asynchronous ISA Adapter (# 2933)" on page 186 explains how to install and configure this adapter.

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.

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.

The new RANs being announced provide 16 EIA-232 (# 8137) or 16 RS-422 (# 8138) ports, yielding a total of 128-ports in a maximum configuration. Existing 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 the best results, keep the previous and new RANs on separate synchronous channels.

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).

5.1.8 Miscellaneous Communications Adapters

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

PCI Digital Trunk Quad Adapter (# 6309)

The Digital Trunk Quad Adapter attaches to a Single Digital Trunk Processor (# 9291) or a Multiple Digital Trunk Processor (# 9295). It requires one PCI slot and allows up to 96 T1 or 120 CEPT (a European telephone standard) telephone channels to be attached. A maximum of two adapters per RS/6000 are supported. IBM DirectTalk for AIX Version 2.1 (5755-B81) is a corequisite for this adapter.

5.2 Graphics Adapters

In this section, each of 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 several industry-leading 2D accelerators designed for any RS/6000 workstation that 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 recently announced graphics accelerators:

- MVP Dual Monitor
- GXT110P
- GXT120P
- GXT250P/255P
- GXT500P/550P
- GXT800P/800P with Texture Option
- GXT1000 internal

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

5.2.1 MVP Power Multi-Monitor Graphics Adapter (# 2837)

The MVP Power Multi-Monitor Graphics Accelerator supports two displays on a single adapter, thus conserving an expansion slot. This allows the viewing of multiple data sets simultaneously and 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 seamless between the two monitors for easy operation.

The MVP Power Multi-Monitor Graphics Accelerator is a full-length PCI card with a single SCSI-type output connector. An included “Y” cable provides adaptation to two standard 15-pin DIN connectors. The accelerator supports multiple display resolutions (from 640x480 to 1280x1024) that are software selectable. It satisfies the ISO 9241 standard of visual quality at 1280x1024 pixel resolution, refreshing the screen at up to 85 Hz. When used with compliant monitors, this standard provides improved viewing, reduced-flicker, minimized reflections, and sharper characters.

5.2.2 POWER GXT110P (# 2839)

The POWER GXT110P graphics accelerator provides exceptional value for entry-level 2D graphics on the RS/6000 Model E30 (as of May 15, 1998, the GXT110P is unavailable for several RS/6000 models, the main exception being the Model E30). The GXT110P has a single 8-bit frame buffer which allows the display of 256 simultaneous colors from a palette of 256,000 colors and one hardware color map.

The GXT110P is designed to support high-resolution monitors with display resolutions from 640x480 up to 1280x1024. The accelerator supports refresh rates from 60 Hz to 85 Hz (resolution dependent) and monitors that comply with Part 3 of the ISO 9241 ergonomic standards at a resolution of 1024x768 and a refresh rate of 75 Hz or 85 Hz.

The GXT110P accelerates drawing features such as pattern fill and rectangular and non-rectangular clipping for Xlib. The accelerator provides video playback support such as pixel scaling and color space conversion under Ultimeia Services (UMS).

The POWER GXT110P occupies one 32-bit PCI slot. Only one POWER GXT110P may be used in the Model F40 (7025).

The GXT110P does *not* support IBM’s Softgraphics for entry-level 3D graphics applications.

5.2.3 POWER GXT120P (# 2838)

The POWER GXT120P is a versatile, low-priced graphics accelerator for RS/6000 PCI bus Models F40, F50, and H50. It accelerates X Windows graphics applications, such as displaying business graphics, data base status, 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 1280x1024 8-bit pixels.

The GXT120P graphics accelerator is a half-length PCI card with a standard 15-pin DIN output connector. See accompanying table for multiple card combinations.

A maximum of four GXT120Ps can be added to a Model F40 and two to the Model F50 and H50.

5.2.4 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 Models E30 (limited to the GXT250P) and F40. The POWER GXT250P provides a high level of 8-bit 2D graphics performance, and 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 which 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 like

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.

5.2.5 POWER GXT500P (# 2854) and POWER GXT550P (# 2855)

The POWER GXT500P and POWER GXT550P graphics accelerators offer exceptional and affordable, mid-range 3D graphics for the PCI-based RS/6000 Model F40 workstation. Coupled with IBM's implementation of many of the industry's most popular APIs, these accelerators are 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 GXT500P and GXT550P further establish IBM as a leader in providing price/performance solutions for a wide range of demanding customer application requirements. The GXT500P and GXT550P combine 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 GXT500P and GXT550P provide hardware acceleration for the OpenGL, PHIGS, and IBM GL 3.2 3D APIs, and additionally support the PEX API through software. IBM is a leader in the industry by offering native support of OpenGL and PHIGS on the same accelerator. These adapters accelerate 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 GXT500P and GXT550P offer 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 GXT500P supports 8-bit and 12-bit double-buffered color. Both accelerators provide 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. The POWER GXT500P and GXT550P provide 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 GXT500P and the GXT550P graphics accelerators are each 32-bit PCI cards that plug into a single PCI slot. They support 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 Model F40 workstation, make the POWER GXT500P and GXT550P ideal solutions for demanding 3D application needs.

5.2.6 POWER GXT800P (# 2853) and with Texture Option (# 2859)

The POWER GXT800P graphics accelerator provides a new level of powerful and advanced 3D graphics for the PCI-based RS/6000 Models F40 and F50 workstations. The GXT800P offers an optional configuration that accelerates texture mapping in hardware, generating more realistic images at interactive speeds.

The GXT800P is particularly well suited for customers requiring the next step in performance above that provided by the GXT500 family of mid-range graphics accelerators. Coupled with IBM's implementation of many of the industry's most popular APIs, this accelerator is an excellent fit for today's most demanding 3D graphics applications in fields such as mechanical design and analysis, petroleum exploration and production, molecular modeling, and scientific research. The GXT800P further establishes IBM as a leader in providing advanced graphics solutions for a wide range of demanding customer application requirements.

The GXT800P utilizes a 5-way rendering engine that processes advanced 3D graphics in parallel, providing the throughput required to work with complex geometries at interactive speeds. When this is combined with the latest PowerPC microprocessors in the RS/6000 Models F40 and F50, the GXT800P 3D graphics accelerator delivers increasing performance as you upgrade to more powerful systems.

The optional texture map configuration of the POWER GXT800P provides native hardware acceleration for advanced texture mapping through OpenGL. This configuration supports MIP-mapped textures up to 512x512 texels in size, dramatically increasing the performance of OpenGL texture-mapped rendering and providing users with interactive viewing of realistically rendered 3D models. For larger texture maps, both configurations support up to 1024x1024 textures and a 64-bit accumulation buffer through software. The hardware texture mapping option helps geoscientists map the latest satellite imagery onto complex oil reservoir models, which greatly aids in the visual analysis of the reservoir.

The GXT800P provides native hardware acceleration for the OpenGL, PHIGS, and GL 3.2 3D APIs, and additionally, supports the PEX API through software. The GXT800P accelerates advanced 3D graphics functions such as Gouraud shading, antialiasing, hidden surface removal, depth-cueing/fogging, and transparency. This should enable your 3D applications to run more quickly and interactively.

The GXT800P offers a highly flexible frame buffer that can be dynamically configured to provide a broad set of color and feature options. When using OpenGL, PHIGS and GL 3.2, the GXT800P supports 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 which enhance the speed of a graphical user interface (GUI), 8-bit stencil buffers, and 24-bit Z-buffer for hidden surface removal operations. The POWER GXT800P 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 GXT800P graphics accelerator is packaged as a single 32-bit PCI card that fits inside supported RS/6000 workstations providing powerful 3D graphics in a compact desktop configuration. However, the POWER GXT800P attaches through a 32-bit or 64-bit PCI bus occupying three slots.

The GXT800P 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. The advanced features and powerful performance of the GXT800P, coupled with IBM's powerful RS/6000 Models F40 and F50 workstations, make this accelerator the ideal solution for your most demanding 3D application needs.

5.2.7 POWER GXT1000 Internal (# 7252)

The POWER GXT1000 graphics accelerators are Class III products that offer advanced graphics function and performance for the RS/6000 workstation family. They are one of the industry's first product offerings designed to provide hardware acceleration for multiple 3D APIs including OpenGL, PEX, PHIGS, and IBM GL 3.2. They also utilize a dedicated datapath for X Window System calls, providing exceptional 2D performance for a 3D accelerator.

This adapter is no longer available after July 24, 1998. Check your online sales manual for further restrictions.

Enhancements included in the release of PEX/PHIGS Version 4.1.x for AIX and OpenGL Version 4.1.x for AIX licensed program products provide customers with significant performance improvements for OpenGL and graPHIGS. GraPHIGS applications can run as much as 58 percent faster and OpenGL applications as much as 90 percent faster, through the combination of these product releases of AIX Version 4.1.x and the GXT1000 graphics accelerators. IBM continues to be the only vendor to provide the ability to concurrently support both of these important industry-standard APIs natively on the graphics accelerator on a single graphics workstation.

The GXT1000 accelerators have a highly configurable frame buffer that provides onboard support for double 8-bit, 12-bit, and 24-bit color buffers, a 24-bit z-buffer, alpha blending, accumulation, and texture mapping. A double-buffered overlay buffer and a stencil buffer are also standard on these offerings.

The GXT1000 internal accelerator (# 7252) is available only for the RS/6000 Model F40 (7025) and is placed behind the cover. The PCI-attachment adapter occupies one PCI slot. Note that installation of a GXT1000 internal accelerator in the Model F40 limits the customer to one *6-pack* of disks due to power supply constraints.

The GXT1000 internal accelerator supports stereo viewing. Stereo viewing adds depth to a 3D image rendered on the 2D surface of a monitor. The method supported is the CrystalEyes system by StereoGraphics Corp. Resolution is 1280x496 spaced across the full screen. For correct results, be aware that the stereo application must run in a full-screen window. Also, because stereo uses stored left and right images from the upper and lower

halves of the frame buffer, respectively, the application will need to adjust the frame buffer contents accordingly. The POWER GXT1000 is stereo-ready (sync-and power-supplied) and requires attachment of a stereo device such as the StereoGraphics CrystalEyes Stereoscopic System (which includes the required LCD glasses, emitter, and StereoGraphics cable).

The GXT1000 internal graphics accelerator supports a separately orderable Video Output Option (VOO) feature # 7254. The VOO feature is used to transmit an image to a VCR or standard TV monitor. VOO translates high-resolution RGB to either NTSC (including Y/C) or PAL outputs. For NTSC output, VOO translates the 640x480 top-left pixels of the root window. For PAL output, the 768x575 top-level pixels are translated. When translating data, only the VOO monitor is active; the RGB or HFT monitor appears blank. VOO translates data as WYSIWYG (What You See Is What You Get); however, the default color lookup table can be bypassed in favor of a customer-defined one if desired. For correct results, be aware that the application must arrange windows because VOO only captures the top-left 640x480 pixels for NTSC output or the top-left 768x575 pixels for PAL output. The VOO feature also supports Genlock, which enables applications to synchronize the VOO session with, for example, an audio and/or video (external) source. The VOO and Genlock connectors require standard 75-ohm coaxial cables with BNC connectors which must be provided by the customer.

The GXT1000 internal graphics accelerator provides support for state-of-the-art functions like depth of field, motion blur, scene antialiasing, and texture mapping. The POWER GXT1000 internal provides the same level of function as the POWER GXT1000 Model 002, including the texture memory upgrade for supporting larger textures. The GXT1000 internal supports a 256x256 32-bit texture map in the base configuration and supports a maximum of one 1024x1024 32-bit texture map with the addition of the Advanced Graphics feature # 7253.

The POWER GXT1000 accelerators support a display resolution of 1280x1024. This includes monitors that comply with Part 3 of the ISO 9241 ergonomic standard, providing customers with improved viewing and physical comfort, flicker-free images, minimized reflections, and sharper characters.

The POWER GXT1000 accelerators made IBM one of the first vendors to offer graphics accelerators optimized for OpenGL. They provide customers with the advanced graphics function and performance needed by today's advanced 3D graphics applications.

5.2.8 Graphics Adapter Feature Comparison

Table 29 and Table 30 on page 168 show the specifications and features for the RS/6000 graphics accelerators currently available.

Adapter	MVP POWER Multi Monitor	POWER GXT110P	POWER GXT120P	POWER GXT250P	POWER GXT255P	POWER GXT500P
Feature code•	# 2837	# 2839	# 2838	# 2851	# 2852	# 2854
Class	2D only			2D; 3D (Soft)	2D; 3D (Soft)	3D
PCI slots occupied•	1			1	1	1
Max/system•	1•			2 AIX	2 AIX	1
Resolution	640 x 480 800 x 600 1024 x 768 1280 x 1024			1024 x 768 1280 x 1024•	1024 x 768 1280 x 1024 1600 x 1280•	1024 x 768 1280 x 1024
ISO compliant	Yes			Yes	Yes	Yes
Frame buffer	8-bit single			8-bit single (8-bit double for 1024 x 768)	8-bit double; 24-bit single	8/12-bit double; 24-bit single
Color palette	256,000			16.7 million	16.7 million	16.7 million
Concurrent colors	256			256	Up to 16.7 million	Up to 16.7 million
Z-buffer	No			No•	No•	24-bit
HW color map	1			3	3	4
Overlay buffer	0-bit			0-bit	8-bit ••	8-bit
Dithering	No			Yes	Yes	Yes
Depth cueing	No			No	No	Yes
Anti-aliasing	No			No	No	Yes
Gouraud shading	No			No•	No•	Yes
Local lighting	No			No•	No•	No•
NURBs	No			No•	No•	No•

<i>Table 29 (Page 2 of 2). Entry-Level Graphics Adapters</i>						
Adapter	MVP POWER Multi Monitor	POWER GXT110P	POWER GXT120P	POWER GXT250P	POWER GXT255P	POWER GXT500P
Texture mapping	No			No•	No•	HA•
Alpha buffer	No			No•	No•	8-bit
Stencil buffer	No			No•	No•	8-bit
Accumulation buffer	No			No•	No•	No•
Supported APIs	Xlib,			Xlib, OpenGL, PHIGS/PEX	Xlib, OpenGL, PHIGS/PEX	Xlib, OpenGL, PHIGS/PEX, GL3.2

Table 30. 3D Graphics Adapters

Adapter	POWER GXT550P	POWER GXT800P	POWER GXT800P Texture Option	POWER GXT1000 (Internal)
Feature code	# 2855	# 2853	# 2859	# 7252
Class	3D	3D	3D	3D
PCI slots occupied	1	3	3	1•
Max/system•	1	1	1	1
Resolution	1024 x 768 1280 x 1024	1280 x 1024	1280 x 1024	1280 x 1024
ISO compliant	Yes	Yes	Yes	Yes
Frame buffer	8/24-bit double	8/24-bit double	8/24-bit double	8/12/24-bit double
Color palette	16.7 million	16.7 million	16.7 million	16.7 million
Concurrent colors	Up to 16.7 million			
Z-buffer	24-bit	24-bit	24-bit	24-bit
HW color map	4	4	4	19
Overlay buffer	8-bit	8-bit	8-bit	8-bit double
Dithering	Yes	Yes	Yes	No
Depth cueing	Yes	Yes	Yes	Yes
Anti-aliasing	Yes	Yes	Yes	Yes
Gouraud shading	Yes	Yes	Yes	Yes
Local lighting	No•	No•	No•	No•
NURBs	No•	No•	No•	Yes
Texture mapping	HA•	HA•	Yes	Yes
Alpha buffer	8-bit double	8-bit double	8-bit double	8-bit double
Stencil buffer	8-bit	8-bit	8-bit	8-bit
Accumulation buffer	No•	No•	No•	Yes
Supported APIs	Xlib, OpenGL, GL3.2, PHIGS/PEX	Xlib, OpenGL, GL3.2, PHIGS/PEX	Xlib, OpenGL, GL3.2, PHIGS/PEX	Xlib, OpenGL, GL3.2, PHIGS/PEX

Notes:

1. Accelerators with the same feature code may relate to different part numbers, depending on the RS/6000 system into which they are installed. Be sure you order the correct part number.
2. The GXT1000 Internal (# 7252) accelerator is available with the advanced graphics option (# 7253) and the video output option (V00, # 7254).
3. For placement of adapter cards and fit, also refer to *PCI Graphics Adapter Placement Reference*, SA38-0538, part number 93H3514.

4. GXT1000 Internal is available only for the RS/6000 Model F40 (no longer be available after July 24, 1998) and is placed behind the cover. The PCI-attachment adapter occupies one PCI slot. Also note that installation of a GXT1000 Internal in the Model F40 limits you to one "6-pack" of disks due to power supply constraints.
5. You can use only one POWER GXT110P graphics accelerators (for AIX) in the Model F40 (7025). The adapter is available only in the E30.
6. An RPQ to support up to two POWER GXT110P accelerators per system will be available. Refer to the announcement publications.
7. Through Softgraphics, this graphics accelerator supports Gouraud shading, lighting, NURBs, texture mapping, alpha buffer, stenciling, and accumulation buffering.
8. Texture mapping is accomplished in software but with hardware acceleration (HA) by the graphics board.
9. With AIX only.
10. If not using the 24-bit frame buffer.

5.2.9 Multiple Adapter Support

Currently, a maximum of two graphics adapters are supported in the RS/6000 Models E30, F50, and H50. A maximum of four adapters are supported on the F40. Check the adapter documentation to determine if there are model specific limitations. This section discusses how to configure a second adapter.

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
```

Changing the Default Display

In the above 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.

5.2.10 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 Xclient will be displayed. This is shown in Table 31.

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.

In normal environments, having a second screen is not very useful. As 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 a piece of 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 cluttering up the desktop.

5.3 Adapter and Device Configuration on PCI-Based RS/6000 Servers

This section includes step-by-step procedures on how to configure devices on the PCI-based RS/6000 servers and provides cable types and cabling information that is necessary on most adapters.

5.3.1 Device Types

PCI-based RS/6000 servers support attachment of devices using the following buses:

- SCSI
- PCI
- ISA

Devices on PCI-based RS/6000 systems might 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 or ISA 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, `cfgmgr`, 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. Non-native ISA devices will have to be configured manually and may even need to change some of the device's predefined or customized attribute values, especially when configuring two or more ISA adapters of the same type.

Plug and Play ISA adapters provide dynamic detection and modification of settings on the card. They require the device drivers and the operating system to support the Plug and Play architecture.

Note

AIX does support EISA I/O buses with ISA adapter cards and integrated devices, but it does not support EISA adapter cards.

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 to 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 of Micro Channel devices on Micro Channel-based RS/6000 systems and PCI devices on PCI-based RS/6000 systems whenever the Configuration Manager (cfgmgr) program is run at system boot and run time.

As for non-native ISA devices, the information data contained in the predefined part of the configuration database is not sufficient to perform automatic, conflict-free ISA device configuration. Thus, the user needs to manually customize some values to be used by the ISA device (for example, interrupt level, shared memory address and so forth) when configuring the device for the first time. This is explained in detail later in this section.

The Device Configuration Database that is currently used by the system is located in the directory that is specified by the ODMDIR environment variable, 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.

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. In the case of ISA adapters, these numbers will not actually correspond to the physical slot number. They will depend on the order in which the ISA cards are defined or configured.
- 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.

The following list shows an example of PCI-based RS/6000 systems' location codes.

sys0	Available 00-00	System Object
sysplanar0	Available 00-00	System Planar
bus0	Available 00-00	PCI Bus
mem0	Available 00-00	Memory
proc0	Available 00-00	Processor
L2cache0	Available 00-00	L2 Cache
fda0	Available 01-A0	Standard I/O Diskette Adapter
fd0	Available 01-A0-00-00	Diskette Drive
ppa0	Available 01-B0	Standard I/O Parallel Port Adapter
sa0	Available 01-C0	Standard I/O Serial Port 1
tty0	Available 01-C0-00-00	Asynchronous Terminal
sal	Available 01-D0	Standard I/O Serial Port 2
sioka0	Available 01-E0	Keyboard Adapter
kbd0	Available 01-E0-00-00	PS/2 keyboard
sioma0	Available 01-F0	Mouse Adapter
mouse0	Available 01-F0-00-00	3 button mouse
tok0	Available 01-01	IBM 16/4 PowerPC Token-Ring Adapter (isa)
ampx0	Available 01-02	X.25 CoProcessor/1 Adapter
bus2	Available 04-A0	PCI Bus
bus1	Available 04-B0	ISA Bus
scsi0	Available 04-C0	Wide SCSI I/O Controller
cd0	Available 04-C0-00-4,0	SCSI Multimedia CD-ROM Drive
rmt0	Available 04-C0-00-5,0	4.0 GB 4 mm Tape Drive
hdisk0	Available 04-C0-00-8,0	2.2 GB 16 Bit SCSI Disk Drive
hdisk1	Available 04-C0-00-9,0	2.2 GB 16 Bit SCSI Disk Drive
hdisk2	Available 04-C0-00-10,0	2.2 GB 16 Bit SCSI Disk Drive
scsi1	Available 04-01	Wide SCSI I/O Controller
rmt1	Defined 04-01-00-4,0	2.3 GB 8 mm Tape Drive
gga1	Available 04-02	IBM Power Series S15 Graphics Adapter
tok0	Available 04-03	IBM PCI Tokenring Adapter (14101800)
ent0	Available 04-04	IBM PCI Ethernet Adapter (22100020)
gga0	Available 04-05	IBM Power Series S15 Graphics Adapter

Figure 48. Devices' Location Codes

Configuration Information

Information about the ISA and PCI adapters can be obtained by running the `lscfg` command (see 7.4, “Configuration Information” on page 241).

To determine how much memory you have in your machine, run the following command:

```
lssattr -E1 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 `lssattr` command:

```
lssattr -E1 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 as important as it is in Micro Channel devices. 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.

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` program is run.

The easiest way to configure “non-graphic” PCI adapters (such as LAN adapters and so on) installed after IPL is to proceed in the same way as with Micro Channel adapters. For example, using the SMIT support, 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.

Important

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 158.

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.

Configuring SCSI Devices

All PCI servers have at least one integrated SCSI-2 Fast Interface. On the RS/6000 Model E30 and RS/6000 Model F40, only internal devices are supported, but PCI SCSI-2 adapters can be installed to attach external SCSI devices.

The process of configuring SCSI devices on PCI-based RS/6000 systems is exactly the same as it is in Micro Channel-based 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.

5.3.3 ISA Adapter Configuration

New ISA adapters are not automatically configured and the AIX device drivers are not automatically installed when they are inserted into a system. This section contains the information needed to better understand the process involved when adding an ISA adapter to your system.

When configuring ISA adapters, you will have to configure some adapter hardware parameters. The values you assign to the parameters are used by AIX to communicate with the adapter; thus, it is important they are selected carefully to perform a conflict-free configuration.

AIX manages five resources for ISA adapters. These resources are:

- Bus I/O ranges
- Bus memory address ranges
- System interrupts (IRQ)
- Direct Memory Access (DMA) channels
- Bus Memory DMA Address ranges

The first four resources are allocated by the AIX configuration subsystem when the controller is configured. The last resource is allocated by the AIX ISA DMA services when a DMA transfer is mapped. For integrated ISA adapters, the four required resources are provided in residual data with RSPC PNP (Plug and Play) descriptors.

In the case of attached ISA adapters, parameters such as interrupt level or DMA arbitration level cannot always be set dynamically. Some adapters are configured using jumpers or DIP switches that physically reside on the adapter. The system administrator must first set these parameters to avoid conflicts with other adapters before physically installing the adapter into the system.

For this reason, before configuring any ISA adapters, you should carefully plan and document the configuration. This topic includes procedures that will help you to document your current configuration.

The process of an ISA adapter configuration involves the following steps:

1. Record parameter settings of ISA adapters already configured.
2. Select parameter values for new ISA adapters.
3. Install the adapter device driver.
4. Use `smit isa` to define the adapter.
5. Set the proper DIP switches and physically install the adapter.
6. Configure the adapter using `smit isa`.

Recording Settings of ISA Adapters Already Configured

When configuring new ISA adapters, you will have to provide parameter values for these adapters. These parameter values should not already be in

use by other adapters that may already exist in your system. For this reason, and to better perform problem analysis, you should record the parameter settings for any ISA adapter already installed in your system.

Having the previously installed adapter parameters documented allows better planning for the installation of additional ISA adapters and enables you to start with a set of conflict-free definitions. To obtain a list of ISA adapters already configured in your system, enter the following command:

```
lsdev -Cc adapter -s isa
```

```
ampx0 Available 01-01 X.25 CoProcessor/1 Adapter
sa2 Available 01-02 IBM 8-Port Async, EIA-232 (ISA)
```

Figure 49. List of Configured ISA Adapters

The following procedures show you how to obtain the parameter settings of already installed ISA adapters. If AIX is up and running, you may first want to record the settings at AIX level for all your ISA adapters before shutting down and powering-off the system. If the system is powered off, you may first inspect the settings at hardware level by looking at the switches, or inspect them at firmware level by using System Management Services.

You can display the relevant parameters of configured adapters either by using the `odmget`, the `lsattr` or the `lsresource` command.

The `lsresource` command lets you display the currently assigned values for the bus resource attributes for the device specified by the given device logical name (-l Name). Use the `-a` flag to display the currently assigned values for all bus resource attributes for all devices under the same parent bus as the specified device.

For example, to display the attribute values set for the device `ampx0`, you can enter the following command as shown in Figure 50.

```
# lsresource -l ampx0 -s isa
```

TYPE	DEVICE	ATTRIBUTE	S G CURRENT VALUE
B	ampx0	bus_mem_addr	0x000e0000 - 0x000e1fff
O	ampx0	bus_io_addr	0x000002a0 - 0x000002a6
N	ampx0	bus_intr_lvl	11 (A0)

Figure 50. Attributes Listed Using the `lsresource` Command

Or use the command as shown in Figure 51 on page 180.

```
# lsattr -l amp0 -E -H
```

attribute	value	description	user_settable
bus_intr_lvl	11	Parent adapter	True
intr_priority	2	PORT number	False
bus_io_addr	0x2A0	Connection address	True
bus_mem_addr	0xE0000	Description	True
window_size	0x2000	Status	False

Figure 51. Attributes Listed Using the lsattr Command

Selecting ISA Adapter Hardware Parameters

Your system will work without problems only if there are no conflicts among the parameters for the different ISA adapters in the system. When selecting the parameters for the new ISA adapter, you can use the information collected at the previous step and the information provided by the Installation Guide booklet for the adapter. Try to use the predefined default parameters whenever possible.

The number of parameters to set depends on the type of adapter. For some adapters, such as ISA X.25 Coprocessor/1 Adapter, you have to set values for Interrupt Level, Bus I/O Address, and Bus Memory Address, and for others, such as the ISA 8-Port Asynchronous Adapter, you have to specify values only for the Bus I/O Address.

Refer to the sections describing each ISA adapter configuration to know which parameters have to be specified.

Installing Device Drivers

If not yet installed, identify and install the device driver for the new ISA adapter. To install the software using the SMIT support, you might enter the fastpath:

```
smit isa
```

Then select the **Install ISA Adapter Software** option.

The required filesets to install ISA adapters device drivers and diagnostics supported by AIX Version 4 are listed as follows:

- devices.isa.c1x** X.25 CoProcessor/1 Adapter Software
- devices.isa.cxia** ISA 8-Port Asynchronous Adapter Software
- devices.isa.cxia128** ISA 128-Port Asynchronous Adapter Software
- devices.isa.mm2** 4-Port Selectable Multiport Model 2 Device Driver

devices.isa.ethernet IBM ISA Ethernet Adapter Software

devices.isa.tokenring IBM 16/4 PowerPC Token-Ring Adapter (isa) Software

Defining ISA Adapters at AIX Level

A SMIT menu is provided with AIX Version 4.1.4, or later, that supports configuration of ISA adapters.

A second method to configure ISA adapter attributes is accomplished by using the `mkdev` command. In this case, you can specify the attributes directly on the command line. When possible, we suggest you to use the SMIT menu.

The SMIT support for some of the ISA adapters allows you to define the adapter to the operating system before physically installing it. The `smit isa` fastpath provides a menu to configure and manage ISA adapters, as shown in Figure 52.

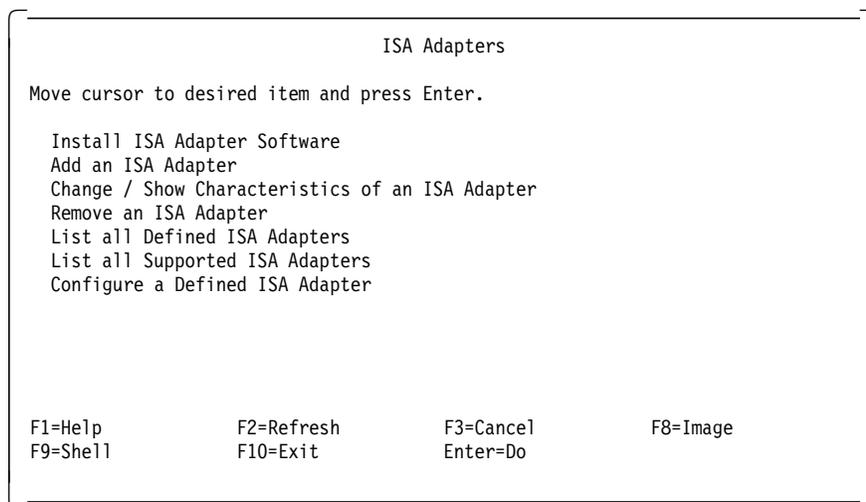


Figure 52. SMIT ISA Menu

Once there, select **Add an ISA Adapter**. The follow-up screen is not the same for all adapters. Some adapters require more parameters to be specified than others. Moreover, some menus allow you to define the adapter without configuring it. If that is the case for your adapter, specify the selected parameters, and set the **Define device only, do not configure** menu field to yes.

If the adapter definition is successful, it means that your selected parameters could be resolved by the system and are not in conflict with any

other ISA adapters already configured. If the definition process fails, the `lsresource` command will display information on the SMIT menu about which parameters could not be resolved and will display recommended values for these parameters.

If the **Add an ISA Adapter** menu does not provide the field **Define device only, do not configure**, you can still execute the configuration. Since the adapter is not yet physically installed, the configuration will fail. However, if no conflict error messages are displayed on the SMIT menu by the `lsresource` command, it means that your selected parameters values would be the right ones to configure the adapter once it is physically installed in the system.

Setting the Adapter DIP Switches

Once the values parameters have been selected and you have tested and found there are no conflicts with ISA adapters already installed on the system by using the **Add an ISA adapter** menu, you can select these values by setting the DIP switches on the adapter as described in its Installation Guide booklet.

Matching Parameters:

It is very important to understand that some parameters, such as the IRQ and I/O Port, are set at both the hardware and AIX level, and the settings must be identical.

Refer to the following sections describing each ISA adapter configuration to know how to set DIP switches on each particular adapter.

Shut down and power-off the system, and then physically install the adapter in an ISA free slot.

Making ISA Adapters Available

After physically installing the adapter, boot the system again, and configure the ISA adapter using `smit isa fastpath`.

If the ISA adapter is in the defined state, you can select **Configure a Defined ISA Adapter** from the menu to set it to the Available state.

Otherwise, if the ISA adapter could not be set to the defined state since the SMIT menu does not allow the adapter definition only, select **Add an ISA Adapter** to set the parameter values, and configure it.

5.3.4 8-Port EIA-232 Asynchronous ISA Adapter (# 2931)

This section provides information needed for installing and configuring the 8-Port Asynchronous ISA Adapter. This adapter can be ordered as an optional feature on PCI-based RS/6000 servers.

The 8-Port Asynchronous ISA Adapter is a multichannel, intelligent serial communications feature which supports speeds of 155 Kbps for each asynchronous port. It contains the following components:

- 8-port asynchronous card
- 8-port connector box
- Optional media containing device driver or diagnostic software

The device driver for the 8-Port Asynchronous ISA Adapter is included in the `devices.isa.cxia` fileset.

Note

Refer to the *8-Port Asynchronous ISA Adapter Installation Guide* shipped with the adapter for detailed information on this adapter.

DIP-Switch Setting

The only parameter you have to set to install the adapter into the system is the *Bus I/O Address*. Its possible values are set with four switches at the rear of the adapter, and the possible switch settings are described in the following table.

Hex Addr	Address Switch			
	1	2	3	4
104	Off	Off	On	On
114	Off	On	Off	On
124	Off	On	On	On
204	On	Off	Off	On
224	On	Off	On	On
304	On	On	Off	On
324	On	On	On	On

Note:

The *on* and *off* DIP-switch selections are set by switching them *down* and *up*, respectively. The mapping is:

off = up
on = down

Configuring the Adapter

The following procedures apply to every PCI-based RS/6000 system that supports the 8-Port Asynchronous ISA Adapter.

1. Check if there are any ISA adapters that are already installed in your computer by using the following command:

```
lsdev -Cc adapter -s isa
```

Default Parameters

If the 8-Port Asynchronous ISA Adapter is the first ISA adapter of this type installed into your system, you may use the default Bus I/O Address of 0x324.

2. If any ISA adapter is installed, enter the command:

```
lsresource -l bus0 -a | grep bus_io
```

This command generates a list of adapters using bus I/O that looks similar to the following:

0	fda0	bus_io_addr		0x000003f0 - 0x000003f5
0	fda0	bus_io_addr2		0x000003f7
0	ppa0	bus_io_addr		0x000003bc - 0x000003be
0	sa0	bus_io_addr		0x000003f8 - 0x000003ff
0	sa1	bus_io_addr		0x000002f8 - 0x000002ff
0	sioka0	bus_io_addr	1	0x00000060
0	sioma0	bus_io_addr	1	0x00000060
0	ampx0	bus_io_addr		0x000002a0 - 0x000002a6
0	sa2	bus_io_addr		0x00000104 - 0x00000107
0	scsi0	bus_io_addr		0x0000f000 - 0x0000f0ff
0	scsi1	bus_io_addr		0x01000000 - 0x010000ff

3. Select a Bus I/O Address that is not currently in use.
4. Install the 8-Port Asynchronous ISA Adapter device driver and diagnostic by choosing:

```
smit isa
```

then select the **Install ISA Adapter Software** option.

To install the fileset device.isa.cxia.

5. After installing the software, you might enter:

```
smit isa
```

then select the **Add an ISA Adapter** option.

6. Select **pcxr isa IBM 8-Port Async, EIA-232 (ISA)** from the SMIT pop-up menu.
7. Select **bus1 Available ISA Bus** from the following SMIT pop-up menu.
8. From the menu described in Figure 53, select the correct Bus I/O Address, and press Enter.

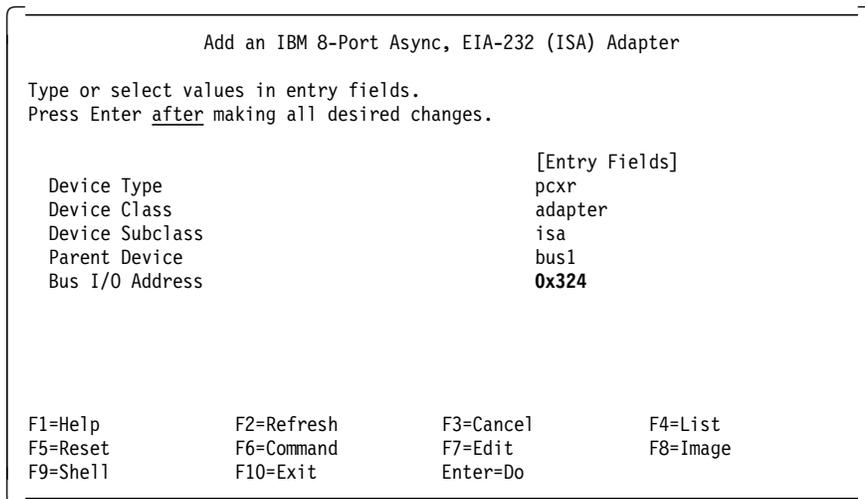


Figure 53. SMIT Configuration Menu for Adding an 8-Port Asynchronous ISA Adapter

This configuration menu does not provide the capability to only define the adapter without attempting to configure it. So, the configuration operation will fail because the adapter is not yet physically installed on the system. If no bus I/O address conflicts are detected, the following message will be displayed:

```
sa2 Defined
Method error (/usr/lib/methods/cfgcxma):
    0514-040 Error initializing a device into the kernel.

sa2 deleted
```

This error does not matter. You should use this procedure just to make sure that the selected adapter parameter is correct.

9. If the Bus I/O Address is in conflict with the parameter setting for another ISA adapter already installed, the following message will be displayed:

```
sa3 Defined
lsresource : The attribute(s) for some device(s) in the system
could not be resolved. To resolve conflicts, attribute(s) need to be
modified. A suggested value for each attribute is provided.

DEVICE          ATTRIBUTE          CURRENT  SUGGESTED DESCRIPTION
-----
sa3             bus_io_addr        0x324   0x304   Connection address

Method error (/usr/lib/methods/cfgcxma):
    0514-040 Error initializing a device into the kernel.

sa3 deleted
```

Change the parameter value as suggested by the `lsresource` command, and press Enter.

10. Once no parameter conflicts are reported, you can set the DIP switches on the adapter card as described in Table 32 on page 183 to reflect the selected Bus I/O Address's correct value.
11. Shut down and power-off the system.
12. Physically install the adapter on an ISA slot.
13. Power-on and reboot the system.
14. Configure the adapter using the fastpath `smit isa`, and repeat the procedure from step 5 to step 8.
15. Verify the results with the following command:

```
lsdev -Cc adapter -s isa
```

If the configuration was successful, the adapter will be available, and the system will display information similar to the following:

```
sa3 Available 01-01 IBM 8-Port Async, EIA-232 (ISA)
```

16. Once the adapter status is Available, you may proceed to configure the serial devices on its ports by using the `smit devices fastpath`.

5.3.5 128-Port EIA-232 Asynchronous ISA Adapter (# 2933)

This section provides information needed for installing and configuring the 128-Port Asynchronous ISA Adapter. This adapter can be ordered as an optional feature on PCI-based RS/6000 servers.

The 128-Port Asynchronous ISA Adapter provides support for attaching a maximum of 128 EIA-232D asynchronous serial devices (such as terminals, modems and printers) to a system unit when used with one or more optional Remote Asynchronous Nodes 16-Port EIA-232. It contains the following components:

- 128-port asynchronous card
- Two terminators
- Asynchronous wrap plug

The device driver for the 128-Port Asynchronous ISA Adapter is included in `devices.isa.cxial28` fileset.

Note

Refer to the *128-Port Asynchronous ISA Adapter Installation Guide* shipped with the adapter for detailed information on this adapter.

DIP-Switch Setting

The only parameter you have to set to install the adapter into the system is the *Bus I/O Address*. Its possible values are set with four switches at the rear of the adapter, and the possible switch settings are described in the following table.

Table 33. 128-Port Asynchronous ISA Adapter DIP Switches

Hex Addr	Address Switch			
	1	2	3	4
108	Off	Off	On	On
118	Off	On	Off	On
128	Off	On	On	On
208	On	Off	Off	On
228	On	Off	On	On
308	On	On	Off	On
328	On	On	On	On

Note:

The *on* and *off* DIP-switch selections are set by switching them *down* and *up*, respectively. The mapping is:

off = up
on = down

Configuring the Adapter

The following procedures apply to every PCI-based RS/6000 system that supports the 128-Port Asynchronous ISA Adapter.

1. Check if there are any ISA adapters that are already installed in your computer by using the following command:

```
lsdev -Cc adapter -s isa
```

Default Parameters

If the 128-Port Asynchronous ISA Adapter is the first ISA adapter of this type installed into your system, you may use the default Bus I/O Address of 0x328

2. If any ISA adapter is installed, enter the command:

```
lsresource -l bus0 -a | grep bus_io
```

This command generates a list of adapters using bus I/O that looks similar to the following:

0	fda0	bus_io_addr	0x000003f0 - 0x000003f5
0	fda0	bus_io_addr2	0x000003f7
0	ppa0	bus_io_addr	0x000003bc - 0x000003be
0	sa0	bus_io_addr	0x000003f8 - 0x000003ff
0	sa1	bus_io_addr	0x000002f8 - 0x000002ff
0	sioka0	bus_io_addr	1 0x00000060
0	sioma0	bus_io_addr	1 0x00000060
0	ampx0	bus_io_addr	0x000002a0 - 0x000002a6
0	sa2	bus_io_addr	0x00000104 - 0x00000107
0	scsi0	bus_io_addr	0x0000f000 - 0x0000f0ff
0	scsi1	bus_io_addr	0x01000000 - 0x010000ff

3. Select a Bus I/O Address that is not currently in use.
4. Install the 128-Port Asynchronous ISA Adapter device driver and diagnostic by choosing:

```
smit isa
```

then select the **Install ISA Adapter Software** option.

To install the fileset device.isa.cxial28.

5. After installing the software, you might enter:

```
smit isa
```

then select the **Add an ISA Adapter** option.

6. Select **cxia isa IBM 128-Port Async, EIA-232 (ISA)** from the SMIT pop-up menu.
7. Select **bus1 Available ISA Bus** from the following SMIT pop-up menu.
8. From the menu described in Figure 54, select the correct Bus I/O Address, and press Enter.

```

                                Add an IBM 128-Port Async, EIA-232 (ISA) Adapter

Type or select values in entry fields.
Press Enter after making all desired changes.

                                [Entry Fields]
Device Type                       cxia
Device Class                      adapter
Device Subclass                   isa
Parent Device                     bus1
Bus I/O Address                   0x328
Line 1 Speed                      1.2M
Line 1 Cable Type                 8
    Node 1 Comm Mode              direct
    Node 2 Comm Mode              direct
    Node 3 Comm Mode              direct
    Node 4 Comm Mode              direct
Line 2 Speed                      1.2M
Line 2 Cable Type                 8
    Node 1 Comm Mode              direct
    Node 2 Comm Mode              direct
    Node 3 Comm Mode              direct
    Node 4 Comm Mode              direct

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F9=Shell     F10=Exit       Enter=Do

```

Figure 54. SMIT Configuration Menu for Adding an 128-Port Asynchronous ISA Adapter

This configuration menu does not provide the capability to only define the adapter without attempting to configure it. So, the configuration operation will fail because the adapter is not yet physically installed on the system. If no Bus I/O Address conflicts are detected, the following message will be displayed:

```

cxia0 Defined
Method error (/usr/lib/methods/cfgcxma):
    0514-040 Error initializing a device into the kernel.

cxia0 deleted

```

It does not matter. You should use this procedure just to make sure that the selected adapter parameter is correct.

9. If the Bus I/O Addresses are in conflict with another parameter set for another ISA adapter already installed, the following message will be displayed:

```
cxia0 Defined
lsresource : The attribute(s) for some device(s) in the system
could not be resolved. To resolve conflicts, attribute(s) need to be
modified. A suggested value for each attribute is provided.

DEVICE          ATTRIBUTE          CURRENT  SUGGESTED DESCRIPTION
-----
cxia0           bus_io_addr      0x328   0x308   Connection address

Method error (/usr/lib/methods/cfgcxma):
0514-040 Error initializing a device into the kernel.

cxia0 deleted
```

Change the parameter value as suggested by the `lsresource` command, and press Enter.

10. Once no parameter conflicts are reported, you can set the DIP switches on the adapter card, as described in Table 33 on page 187, to reflect the selected Bus I/O Address's correct value.
11. Shut down and power-off the system.
12. Physically install the adapter on an ISA slot.
13. Power on and reboot the system.
14. Configure the adapter using the fastpath `smit isa`, and repeat the procedure from step 5 to step 8.
15. Verify the results with the following command:

```
lsdev -Cc adapter -s isa
```

If the configuration was successful, the adapter will be available, and the system will display information similar to the following:

```
cxia0 Available 01-02 IBM 128-Port Async, EIA-232 (ISA)
```

16. Once the adapter status is Available, you may proceed to configure the serial devices on its ports using the `smit devices fastpath`.

5.3.6 4-Port Multi-Protocol ISA Adapter (# 2701)

This section provides information needed for installing and configuring the 4-Port Multi-Protocol ISA Adapter. This adapter can be ordered as an optional feature on PCI-based RS/6000 servers.

The 4-Port Multi-Protocol ISA Adapter provides asynchronous support and synchronous communication. The adapter consists of a base card and a daughter card. The two cards are physically connected and require a single ISA bus card slot.

The device driver for the 4-Port Multi-Protocol ISA Adapter is included in the `devices.isa.mm2` filesset.

Note

Refer to the *4-Port Multi-Protocol ISA Adapter Installation Guide* shipped with the adapter for detailed information on this adapter.

Cable Selections for 4-Port Multiprotocol Interface

This adapter can support four different physical interfaces by using an IBM 4-Port Multiprotocol Interface Cable (feature #2705).

This cable is 3.1 meters (10 feet) long. It has a 78-pin connector on one end and a molded distribution block on the other end that contains six 25-pin connectors and three 15-pin connectors for the different physical interfaces supported.

The following table shows the supported interfaces for each port.

Port Number	Supported Interfaces			
# 0	EIA-232	EIA-422	V.35	X.21
# 1	EIA-232	-	V.35	-
# 2	EIA-232	EIA-422	-	-
# 3	EIA-232	-	-	-

The maximum cable lengths that are supported for the various communications protocols are described in the following table.

Protocol	Length (meters)	Length (feet)
RS-232	15.2	50
V.35	15.2	50
RS-422	122	400
X.21	122	400

DIP-Switch Setting

The 4-Port Multi-Protocol ISA Adapter has a 10-position DIP switch which is used to set the following parameters:

Parameter	Switch Number(s)
Interrupt Level	1-2-3
Memory Size	4
Bus I/O Address	5-6-7-8
Edge Connector	9
Bus Width	10

You do not need to set every DIP switch on the adapter to configure it. The only DIP switches you have to change are those referring to *Interrupt Level* and *Bus I/O Address*.

The DIP switches indicating the *Memory Size*, the *Edge Connector* and the *Bus-Width* can be left in the default position set by the factory.

The following table shows the parameters values you may want to select to define the Interrupt Level and the Bus I/O Address for a conflict-free configuration of your first, second and third 4-Port Multiprotocol adapters.

Dip Switch #	Position			Comment
	First Adapter	Second Adapter	Third Adapter	
1	on	off	on	IRQ Level <ul style="list-style-type: none"> • First Card: 7 • Second Card: 9 • Third Card: 10
2	off	off	on	
3	on	on	off	
4	off	off	off	Memory Size a
5	off	off	off	I/O Port Address <ul style="list-style-type: none"> • First Card: 06A0 • Second Card: 0EA0 • Third Card: 016A0
6	on	off	on	
7	on	on	off	
8	on	on	on	
9	on	on	on	Edge Connector b
10	off	off	off	Bus-Width c

Notes:

Refer to *IBM RS/6000 Adapters, Devices, and Cable Information for Multi Bus Systems*, SA38-0516, for detailed information on DIP-switch positions for setting other parameter values.

The *on* and *off* DIP-switch selections are set by switching them *down* and *up*, respectively. The mapping is:

off = up
on = down

a This switch is factory set to indicate the size of RAM installed on the adapter. Verify that the position of this switch is *off* to reflect a RAM size of 1 MB.

b Indicates whether a one-edge (62-pin) or a two-edge (62-pin and 36-pin) connector is used to hold your adapter.
on = two-edged / off = one-edged

c Sets the bus width and should generally be set to 8-bit.
on = 8-bit / off = 16-bit

Configuring the Adapter

The following procedures apply to every PCI-based RS/6000 system that supports the 4-Port Multi-Protocol ISA Adapter.

The configuration for this adapter requires you to specify the *Interrupt Level*, the *Bus I/O Address* and the *Bus Memory Address* to be used with the adapter.

1. Check if there are any ISA adapters that are already installed in your computer by using the following command:

```
lsdev -Cc adapter -s isa
```

Default Parameters

If this is the first adapter of this type on the system, you may select the default values which are:

- Interrupt Level = 7
- Bus I/O Address = 0x6a0
- Bus Memory = 0X0e2000

and go directly to step number 9.

2. If any ISA adapter is installed, check for available *Interrupt Level* values by entering the following command:

```
lsresource -l bus0 -a | grep bus_intr
```

This command generates a list of adapters using different Interrupt Levels that looks similar to the following:

N	fda0	bus_intr_lvl	6 (A0)
N	sa0	bus_intr_lvl	4 (A0)
N	sa1	bus_intr_lvl	3 (A0)
N	sioka0	bus_intr_lvl	1 (A0)
N	sioma0	bus_intr_lvl	12 (A0)
N	ampx0	bus_intr_lvl	11 (A0)
N	apm0	bus_intr_lvl	7 (A0)
N	scsi0	bus_intr_lvl	13 (A0)

3. Select an Interrupt Level that is not currently in use.
4. If any ISA adapter is installed, check for available *Bus I/O Address* values by entering the following command:

```
lsresource -l bus0 -a | grep bus_io
```

This command generates a list of adapters using bus I/O that looks similar to the following:

0	fda0	bus_io_addr		0x000003f0 - 0x000003f5
0	fda0	bus_io_addr2		0x000003f7
0	ppa0	bus_io_addr		0x000003bc - 0x000003be
0	sa0	bus_io_addr		0x000003f8 - 0x000003ff
0	sa1	bus_io_addr		0x000002f8 - 0x000002ff
0	sioka0	bus_io_addr	1	0x00000060
0	sioma0	bus_io_addr	1	0x00000060
0	ampx0	bus_io_addr		0x000002a0 - 0x000002a6
0	apm0	bus_io_addr		0x000006a0 - 0x000006a6
0	sa2	bus_io_addr		0x00000104 - 0x00000107
0	scsi0	bus_io_addr		0x0000f000 - 0x0000f0ff
0	scsi1	bus_io_addr		0x01000000 - 0x010000ff

5. Select a Bus I/O Address that is not currently in use.
6. If any ISA adapter is installed, check for available *Bus Memory* values by entering the following command:

```
lsresource -l bus0 -a | grep bus_mem
```

This command generates a list of adapters using different Bus Memory Addresses that looks similar to the following:

B	ampx0	bus_mem_addr		0x000e0000 - 0x000e1fff
B	apm0	bus_mem_addr		0x000e2000 - 0x000e3fff
B	gga0	bus_mem_start		0x01000000 - 0x01ffffff

7. Select a bus memory address that is not currently in use.

The valid bus memory values for use with the 4-Port Multi-Protocol ISA Adapter start at 0x0e0000 and increment by 0x002000, such as 0x0e2000, 0x0e4000, 0x0e6000, etc. The previous example shows ampx0 is using the first bus memory range of 0x0e0000, and apm0 is using the second range of 0x0e2000.
8. Install the 4-Port Multi-Protocol ISA Adapter device driver and diagnostic by choosing:

```
smit isa
```

then select the **Install ISA Adapter Software** option.
to install the fileset device.isa.mm2.
9. After installing the software, you might enter:

```
smit isa
```

then select the **Add an ISA Adapter** option.
10. Select **Portmaster isa Multiport Model 2 Adapter** from the SMIT pop-up menu.
11. Select **bus1 Available ISA Bus** from the following SMIT pop-up menu.

12. From the menu described in Figure 55 on page 196, select the correct Interrupt Level, Bus I/O Address and Bus Memory values, and press Enter.

```

                                Add a Multiport Model 2 Adapter

Type or select values in entry fields.
Press Enter after making all desired changes.

Description                                [Fields]
Parent Device                               Multiport Model 2 Adapter
ISA Interrupt Level                         bus1
Bus IO Address                             [7]
Bus Memory Address                          [0x6A0]
Define device only, do not configure        [0xE2000]
                                           yes

F1=Help      F2=Refresh      F3=Cancel      F4=List
F5=Reset     F6=Command     F7=Edit       F8=Image
F9=Shell     F10=Exit        Enter=Do

```

Figure 55. SMIT Configuration Menu for Adding an 4-Port Multi-Protocol ISA Adapter

This configuration menu also provides the capability to select the **Define device only, do not configure** option. Change this field value to **yes**. It will allow you to define the adapter without configuring it. If any parameter conflicts were detected, an error message similar to the following is displayed:

```

apml Defined
lsresource : The attribute(s) for some device(s) in the system
could not be resolved. To resolve conflicts, attribute(s) need to be
modified. A suggested value for each attribute is provided.

DEVICE          ATTRIBUTE          CURRENT  SUGGESTED DESCRIPTION
-----
apml            bus_intr_lvl  7        9      Interrupt Level
apml            bus_io_addr   0x6a0    0xea0  Connection Address

apml deleted

```

Change the parameter value as suggested by the `lsresource` command, and press Enter.

13. Once no parameter conflicts are reported, you can set the DIP switches on the adapter card as described in “DIP-Switch Setting” on page 192, to reflect the selected parameters correct values.
14. Shut down and power-off the system.
15. Physically install the adapter on an ISA slot.
16. Power on and reboot the system.
17. Configure the adapter using the fastpath:

```
smit isa
```

then select the **Configure a Defined ISA Adapter** option.

18. Verify the results with the following command:

```
lsdev -Cc adapter -s isa
```

If the configuration was successful, the adapter will be available, and the system will display information similar to the following:

```
apm1 Available 01-03 IBM Multiport Model 2 Adapter
```

19. Once the adapter status is Available, you may proceed to configure its ports by using the `smit devices` fastpath.

5.3.7 X.25 Interface Co-Processor ISA Adapter (# 2961)

This section provides information needed for installing and configuring the X.25 Interface Co-Processor ISA Adapter. This adapter can be ordered as an optional feature on PCI-based RS/6000 servers.

The X.25 Interface Co-Processor ISA Adapter, with supporting software, enables the attachment of an ISA-bus-compatible system to an X.25 packet-switched network so that the computer can operate as a packet terminal.

In AIX Version 4, support for X.25 communication is provided with the IBM AIXLink/X.25 Version 1.1 LPP (feature code 5696-926).

The device driver for the X.25 Interface Co-Processor ISA Adapter is included in `devices.isa.c1x` fileset that can be found in the AIX Version 4 base operation system shipping.

Note

Refer to the *X.25 Interface Co-Processor ISA Adapter Installation Guide* shipped with the adapter for detailed information on this adapter.

DIP-Switch Setting

The X.25 Interface Co-Processor ISA Adapter has a 10-position DIP switch which is used to set the following parameters:

Table 38. DIP Switches on the X.25 Interface Co-Processor ISA Adapter

Parameter	Switch Number(s)
Interrupt Level	1-2-3
Memory Size	4
Bus I/O Address	5-6-7-8
Edge Connector	9
Bus Width	10

You do not need to set every DIP switch on the adapter to configure it. The only DIP switches you have to change are those referring to *Interrupt Level* and *Bus I/O Address*.

The DIP switches indicating the *Memory Size*, *Edge Connector* and *Bus-Width* can be left in the default position set by the factory.

The following table shows the parameter values you may want to select to define the *Interrupt Level* and the *Bus I/O Address* for a conflict-free configuration of your first, second and third X.25 Co-Processor Adapters.

Table 39 (Page 1 of 2). X.25 Adapter: DIP-Switch Suggested Settings

Dip Switch #	Position			Comment
	First Adapter	Second Adapter	Third Adapter	
1	off	off	on	IRQ Level <ul style="list-style-type: none">• First Card: 11• Second Card: 9• Third Card: 7
2	on	off	off	
3	off	on	on	
4	on	on	on	Memory Size a

Table 39 (Page 2 of 2). X.25 Adapter: DIP-Switch Suggested Settings				
Dip Switch #	Position			Comment
	First Adapter	Second Adapter	Third Adapter	
5	on	off	on	I/O Port Address <ul style="list-style-type: none"> • First Card: 02A0 • Second Card: 06A0 • Third Card: 0AA0
6	on	on	off	
7	on	on	on	
8	on	on	on	
9	on	on	on	Edge Connector b
10	on	on	on	Bus-Width c

Notes:

Refer to *IBM RS/6000 Adapters, Devices, and Cable Information for Multi Bus Systems*, SA38-0516, for detailed information on DIP-switch positions for setting other parameter values.

The *on* and *off* DIP-switch selections are set by switching them *down* and *up*, respectively. The mapping is:

off = up
on = down

a This switch is factory set to indicate the size of RAM installed on the adapter. Verify that the position of this switch is *on* to reflect a RAM size of 512 KB.

b Indicates whether a one-edge (62-pin) or a two-edge (62-pin and 36-pin) connector is used to hold your adapter.
on = two-edged / off = one-edged

c Sets the bus width and should generally be set to 8-bit.
on = 8-bit / off = 16-bit

Configuring the Adapter

The following procedures apply to every PCI-based RS/6000 system that supports the X.25 Interface Co-Processor ISA Adapter.

The configuration for this adapter requires you to specify the Interrupt Level, the Bus I/O Address and the Bus Memory Address to be used with the adapter.

1. Check if there are any ISA adapters that are already installed in your computer by using the following command:

```
lsdev -Cc adapter -s isa
```

Default Parameters

If this is the first adapter of this type on the system, you may select the default values:

- Interrupt Level = 11
- Bus I/O Address = 0x2a0
- Bus Memory = 0X0e0000

and go directly to step number 9.

2. If any ISA adapter is installed, check for available Interrupt Level values by entering the following command:

```
lsresource -l bus0 -a | grep bus_intr
```

This command generates a list of adapters using different Interrupt Levels that looks similar to the following:

N	fda0	bus_intr_lvl	6 (A0)
N	sa0	bus_intr_lvl	4 (A0)
N	sa1	bus_intr_lvl	3 (A0)
N	sioka0	bus_intr_lvl	1 (A0)
N	sioma0	bus_intr_lvl	12 (A0)
N	ampx0	bus_intr_lvl	11 (A0)
N	scsi0	bus_intr_lvl	13 (A0)

3. Select an Interrupt Level that is not currently in use.
4. If any ISA adapter is installed, check for available Bus I/O Address values by entering the following command:

```
lsresource -l bus0 -a | grep bus_io
```

This command generates a list of adapters using bus I/O that looks similar to the following:

0	fda0	bus_io_addr	0x000003f0 - 0x000003f5
0	fda0	bus_io_addr2	0x000003f7
0	ppa0	bus_io_addr	0x000003bc - 0x000003be
0	sa0	bus_io_addr	0x000003f8 - 0x000003ff
0	sa1	bus_io_addr	0x000002f8 - 0x000002ff
0	sioka0	bus_io_addr	0x00000060
0	sioma0	bus_io_addr	0x00000060
0	ampx0	bus_io_addr	0x000002a0 - 0x000002a6
0	sa2	bus_io_addr	0x00000104 - 0x00000107
0	scsi0	bus_io_addr	0x0000f000 - 0x0000f0ff
0	scsi1	bus_io_addr	0x01000000 - 0x010000ff

5. Select a Bus I/O Address that is not currently in use.

6. If any ISA adapter is installed, check for available Bus Memory values by entering the following command:

```
lsresource -l bus0 -a | grep bus_mem
```

This command generates a list of adapters using different Bus Memory Addresses that looks similar to the following:

B	ampx0	bus_mem_addr	0x000e0000 - 0x000e1fff
B	apm0	bus_mem_addr	0x000e2000 - 0x000e3fff
B	gga0	bus_mem_start	0x01000000 - 0x01ffffff

7. Select a Bus Memory that is not currently in use.

The valid Bus Memory values for use with the X.25 Interface Co-Processor ISA Adapter start at 0x0e0000 and increment by 0x002000. For example, they are 0x0e2000, 0x0e4000, 0x0e6000, and so on. The previous example shows ampx0 is using the first Bus Memory range of 0x0e0000, and apm0 is using the second range of 0x0e2000.

8. Install the X.25 Interface Co-Processor ISA Adapter device driver and diagnostic by choosing:

```
smit isa
```

then select the **Install ISA Adapter Software** option.

to install the fileset device.isa.c1x.

9. After installing the software, you might enter:

```
smit isa
```

then select the **Add an ISA Adapter** option.

10. Select **articmpx isa X.25 Coprocessor/1 Adapter** from the SMIT pop-up menu.

11. Select **bus1 Available ISA Bus** from the following SMIT pop-up menu.

12. From the menu described in Figure 56 on page 202, select the correct Interrupt Level, Bus I/O Address and Bus Memory values, and press Enter.

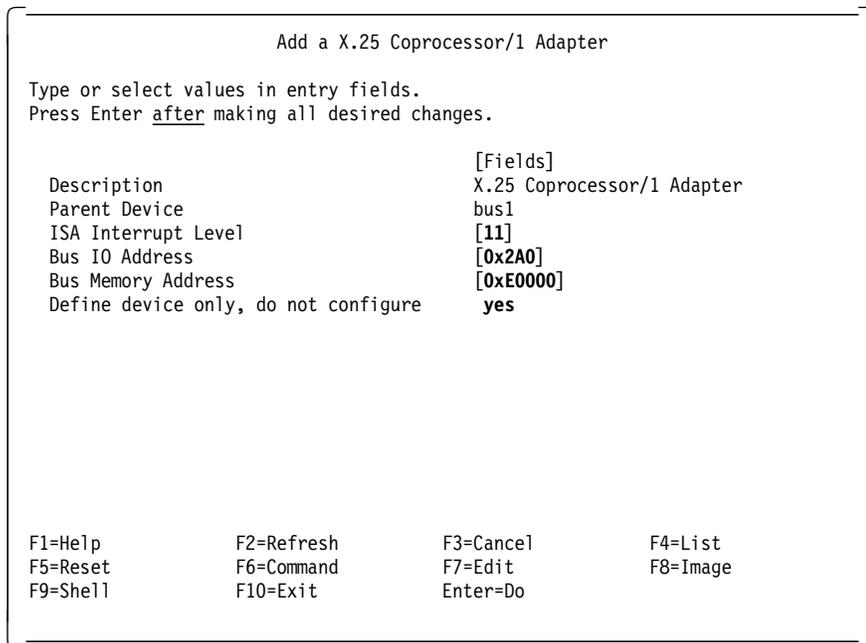
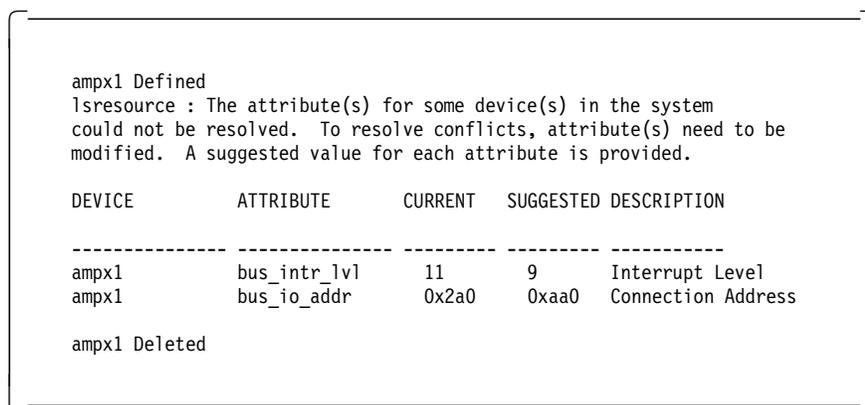


Figure 56. SMIT Configuration Menu for Adding an X.25 Interface Co-Processor ISA Adapter

This configuration menu also provides the capability to select the **Define device only, do not configure** option. Change this field value to **yes**. It will allow you to define the adapter without configuring it. If any parameter conflicts were detected, an error message similar to the following is displayed:



Change the parameter value as suggested by the `lsresource` command, and press Enter.

13. Once no parameter conflicts are reported, you can set the DIP switches on the adapter card, as described in “DIP-Switch Setting” on page 198, to reflect the selected parameters correct values.

14. Shut down and power-off the system.

15. Physically install the adapter on an ISA slot.

16. Power on and reboot the system.

17. Configure the adapter using the fastpath:

```
smit isa
```

then select the **Configure a Defined ISA Adapter** option.

18. Verify the results with the following command:

```
lsdev -Cc adapter -s isa
```

If the configuration was successful, the adapter will be available, and the system will display information similar to the following:

```
ampx1 Available 01-04 X.25 Coprocessor/1 Adapter
```

19. Once the adapter status is Available, you may proceed to configuring and setting up a X.25 network. Refer to the *AIX/6000 X.25 Cookbook*, which explains the X.25 concepts and provides procedures to set up an X.25 network using the AIXLink/X.25 LPP software.

Chapter 6. Diagnostics and System Configuration Management

One of the most important features provided by the PCI-based RS/6000 servers is support of the AIX diagnostics programs. Servers, by their nature, should provide enhanced Reliability, Availability and Serviceability (RAS) features. Access to AIX diagnostics, stand-alone and online, helps provide these features

In addition to providing diagnostics support for new PCI and ISA devices available for the PCI-based RS/6000 servers, a major change had to be made to the firmware. On Micro Channel-based RS/6000 systems, diagnostics can be loaded by switching the physical key to the service position and booting the system. If using online diagnostics, AIX will determine that the key is in the Service position and boot diagnostics. Support for booting diagnostics from CD-ROM is also provided. PCI-based RS/6000 servers do not have a physical key; so an extra function had to be added to the firmware to allow a logical representation of the keyswitch which would allow the system to either boot in *normal* or *service* mode. The E30, F40, F50, and H50 have an "electronic keyswitch" function (see 7.1, "Electronic Key-Switch Function" on page 231) 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 CD-ROM.

Two methods of running diagnostics are supported on the PCI-based RS/6000 servers:

- **Stand-Alone Diagnostics**
- **Online 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 installed. Diagnostics support is, however, automatically installed when you install from an AIX Version 4 server CD-ROM. The Online Diagnostics test list will display entries for all the devices in your system which have diagnostics support.

- **Selecting a Console Display**

When you run Stand-Alone Diagnostics, you will have to select the console display. The diagnostics 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 the user to enter a valid terminal type, and the menu will be re-displayed until a valid type is entered or the user exits from the option.

- **Resetting the Terminal**

If the user enters 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 which results from this resetting action varies with the mode in which the system is being run:

- Online Concurrent or Single-User Mode - the command prompt appears.
- Stand-Alone Mode or Online Service Mode - the terminal type will be reset to “dumb,” the diagnostics operation instruction panel will be displayed, and the user 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 3161, 3164 and 3151 are listed in the *IBM RS/6000 Diagnostics Information for Multiple Bus Systems*, SA38-0509.

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 Online 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 Online 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 which 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 (Online 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.
- ISA adapters can only be tested if they have been configured to the system. If you boot diagnostics from CD-ROM, you will have to configure the ISA adapters before you can test them.

6.2.1 Booting the Stand-Alone Diagnostics CD-ROM

To run 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. (The CD-ROM drive is the second device in this list and, provided there is not a bootable diskette in the diskette drive, the system will boot the diagnostics CD-ROM which you have inserted).
5. Enter any requested password (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.

If your system contains ISA adapters, you will not be able to automatically test them. You will have to configure the adapters to the system using the ISA Configuration menu found in Service Aids. The ISA Configuration Service Aids will display all the ISA adapters supported by the diagnostics which you are running. Diagnostic support for ISA adapters not shown in the list may be supported from a supplemental diskette.

Configuring ISA Adapters from Diagnostics

As mentioned above, when booting into diagnostics from CD-ROM, you can only test the ISA adapters in your system after they have been configured. The configuration is only for diagnostics purposes and will not affect the adapters configuration in AIX.

The following procedure explains how to configure an adapter using diagnostics. For the purpose of the explanation, we have used an 8-port adapter as an example.

1. Before you boot into diagnostics, note the ISA adapter parameters which are set for the adapter(s) you wish to test. These will either be the parameters you entered to install the card (see 5.3.3, "ISA Adapter Configuration" on page 177) or parameters you have set on the card if the card has not yet been configured by AIX. For the example, the only parameter we had to record was the Bus I/O Address.
2. Boot into diagnostics from the CD-ROM as explained above.
3. Select **Service Aids**.
4. Select **ISA Adapter Configuration Service Aid**. This will provide you with a menu like the one shown in Figure 57 on page 209.

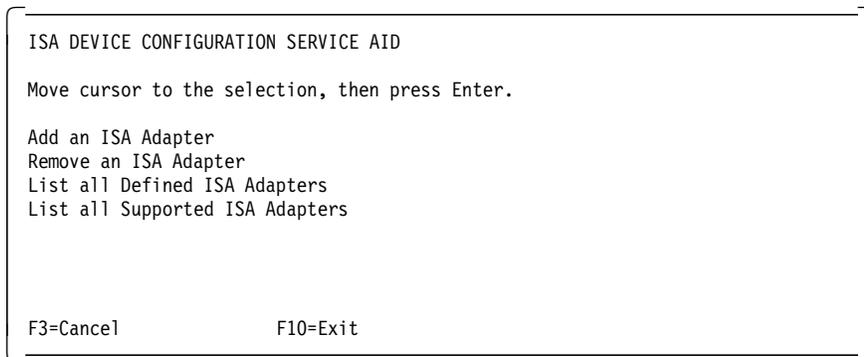


Figure 57. ISA Adapter Service AID Diagnostic Menu

5. From the menu, select **Add an ISA Adapter**.
6. Choose the desired adapter.
7. Enter the parameters, as requested, as shown in Figure 58. For our 8-port card, we only had to enter the Bus I/O Address.

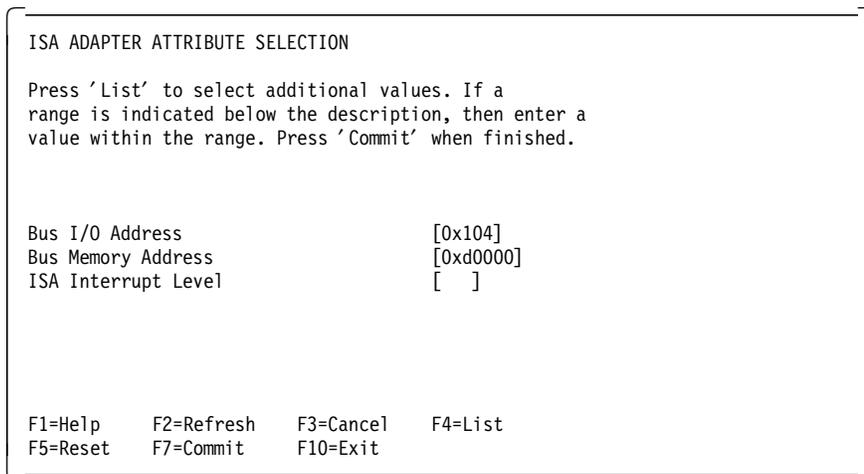


Figure 58. ISA Adapter Attribute Selection Diagnostic Menu

After selecting **Commit**, a message similar to the following will appear:

```

Adding ISA Adapter 'IBM 8-Port Async, EIA-232 (ISA)'
Please Standby

```

If the addition was successful, you should see a message similar to:

```

ISA Adapter device 'sa2'
Configuration complete

```

8. Return to the main diagnostics menu.

You should now be able to test the ISA adapter in addition to other devices in your system.

6.2.2 Running Stand-Alone Diagnostics in Concurrent Mode

There might be a situation where Online 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 which 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.

If you wish to test ISA adapters which have not yet been configured, use SMIT to configure the adapter before running diagnostics.

6.3 Online Diagnostics

The Online Diagnostics are an AIX Version 4 installable package. They are packaged under `bos.diag`. There is a separate, installable option for each device or adapter supported by diagnostics. If you install your system from an AIX Version 4 for Servers CD-ROM, diagnostics will be automatically installed on your system, including diagnostics support for any PCI devices which 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 to add devices after IPL) specifying an installation device, diagnostics support for the device (if available) will be automatically installed. When you install device support for ISA adapters, you can choose whether to install diagnostic support for each adapter. Running Online Diagnostics provides

you with access to the AIX Error Log and the AIX Configuration Data, but again, you are limited to which devices you can test.

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:

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.

Online Diagnostics can be run in the following three modes:

- Concurrent Mode
- Service Mode
- Single-User Mode

These three modes are discussed in the following sections.

6.3.1 Concurrent Mode

The Concurrent mode provides a way to run Online 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 Online 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.

6.3.2 Service Mode

Running diagnostics in Service mode, provides the most thorough test of the system resources available with Online Diagnostics. As in the Concurrent mode, however, there are resources which cannot be tested. These resources are the SCSI adapter and disks used for paging. Any ISA adapters which have been configured to the system with AIX can be tested (providing diagnostics support is installed for the adapters). 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. Reboot your system.
2. When the POST icons are displayed, 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.
3. Enter any requested password (Power-On or Privileged-Access).
4. Follow any instructions to select a console.

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 and follow the instructions displayed to shut down AIX. Failure to shut down AIX could result in damage to data on your disks.

6.3.3 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 your system down into a Single-User mode which will allow testing 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 a message indicates the system is in single-user mode, enter the `diag` command.
5. When the Diagnostic Operating Instructions area is displayed, follow the instructions to check the desired resources.

After you have completed testing the system, you can shut down and reboot the system.

6.4 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 coded in the factory.

The firmware uses the console and the speaker as 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.

The firmware is organized into three major sections:

Boot Block	Receives control at power-on and is able to perform a soft firmware recovery from diskette when necessary
Firmware Core	Brings the system up to where the compressed image has been decompressed and running
Virtual Disk	Contains the firmware modules required to support each hardware subsystem

6.4.1 Firmware Boot Sequence

This section describes the firmware booting process on PCI-based RS/6000. This is also shown on Figure 59.

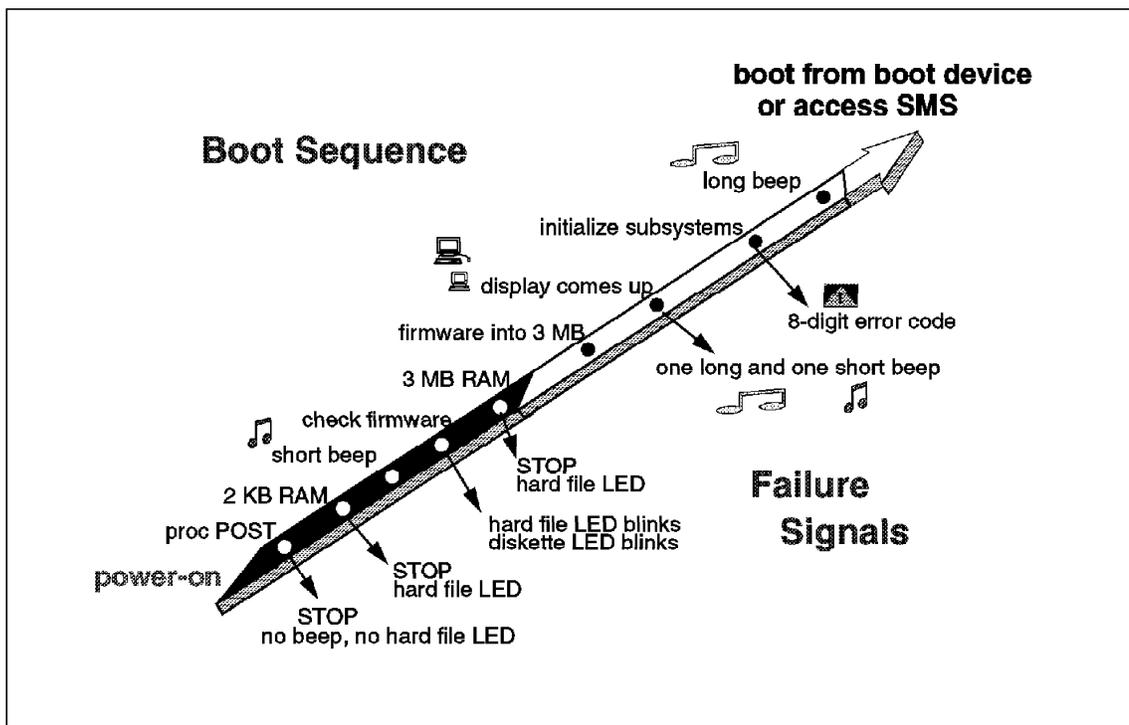


Figure 59. Firmware Boot Sequence

When the system is powered on, the hardware passes control to the storage address 0xFFFF00100, the firmware entry point.

1. The firmware's first task is to perform a processor POST (Power-On Self-Test).

An uncompressed instruction resides on the entry point of the firmware. If the processor fails to execute this branch instruction, the next instruction is executed.

- If it fails, the system stops.
2. The 60x registers are initialized.
 - If it fails, the system stops and appears to be doing nothing, but cooling fans might be observed running.
 3. The firmware initializes the memory controller, and 2 KB of good memory is found.

- If it fails, the hard file LED is turned on, and the system stops.
4. A **short beep** is given on the speaker.
 5. A Cyclic Redundancy Check (CRC) is performed on the compressed firmware.
 - If the firmware is damaged, run the firmware recovery procedure (see 6.4.5, “Firmware Recovery” on page 219).
 6. The firmware core copies a particular part of the firmware code into the small amount of memory. This firmware code runs the memory POST and configures 3 MB of good memory.
 - If it fails, the hard file LED is turned on, and the system stops.

Passed Critical Phase

After 3 MB of good memory has been successfully configured, the system will not stop. Any subsequent failures will be logged and displayed as 8-digit error numbers.

7. The compressed firmware is copied into the larger amount of good memory (3 MB) and uncompressed.
8. A new C environment is set up in the larger area, along with a user stack, an interrupt stack, a debugger stack, and a TOC (Table of Contents) pointer.
9. The individual subsystems are initialized:
 - The Virtual Disk is brought online
 - NVRAM initialization
 - PCI set up
 - Interrupts initialization
 - Real-Time Clock initialization
 - Keyboard/mouse controller initialization
 - Password initialization
 - Debugger initialization
10. The Virtual Disk (VDISK) file system is loaded and initialized.
11. All device drives from the VDISK file system are loaded.
 - Video: initialize ASCII terminal or VGA adapter
 - Initialize Async Console Support (ASCII terminal S1 serial port)
 - Execute the VGA adapter POST. If there is an error, then one **long beep** and one **short beep** is sounded on the speaker.

- Display the “brand” logo on the graphic console or a POST indicator on an ASCII terminal.
 - Keyboard icon appears; initialize the keyboard and mouse.
 - The window for pressing F1 (GUI SMS interface) or F4 (ASCII) keys is now open.
 - The diskette icon appears; initialize the controller.
 - Extended memory icon appears; test the remainder of memory. Memory modules failures are displayed with the 8-digit error code 00020000.
 - The SCSI icon appears; initialize the adapters: SCSI, Ethernet, pci_c970, trpci, tcpip, and initialize network when present.
 - The last icon is the boot disk icon.
12. One **long beep** is sounded on the speaker.
13. If Power-On Password or Privileged Access Password options are in effect, then the Password screen appears. There are three attempts allowed to enter the correct password; otherwise the 8-digit error code 00016005 is displayed, and the system stops. The only way to recover the system is to power-off and power-on again.

6.4.2 Firmware Passwords

The user 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
- Privileged Access Password
- Unattended start mode

Do Not Have the Password?

In case you do not have a machine’s password, the only way to get access to the system is by removing the system’s battery. You must be aware that this procedure will erase all firmware configuration data maintained in NVRAM, such as the error log and any configured IP addresses.

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. The user is given three attempts to enter the correct password. If the user fails 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 the user presses one of the keys to access SMS, the system will check to see if a Privileged Access password exists; if it does, the user is prompted to enter the Privileged Access password. The user is given three attempts to supply the correct password. If the user fails 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 user 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 a user to enter the power-on password.

6.4.3 Firmware Flash Update

There are several situations where a firmware flash update may be required. For example, after a field EC (engineering change) or to provide firmware support for a new adapter.

The firmware flash update can be performed by the user in Normal mode or by the system in Recovery mode. The default is the Normal mode update performed by the user from the SMS Tools (Utilities in the ASCII version) menu. Messages and errors are communicated through the menu pop-up interface.

The Recovery mode update is executed automatically when a defective firmware image is detected during early system initialization (see 6.4.1, “Firmware Boot Sequence” on page 214). In this mode, most messages and errors are suppressed; irrecoverable errors are displayed in a serial-interface-compatible format.

To start a firmware update, you have to perform the following steps:

1. Access SMS.

2. Select the **Tools** menu (**Utilities** in the ASCII version).
3. Select the **Update** icon (**Update System Firmware** on ASCII).

Before starting the firmware update process, the system will ask you if you want to create a *firmware recovery* diskette. It is recommended to create the recovery diskette, unless you have created it already. To create the recovery diskette, insert a blank DOS-formatted diskette.

4. After the recovery diskette has been successfully created, follow the instructions to update the firmware.

To update the firmware, the system searches the diskette drive for a flash update file with the extension *.IMG*. If found, the file is loaded. This file is a composite image. It does not contain the entire firmware; thus, when the firmware is updated, only the sectors contained in the *.IMG* file are programmed.

Once loaded, the image is examined to ensure it is a valid firmware image. Flash table/system ID structure, file length and the Cycle Redundancy Check (CRC) are verified. If we have a valid image, the Vital Product Data (VPD) information is duplicated in the recovery block (for recovery purposes in case the image is damaged) and also merged with the composite image.

Now, the update routine starts programming the firmware sectors to be updated.

Attention

A system power-off in this section of code will result in a damaged flash ROM.

Notice that because only the composite image sectors are erased and reprogrammed those sectors are the only ones that could be damaged. Since the recovery block is protected in Normal mode and Recovery mode, the system can still be restored by using the recovery diskette. A system power-down at the wrong time could result in a completely damaged flash ROM and an unusable system.

6.4.4 Firmware Power-Up Problems

When you receive your system, it is advisable to make a firmware recovery diskette. To do this, perform the following steps:

1. Access the System Management Services.
2. Select Option 3, **Utilities**.

3. Select Option 7, **Update System Firmware**. The system will display a message recommending that you create a firmware recovery diskette.
4. Insert a DOS formatted, write-enabled blank diskette, and press Enter.

If your firmware becomes corrupted, for example if you experience a power failure during the update of the firmware, you can use the firmware recovery diskette to repair the system.

To start the recovery operation, insert the firmware recovery diskette into the diskette drive and power-on the system. The process should start automatically. If the system does not restart after a few minutes, the system planar may have to be replaced. There are no other means of recovery other than the recovery diskette.

6.4.5 Firmware Recovery

During the firmware boot process, one of the first tasks is to verify the Cycle Redundancy Check (CRC) of the firmware code. If the update portion of the firmware has a CRC miscompare (is damaged), you will see the following symptoms:

- One **short beep** will be sounded on the speaker.
- The hard disk LED will blink continuously.
- The diskette drive LED is selected or blinking, waiting for a diskette to be inserted.
- The display or tty is blank.

Recovery Procedure: To rewrite the flash memory, you will need a **firmware recovery diskette**.

Note

The recovery diskette must be made in advance at the customer location using the "Update System Firmware" utility provided with the System Management Services (SMS) program.

If you have a firmware recovery diskette, perform the following steps:

1. Insert the firmware recovery diskette into the diskette drive. The diskette LED should turn on while data is being copied into memory.
2. The hard file LED should stop blinking. The system will retry forever when something is wrong.
3. The system will reboot the machine when the recovery is successful.

The firmware flash routine performs some general clean-up and reboots the system. Once the system is rebooted, the new flash image is active and executed.

6.4.6 Power-On Self Test

The following power-on self tests (POSTs) are performed by the firmware on PCI-based RS/6000 systems:

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.

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 good 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 or not).

Graphics Adapter POST

The firmware does test 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 a one long beep, followed by one short beep. The error is logged, and the system boot continues.

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.

SCSI Subsystem POST

The SCSI POST opens the adapter and runs several test scripts which check and exercise the various functions related to the chips on the integrated SCSI adapter, such as DMA, simulate 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 *User's Guide* for the corresponding PCI-based RS/6000 system.

6.5 Function Keys - Key-Switches at Boot Time

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 which determine the type of boot which 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 **1** 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 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** or **8** 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.

6.6 System Management Service (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 which can be used on a graphical display and an ASCII interface which 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 the computer.
2. When the first screen appears, press the **F1** key to use the GUI interface. To use the ASCII interface, press the **F4** key. On an ASCII terminal, you can press either the number **1** or the number **4** key respectively.

Note: You must press keys listed above before the last indicator appears.

6.6.1 SMS Graphical Main Menu

The graphical System Management menu offers the following services:

- | | |
|-----------------|--|
| Start Up | Allows you to define a customized boot list which defines the sequence in which devices are searched by the system for a bootable image. |
| Test | The only test supported here is the memory test. Testing other devices can be performed by using the diagnostics tools as explained in Chapter 6, "Diagnostics and System Configuration Management" on page 205. |
| Tools | 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. |

6.6.2 SMS ASCII Main Menu

The ASCII version of the System Management Services includes an additional function not found in the graphical version, support for the following languages:

- English
- French

- German
- Italian
- Spanish
- Swedish

All menus and options, as well as the button titles, have been translated.

6.6.3 Start-Up Menu

The corresponding menu to the graphical Start Up menu in the ASCII SMS version is the Select Boot Devices menu. The Start Up menu enables you to define a custom boot list which determines the order in which devices are searched by the operating system for a bootable image.

When this menu shows up, the current boot-order list is displayed. The **Default** icon can be used to restore the default settings, or the **Select** icon can be used to configure a new boot-order list. The default settings are:

1. Diskette drive
2. CD-ROM drive
3. Hard disk drive
4. Network (if any network adapter installed)

To configure a device as a boot device, you can use the mouse and click on the device, or use the arrow keys to highlight the device and press the **Spacebar** to select it. A number indicating the order of the startup sequence will appear on the device.

To cancel the configured selections, the **Cancel** icon can be used.

The ASCII interface offers an additional function (option 7. Boot Other Device) which can be used to boot directly from the specified device without having to change the boot-order list.

6.6.4 Test Menu

The corresponding menu to the graphical Test menu in the ASCII SMS version is the Test the Computer menu. The Test menu allows you to test the system memory.

To begin the test, select the **Test** icon. The **Loop Test** icon, when selected, allows the memory to be tested in a continuous loop. The test can be stopped by pressing the **S** key. If a problem is detected during testing, an error code similar to the one below will be displayed:

```
00020000
```

6.6.5 Tools Menu

The corresponding menu to the graphical Tools menu in the ASCII SMS version is the Utilities menu. The Tools menu contains the following utilities:

Power-on	Allows you to set or remove a power-on password.
Super	Allows you to set or remove a Privileged Access password. Privileged Access means access to the System Management Services program.
Remote on	When enabled, the system will boot from the defined boot devices without asking for the power-on password. You must set a power-on password before you can enable the remote mode. Although the system can be booted, the console keyboard will be locked. This mode is ideal for network servers and other computers that operate unattended.
Update	Enables you to update the system firmware and to create a firmware recovery diskette.
Error log	Enables you to view the error log information that is stored if an error occurs.
R IPL	Allows you to select a remote system to load a boot image from. This option must be selected when you want to boot from a NIM Master for installation or maintenance purposes.

6.6.6 Accessing System Management Services with an ASCII Terminal

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 X 80 columns
- Echo mode on

6.7 Service Processor

The service processor shares access to ports S1 and S2 with the server. Thus, the serial ports can be accessed even when the system is powered off (providing there is power to the service processor). Terminals and modems may be attached to these ports, making it possible to perform a number of administrative tasks remotely. The system can also be booted remotely, and a remote user may enter the Systems Management Services (SMS) screens (ASCII only).

The main purposes of the service processor are:

- To enable easier remote administration and maintenance
- To allow for remote power-off and power-on
- To control the system boot process

There are two methods through which the service processor's functions can be customized. The first is a Service Aid in AIX Diagnostics. Although accessible using concurrent AIX diagnostics, the SP functions you can customize there are relatively few. The second way of configuring the service processor is through the SP menus that are only accessible — either locally or remotely — when AIX is not running.

6.7.1 Service Aid Interface

The Service Aid can be accessed from either an ASCII or a graphical terminal by entering `diag` on the AIX command line or by booting off a diagnostics CD-ROM and going into Service Aids.

On the Service Aid Task Selection screen, the administrator has two options: Configure Ring Indicate Power On (RIPO) or Configure Service Processor. These tasks are mutually exclusive: configuring Ring Indicate Power On voids the remote support options of the service processor. The RIPO feature lets you set a number of telephone rings that trigger a system boot in a machine that is powered off. No service processor session is established; so there's no way to immediately confirm that the boot was successful. The call is NOT answered. You will not know the system has booted until you can log in. As a safeguard, if the system is already powered on, the ring will have no effect whatsoever. It is possible to invoke this feature from the command line as long as the `/usr/lpp/diagnostics/bin` subdirectory is in your path. The command is `usp -r`. A menu appears that

turns RIPO on or off and sets the number of rings. The number of rings must fall between 1 and 255.

Other options that can be configured in Service Aids include:

- Enable/Disable Console Mirroring
- Select Modem Line Speed
- Enable/Disable Modem
- Set Up Modem Configuration
- Set Up Dial-Out Phone Numbers
- Set Up Surveillance
- Update SP Flash EPROM

To use the Service Aid menus, the operator must have the AIX root password.

Console mirroring enables someone at a remote location to see the same screen information as the person at the machine. Console mirroring is only possible between ASCII devices that are attached to ports S1 and S2. Thus, if a modem were attached to S1, the customer would need an IBM 3153 or similar terminal attached to the S2 port to implement console mirroring. Console mirroring occurs on the J50 and R50 at run time as well as when the Bring Up Micro Processor (BUMP) is in control of the system. The PCI servers do not support AIX run-time console mirroring through the service processor. Console mirroring may only occur before AIX has been loaded. Do not confuse this with the portmir command that mirrors LFT devices when AIX is running.

The Modem Line Speed should be set to 9600 baud, and the modem itself must be enabled. Under Setup Modem Configuration, you only need to supply two things, the name of a modem description file and the port to which the modem is connected. Sample modem description files are shipped with AIX and stored in `/usr/share/modems`. The *User's Guide*, SA38-0540, gives you guidance on which file to use for several well-known modems. If configuration is successful, a message appears saying that the modem configuration file is being loaded into NVRAM. When the message goes away, the process is complete.

The Setup Dial-out Phone Numbers menu lets you determine whether ports S1 and S2 are accessible to a remote user and what the line speed will be. Its recommended that you remain at 9600 baud because that is the best speed for accessing SMS remotely. The SP can be customized to include information such as the service center phone number or an administrator's

pager number. The SP can use these to alert support personnel to a problem.

Surveillance uses the heartbeat method of ensuring that the system will reboot automatically if it detects a hang. There are three things to note on this screen: surveillance is either turned on or off (the default is off); the frequency of the heartbeat is determined in minutes (from 2 to 128); and the delay before starting the heartbeat can be a number between 0 and 128. Changes made in surveillance take effect at the next system boot.

Update SP Flash EPROM lets you update the Service Processor firmware through diagnostics. You should always be sure that you're using the right level of firmware for your system and that you pay close attention to any README files before applying new firmware.

6.7.2 Service Processor Menus

All of the SP Service Aid tasks can be performed by using the SP menus. The SP menus are accessible from either the S1 or S2 ports using a modem or a TTY. When OK appears in the systems Liquid Crystal Display (LCD), pressing any key will bring up the SP main menu on the system that is powered off. The menu should appear between 15 and 45 seconds after shutdown, depending on whether this is the first time the service processor has been initialized.

Options on the native SP menus include all those found under Service Aids, plus:

- Read Vital Product Data
- View System Environmental Conditions
- Read System POST Errors
- Read SP Error Logs
- View Progress Indicators from Last Boot
- Power-On System
- Power-Off System
- Read NVRAM
- Reset SP
- Setup Reboot/Restart Policy
- Start Talk Mode
- Change General User/Power-On Password
- Change Privileged Access Password

- Select Language
- Enable/Disable Unattended Start Mode

Most of these functions are self-explanatory. However, a few might be worthy of special attention.

The Read System POST Errors and Read SP Error Logs prove to be very useful in problem determination/problem source identification. These are errors separate from any that may be logged in AIX. The POST errors may help determine why a system is having trouble booting. In that connection, being able to View Progress Indicators from Last Boot, including AIX checkpoints, also can be valuable. The administrator can turn on unattended start mode on the Setup Reboot/Restart Policy menu.

Start Talk Mode is for use in console mirroring sessions. It enables the local and remote support people to communicate. This is a full duplex link; so message interference is possible. Alternating between users works best.

Some SP menus can be accessed with only the general access (power-on) password. They are associated with reading vital product data, looking at error logs and powering the system on and off. Most of the more sensitive functions require the operator to know the privileged access password. These passwords can be set in SMS or in the SP menus.

The following functions can also be performed from SMS: read VDP, update the SP Flash EPROM, change passwords, select a language, and enable/disable unattended start mode.

Chapter 7. AIX Common Hardware Reference Platform Support

AIX provides support for specific features of the PCI-based RS/6000 servers. Not all AIX commands or parameters are equal when compared with Micro Channel-based RS/6000 systems. There are also differences between the PReP machines, like the E30 and F40, and the CHRP machines, like the F50 and H50.

This chapter contains information for CHRP-compliant hardware. Information about AIX support of PReP machines can be found in the following two redbooks:

- *A Technical Introduction to PCI-Based RS/6000 Servers*, SG24-4690
- *Managing AIX V4 on PCI-Based RISC System/6000 Workstations (40P/43P)*, SG24-2581.

The following topics are covered in this chapter:

- Soft “key-switch” function
- Debug mode
- Backup and cloning
- Configuration Information (lscfg)
- Network boot support
- Error logging

Refer also to Chapter 5, “Adapters, Overview and Configuration” on page 147, for additional information about ISA and PCI adapter configuration support.

7.1 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 which can be in one of three positions:

- Normal
- Service
- Secure

Although the E30, F40, F50, and H50 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.

7.1.1 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  
  
fd0  
cd0  
rmt0  
hdisk0  
ent0
```

One other point to note is that the RSPC machines cannot use the generic terms such as `scdisk` or `rmt` for devices. All the devices must be explicitly stated, for example `rmt0`. Trying to use generic options will result in an error:

```
bootlist -m normal scdisk  
  
0514-202 bootlist: Invalid device name (scdisk)
```

There are two other ways to alter the boot lists: System Management Services (SMS) and the AIX tool `diag`.

7.2 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 E30, F40, F50, and H50 have a static electronic key which 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.

7.3 Backup and Cloning

This section deals with creating system backups which will be restored on the same or different machines (cloning).

7.3.1 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 which 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 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.

Platform	Description
chrp	CHRP-compliant PCI-based RS/6000 machines
rspc	PReP-compliant PCI-based RS/6000 machines
rs6k	Micro Channel-based RS/6000 machines

7.3.2 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 which 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
```

7.3.3 Restoring Your System Backup

You can restore your `mksysb` by booting from the tape which you have created and following the instructions which 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

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 uniprocessor or Micro Channel-based symmetric multiprocessor machines. Additionally, `mksysb` tapes created on Micro Channel-based uniprocessor and Micro Channel-based symmetric multiprocessor machines will not boot on PCI-based machines.

7.3.4 Cloning

Restoring a `mksysb` backup on a different machine is called *cloning*. 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.

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.

In AIX 4.1, the devices were 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, etc) has made cloning machines almost impossible. Up to this point the only method of cloning was to install all the devices on a machine before you created a `mksysb` image. This is a waste of space because you do not need all the devices, but you do not know what devices you will need until you are actually installing.

There is a new way to do this, 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 development of reproducing problems.

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` fileset 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` fileset 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.

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 -lhd5 | 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
```

On the System You Want to Install

Boot the target machine (the one you want to install on) with the AIX 4.1.4 CD-ROM product media and insert the diskette in the drive. Make sure the tape drive is turned on, but do not insert the tape yet as it might try to boot

from the tape. Press the F1-Key when the console screen comes up; then press Enter when the language screen comes up (to get English). Then, when the:

```
Welcome to Base Operating System
Installation and Maintenance
```

screen comes up, choose option **3**:

```
3 Start Maintenance Mode for System Recovery.
```

Then choose option **4**:

```
4 Install from a System Backup
```

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 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.

7.3.5 Cloning on AIX Version 4.2 and 4.3

You do not need a separate diskette or allocated customization script to clone a system in 4.2. 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.

Set the new variable, SWITCH_TO_PRODUCT_TAPE, to yes to allow bi_main to prompt you to insert the product tape again. If SWITCH_TO_PRODUCT_TAPE is not set to yes, then the mksysb tape will finish installing and reboot without installing additional device support from the product tape. NIM install will also do cloning automatically, from the lpp_sources allocated to the install process. There are three new variables to the bosinst.data file's control_flow stanza that will help with cloning:

```
RECOVER_DEVICES = yes, no          (default - yes)
BOSINST_DEBUG   = no, yes          (default - no)
SWITCH_TO_PRODUCT_TAPE = no, yes   (default - no)
```

The RECOVER_DEVICES variable basically gets you what the cloner script did, which was to remove the /etc/firstboot file so the rda (recover device attributes) program, which tries to reconfigure the devices and network attributes that were present on the original system, will not run. Setting RECOVER_DEVICES to no will prevent rda from running.

Cloning Considerations

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. If you are installing a mksysb on a system it was not created on, use the procedure in "Cloning Your System."

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 Welcome to the Base Operating System Installation and Maintenance screen is displayed.

1. Select the **Start Maintenance Mode for System Recovery** option.
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 nonprompted 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
```

- b. 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`.
- c. 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 `PEX/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.

7.4 Configuration Information

Use `lscfg` command to display 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 an F50, 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      00-00      System Planar
tok0            30-60      IBM PCI Tokenring Adapter (14101800)

Loadable Microcode Level.....WW1B71X
Part Number.....00000000
EC Level.....00D51237
Serial Number.....00062738
FRU Number.....00000000
Manufacturer.....IBM982
Network Address.....08005AB936F6
Displayable Message.....IBM PCI Token-Ring Adapter (14101800)

pci0            00-fef00000  PCI Bus
isa0            10-58       ISA Bus
```

fda0	01-D1	Standard I/O Diskette Adapter
fd0	01-D1-00-00	Diskette Drive
siokma0	01-K1	Keyboard/Mouse Adapter
sioka0	01-K1-00	Keyboard Adapter
sioma0	01-K1-01	Mouse Adapter
ppa0	01-R1	Standard I/O Parallel Port Adapter
sa0	01-S1	Standard I/O Serial Port
tty0	01-S1-00-00	Asynchronous Terminal
sa1	01-S2	Standard I/O Serial Port
sa2	01-S3	Standard I/O Serial Port
scsi0	10-60	Wide SCSI I/O Controller
cd0	10-60-00-4,0	SCSI Multimedia CD-ROM Drive (650 MB)

Manufacturer.....IBM
 Machine Type and Model.....CDRM00203
 ROS Level and ID.....0.D5
 Device Specific.(Z0).....058002028F000018
 Part Number.....93H8050
 EC Level.....E77081
 FRU Number.....93H8055

rmt0	10-60-00-5,0	5.0 GB 8 mm Tape Drive
------	--------------	------------------------

Manufacturer.....EXABYTE
 Machine Type and Model.....IBM-8505
 Device Specific.(Z1).....7Z0A
 Part Number.....8191044
 Serial Number.....693505
 Device Specific.(LI).....00000001
 EC Level.....D48098
 FRU Number.....59H3159
 Device Specific.(Z0).....0180020283000010

hdisk0	10-60-00-8,0	16 Bit SCSI Disk Drive (4500 MB)
--------	--------------	----------------------------------

Manufacturer.....IBM
 Machine Type and Model.....DCHS04U
 FRU Number.....76H2697
 ROS Level and ID.....32393239
 Serial Number.....681181A2
 EC Level.....488638
 Part Number.....93G3159
 Device Specific.(Z0).....000002029F00003E
 Device Specific.(Z1).....RAMUR26D
 Device Specific.(Z2).....09RI
 Device Specific.(Z3).....97234
 Device Specific.(Z4).....0001

```

Device Specific.(Z5).....22
Device Specific.(Z6).....E29611

ses0          10-68-00-15,0      SCSI Enclosure Services Device

Manufacturer.....IBMAUSTN
Machine Type and Model.....HSBP06AUSYSTEM
ROS Level and ID.....0117
Device Specific.(Z0).....0D0003023B004100
Device Specific.(Z1).....FN 93H2166

ent0          10-80            IBM PCI Ethernet Adapter (22100020)

Network Address.....08005ABAC777
Displayable Message.....PCI Ethernet Adapter (22100020)

mem0          00-00            Memory
proc0         00-00            Processor
L2cache0     00-00            L2 Cache
pmc0         00-00            Power Management Controller
PLATFORM SPECIFIC

```

```

Name: IBM,7025-F50
Model: IBM,7025-F50
Node: /
Device Type: chrp

```

```

System:
Machine/Cabinet Serial No...
Machine Type/Model.....7025-F50
Manufacture ID.....IBM966
Version.....RS6K

```

```

I/O Planar:
Part Number.....93H8783
EC Level.....D30105B
Serial Number.....154W168005
FRU Number.....07L6594
Manufacture ID.....1053
Version.....RS6K

```

```

System Planar:
Part Number.....93H6436
EC Level.....D74515B
Serial Number.....17UT148118
FRU Number.....73H1925
Manufacture ID.....IBM053
Version.....RS6K

```

Product Specific.(ZM).....MF=1,MV=01

Op Panel:

FRU Number.....Op Panel
Version.....RS6K
System Info Specific.(YL)...L1

Memory Card:

FRU Number.....Memory Card
Version.....RS6K
System Info Specific.(YL)...P1-M1

Orca M5 CPU:

Part Number.....94H0039
EC Level.....E77410
Serial Number.....157Y350WSU
FRU Number.....93H8945
Product Specific.(MF).....IBM980
Version.....RS6K
Displayable Message.....OrcaM5 CPU DD1.2
Product Specific.(ZC).....PS=0013c9eb00,PB=0009e4f580,SB=0004f27
ac0,NP=02,PF=461,PV=05,KV=01,CL=1

Name: memory-controller

Model: IBM, Cobra-Viper, Rev-id 41
Node: memory-controller@ff000000
Physical Location: P1

Name: IBM,sp

Node: IBM,sp@ie8
Physical Location: P2-X1

SP_CARD_:

Part Number.....93H4228
EC Level.....E76324A
FRU Number.....93H214
Manufacture ID.....3966-1944843
Serial Number.....00000143
ROM Level.(non-alterable)...19971208
Product Specific.(PL)....._01_
ROM Level (alterable).....19971208

The command gives very specific information about the processor card and the service processor.

Orca M5 CPU

Information is given about part numbers and so forth. In this instance, detailed information is given on the following:

- The Processor Speed 332 MHz is given as a Hex-number (PS=0x0013c9eb00=332000000 Hz).
- The Speed of the processor bus is 166 MHz (PB=0009e4f580=166000000 Hz).

IBM,sp

This stanza gives information about the support processor.

7.5 Network Boot Support

If you are booting an over the network, follow the steps:

1. Begin with your machine turned off.
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.

7.6 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 which 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. Combining this problem determination with the Model F40, F50 or H50's hot swap subsystem could eliminate any system down time.

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:

```
errclear 0
```

7.7 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.

7.7.1 Booting from CD-ROM

While AIX supports installation from external CD-ROM, the CD-ROM drive, 7210-001, cannot be used as a boot device.

7.7.2 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.
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
```

In addition, check that the files can be executed without problems, for example:

```
/usr/bin/sh /etc/environment
```

9. Re-create 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.

7.7.3 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 `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. Special Notices

This publication is intended to help technical and non-technical professionals support, configure, and manage the RS/6000 Model E30, F40, F50, and H50. 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 Model E30, RS/6000 Model F40, RS/6000 Model F50, and RS/6000 Model H50 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.

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®	AIX/6000®
AS/400®	BookManager®
DB2®	DirectTalk®
graPHIGS	GXT1000
HACMP/6000	IBM®
IBMLink	InfoExplorer
LANStreamer®	Micro Channel®
Portmaster®	POWER Architecture
Power PC 603®	Power PC 604®
Power Series®	PowerPC®
PowerPC Architecture	PowerPC Reference Platform®
PowerPC 601®	PowerPC 603
PowerPC 604	PowerPC 630
POWER2 Architecture	PS/2®
RISC System/6000®	RS/6000
Serial Storage Architecture®	Service Director®
System/390®	TURBOWAYS®
Ultimedia®	

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 B. Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see "How to Get ITSO Redbooks" on page 257.

- *A Practical Guide to Serial Storage Architecture for AIX*, SG24-4599
- *Implementing PC ServeRAID SCSI and SSA RAID Disk Subsystem*, SG24-2098
- *A Technical Introduction to PCI-Based RS/6000 Servers*, SG24-4690
- *RS/6000 SMP Enterprise Servers Architecture and Implementation*, SG24-2583
- *Managing AIX V4 on PCI-Based RISC System/6000 Workstations (40P/43P)*, SG24-2581
- *RS/6000 AIXLink/X.25 Cookbook*, SG24-4475

B.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

B.3 Other Publications

These publications are also relevant as further information sources:

- *Adapters, Devices, and Cable Information for Multiple Bus Systems (Supplement)*, SN32-9058

- *Adapters, Devices, and Cable Information for Multiple Bus Systems*, SA23-2778
- *IBM RS/6000 Adapters, Devices, and Cable Information for Multiple Bus Systems*, SA38-0516
- *Diagnostic Information for Multiple Bus Systems*, SA38-0509
- *RS/6000 7025 F50 Series Service Guide*, SA38-0541
- *RS/6000 7025 F50 Series User's Guide*, SA38-0540
- *RS/6000 7025 F40 Series Service Guide*, SA38-0515
- *RS/6000 7025 F40 Series User's Guide*, SA38-0514
- *RS/6000 7024 E Series Service Guide*, SA38-0502
- *RS/6000 7024 E Series User's Guide*, SA38-0501
- *RS/6000 Enterprise Server Model H50 Installation and Service Guide*, SA38-0547
- *RS/6000 Enterprise Server Model H50 User's Guide*, SA38-0546
- *RS/6000 Facts and Features*, G320-9878
- *RS/6000 Workgroup Server Model E30*, G221-7031
- *RS/6000 Workgroup Server Model F40*, G221-7017
- *RS/6000 Model F50*, G221-7029
- *RS/6000 Enterprise Server Model H50*, G221-7051
- *RS/6000 High Availability Cluster Server Model HA50*, SA38-0551 (In press, May 1998)
- *PowerPC Microprocessor Common Hardware Reference Platform: A System Architecture*, ISBN 1-55860-394-8
- *AIX Version 4.3 Installation Guide*, SC23-4112
- *AIX Version 4.3 Network Installation Management Guide and Reference*, SC23-4113

B.4 Web Resources

These links will provide you with a stimulating look at UNIX server technology.

- <http://www.tpc.org>
- <http://www.specbench.org>
- <http://www.netlib.no/netlib/benchmark/performance.ps>

- <http://1394TA.org>
- <http://www.chips.ibm.com/products/ppc/>
- <http://aix.boulder.ibm.com/aix.wa/aixfixes>
- <http://www.rs6000.ibm.com/support/micro>
- <http://pitstop.austin.ibm.com/>
- http://pitstop.austin.ibm.com/show_html_fax.sh?clone
- <http://www.us.pc.ibm.com/healthycomputing>

B.5 Tools Disks

IBM employees can get the following packages from MKTTOOLS.

- TOOLS SENDTO USDIST MKTTOOLS GET SSAINST PACKAGE
- TOOLS SENDTO USDIST MKTTOOLS GET SSAFIBRE PACKAGE
- TOOLS SENDTO USDIST MKTTOOLS GET SSAPRFP PACKAGE

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:

- **PUBORDER** — to order hardcopies in United States
- **GOPHER link to the Internet** - type GOPHER.WTSCPOK.ITSO.IBM.COM
- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
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
```

For a list of product area specialists in the ITSO: type the following command:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Web Site on the World Wide Web**
<http://w3.itso.ibm.com/redbooks/>
- **IBM Direct Publications Catalog on the World Wide Web**
<http://www.elink.ibm\link.ibm.com/pb1/pb1>

IBM employees may obtain LIST3820s of redbooks from this page.

- **REDBOOKS category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL
- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm\link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

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	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States) or (+1)001-408-256-5422 (Outside USA)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services** - send note to softwareshop@vnet.ibm.com

- **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

- **Internet Listserver**

With an Internet e-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an e-mail note to announce@webster.ibm.link.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank).

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

AC	alternating current	DIP	direct insertion probe
AC/DC	alternating current/direct current	DMA	direct memory access
AIX	advanced interactive executive (IBM's flavor of UNIX)	DVD	digital video/versatile disk
AMD	Advanced Micro Devices, Inc.	ECC	error checking & correction
ANSI	American National Standards Institute	EISA	extended industry standard architecture
APAR	authorized program analysis report	EL	electro-luminescence
ASCII	American National Standard Code for Information Interchange	EMI	electromagnetic interference
ATM	asynchronous transfer mode	ETHERNET	a communications network (USA, Xerox corp, 1975)
BIOS	basic input/output system (PC and PS/2)	FDDI	fiber distributed data interface (100 Mbit/s fiber optic LAN)
BOS	base operating system	FORTRAN	formula translation (programming language)
CD	compact disc	FRU	field replaceable unit
CD-ROM	(optically read) compact disk - read only memory	GB	gigabyte (10 ⁹ bytes or 1,000,000,000 bytes)
CDE	Common Desktop Environment (from X/Open)	GUI	graphical user interface
CHRP	Common Hardware Reference Platform (PowerPC)	HA	hardware acceleration
CMOS	complementary metal oxide semiconductor	HACMP	high availability cluster multi-processing (AIX)
CPU	central processing unit	HW	hardware
CRC	cyclic redundancy check	I/O	input/output
DC	direct current	ID	identification number
DEC	Digital Equipment Corporation (USA)	IEEE	Institute of Electrical and Electronics Engineers
DIMM	Dual Inline Memory Module	IP	internetwork protocol (OSI)
		IPL	initial program loader
		IRQ	interrupt request

ISA	industry standard architecture	OLTP	online transaction processing
ISBN	international standard book number	P/N	part number
ISDN	integrated-services digital network	PC	personal computer
ISO	International Organization for Standardization	PCI	peripheral component interconnect (Intel bus standard)
ITSO	International Technical Support Organization	PCM	pulse code modulation
JTAG	Joint Test Action Group	PHIGS	programmers hierarchical interactive graphics standard (three-dimensional)
LAN	local area network	POP	power-on password
LCD	liquid crystal display	POST	power-on self-test
LED	light emitting diode	POWER	performance optimization with enhanced RISC (architecture)
LRU	least recently used		
M/T	machine type		
MAC	medium access control (token-ring)	PowerPC	PowerPC, a family of single-chip RISC microprocessors (Apple Computer, Inc., IBM & Motorola, Inc.)
MB	megabyte, 1,000,000 bytes (1,048,576 bytes memory)		
MCAD	mechanical computer-aided design	PTF	program temporary fix
MEM	memory	PVC	permanent virtual circuit
MES	miscellaneous equipment specification (IBM field installed changes to machines)	RAID	redundant array of independent disks
MESI	Modified, Exclusive, Shared, Invalid (cache consistency protocol)	RAM	random access memory
MMF	multimode fiber	RAS	reliability, availability, serviceability
MP	multi processor	RFA	release for announcement
MVP	medium voltage process	RISC	reduced instruction set computer/cycles
NIM	network install manager	ROM	read only memory
NVRAM	non volatile random access memory	SCSI	small computer system interface
ODM	object data manager (AIX)	SDRAM	static dynamic random access memory (DRAM with integrated refresh logic)

SMIT	System Management Interface Tool (see also DSMIT)	STP	shielded twisted pair (cable)
SMOL	sales manual on line	SVC	switched virtual circuit
SMP	symmetric multiprocessors	TOC	table of contents
SMS	software management systems	TPC-C	Transaction Processing Council - C (a benchmark)
SP	service processor	UNIX	an operating system developed at Bell Laboratories (trademark of UNIX System Laboratories, licensed exclusively by X/Open Company, Ltd.)
SPEC	System Performance Evaluation Cooperative (development of computer industry standard benchmarks)	UP	uni-processor
SPI	SCSI Parallel Interface	USB	Universal Serial Bus
SSA	serial storage architecture		

Index

Special Characters

/usp -r, command 226

Numerics

100BaseTx Ethernet 148

10BaseT Ethernet 148

128-port asynchronous adapter

configuring 188

DIP-switch setting 187

general information 157

2-port multiprotocol adapter 157

2.1 GB Hard Disk 99

2.2 GB Hard Disk 99

2443 6-pack SCSI to DASD cable and terminator 60

2444 6-pack SCSI to DASD cable 60, 61, 85

2446 PCI Ultra-SCSI to 6-pack cable 61

2447 6-pack PCI SCSI SE to DASD cable 61, 85

2450 6-pack SSA to PCI SSA adapter cable 63

2452 SSA blank disk module 63, 118

2453 SSA internal to SSA external cable 63

2493 PCI RAID adapter 47, 59, 61, 77, 85, 86, 100, 102, 114

2618 8x CD-ROM 46, 99

2619 20x max. CD-ROM 46, 99

2630 1/4 inch internal tape 66

2638 PCI video capture card 64

2701 4-port multiprotocol ISA adapter

feature number 49, 64, 78, 87, 155, 191

2705 4-port multiprotocol cable 155, 191

2708 Eicon ISDN DIVA PRO 150

2741 SysKonnnect FDDI-LP adapter (SAS) 50, 65, 78, 87, 150, 151, 154

2742 SysKonnnect FDDI-LP adapter (DAS) 50, 65, 78, 87, 150, 151, 154

2743 SysKonnnect FDDI-UP adapter (SAS) 50, 65, 78, 87, 150, 152, 154

2837 MVP multi-monitor adapter 158, 166

2838 GXT120P graphics adapter 86, 159, 166

2839 GXT110P graphics adapter 159, 166

2851 GXT250P graphics adapter 160, 166

2852 GXT255P graphics adapter 160, 166

2853 GXT800P graphics adapter 162, 168

2854 GXT500P graphics adapter 161, 166

2855 GXT550P graphics adapter 168

2859 GXT800P adapter with texture 162, 168

2900 4.5 GB Ultra SCSI hot swap drive 45, 61

2908 9.1 GB Ultra SCSI hot swap drive 45

2913 9.1 GB Ultra SCSI hot swap drive 84

2919 9.1 GB Ultra SCSI hot swap drive

select 84

2920 PCI token ring adapter 64, 78, 86, 148

2931 8-port asynchronous adapter 155, 156, 183

2932 8-port async. adapter EIA 232/RS-422 156

2933 128-port asynchronous controller 156, 186

2943 8-port asynchronous adapter PCI 156

2944 128-port asynchronous adapter PCI 157

2961 X.25 ISA co-processor adapter 49, 64, 78, 87, 149, 197

2962 2-port multiprotocol PCI adapter 49, 64, 78, 87, 157

2963 TURBOWAYS 155 PCI UTP ATM adapter 49, 64, 78, 87, 153

2968 10/100 Mbps Ethernet PCI adapter 49, 64, 78, 86, 148

2979 PCI token ring adapter 49, 64, 78, 86, 147

2985 PCI Ethernet BNC/RJ-45 adapter 49, 64, 78, 86, 149

2986 3Com 10/100 Mbps PCI Ethernet adapter 49, 64, 78, 86, 148, 150, 154

2987 PCI Ethernet AUI/RJ-45 adapter 49, 64, 78, 86, 149

2988 TURBOWAYS 155 PCI MMF ATM adapter 49, 64, 78, 87, 150, 153, 154

2998 TURBOWAYS 25 ATM PCI adapter 49, 64, 78, 87, 154

2D graphics adapters 158

3005 4.5 GB Ultra-SCSI drive select 61, 62, 77

3019 9.1 GB F/W Ultra-SCSI drive select 84

3033 2.2 GB SCSI F/W hot swap drive 45

3071 4.5 GB SSA hot swap drive 62, 77, 118

3072 9.1 GB SSA hot swap drive 62, 63, 77, 84, 118
 3095 2.1 GB SCSI hard drive 45
 3Com Ethernet adapter 148
 3D graphics adapters 158
 4-port adapter
 cable selections 191
 configuring 194
 DIP-switch setting 192
 4.5 GB Hard Disk 99
 5500 fiber optic extender 121
 6-Pack, F40 54
 604e design 31
 60X bus 29
 6137 4 mm internal tape 66
 6142 4/8 GB internal tape drive 46, 100
 6147 5/10 GB 8 mm internal tape 100
 6154 20/40 GB 8 mm internal tape 100
 6159 12/24 GB internal tape drive 46, 100
 6206 PCI single-ended Ultra SCSI adapter 47, 48, 59, 61, 62, 76, 85, 86, 100, 102, 106, 108
 6207 PCI differential Ultra SCSI adapter 47, 48, 59, 62, 77, 85, 86, 100, 102, 111
 6208 PCI SCSI-2 F/W SE adapter 47, 48, 59, 60, 76, 85, 100, 101, 102, 106, 108, 109
 6209 PCI SCSI-2 F/W differential adapter 47, 48, 59, 77, 85, 100, 101, 102, 111
 6215 PCI SSA multi-initiator RAID adapter 62, 77, 86, 119, 120
 6218 PCI SSA 4-port RAID adapter 63
 6222 SSA fast write cache option card 119, 120
 6296 redundant A/C power supply 88
 6309 digital trunk quad PCI adapter 86, 157
 6424 boot drive mounting kit (H50) 85
 6427 boot drive mounting kit (H50) 86
 6509 media-bay mounting kit 61
 6520 first six SCSI hot swap backplane 60
 6523 SCSI 6-pack non hot swap backplane 60, 61
 6527 F40 SSA base backplane 58
 6537 SCSI 6-pack hot swap backplane 60, 61
 6538 SSA select backplane 62, 77
 6539 SSA backplane power cable 62, 77, 86
 6549 additional power supply for SSA 61, 62
 6XX bus 29
 7024-E30, overview 42
 7025-F40, overview 52
 7025-F50, overview 66
 7026-H50, overview 79
 7131 subsystem 123
 7133 subsystem 122
 7252 GXT1000 internal graphics
 accelerator 164, 168
 7253 GXT1000 advanced graphics option 165, 168
 7254 GXT1000 video output option 165, 168
 8-port adapter
 configuring 184
 DIP-switch setting 183
 8-port asynchronous adapter 156
 80C186 processor 155
 8130 remote 16-port async. adapter 156
 8131 4.5 m 128-port async. cable 156
 8132 9 inch 128-port async. cable 156
 8134 non-US remote async. 16-port
 EIA-232 156
 8137 enhanced remote 16-port ESA-232
 node 157
 8138 enhanced remote 16-port RS-422
 node 157
 9.1 GB Hard Disk 99
 9291 single digital trunk processor 157
 9295 multiple digital trunk processor 157
 9398 base 2.2 GB SCSI-2 hot swap drive 60, 61
 9995 internal RAID indicator 61

A

a-law 57
 abbreviations 261
 abstract, publication xiii
 AC power supply 88
 accessing a second display 171
 acronyms 261
 adapter configuration
 adapter slots
 E30 44
 F40 56
 F50 74
 H50 82
 defining ISA adapters 181
 device configuration database 173

- adapter configuration (*continued*)
 - device driver installation 180
 - device location codes 174
 - device types 172
 - ISA 177
 - making ISA adapters available 182
 - non-graphic PCI 176
 - on PCI-based RS/6000 servers 172
 - PCI 176
 - process 178
 - recording settings of already configured ISA adapters 178
 - SCSI devices 177
 - selecting ISA hardware parameters 180
 - setting ISA DIP switches 182
- adapters
 - ATM 153
 - digital trunk quad 157
 - E30, communication 49
 - E30, SCSI 51
 - Ethernet 148
 - F40, communication 64
 - F40, SCSI 59
 - F50, communication 78
 - F50, SCSI 76
 - FDDI 150
 - H50, communication 86
 - H50, SCSI 85
 - ISA IBM 8-Port async EIA-232 185
 - ISDN 150
 - multiple port 154
 - overview 147
 - PCI SCSI-2 fast/wide differential 48
 - PCI SCSI-2 fast/wide RAID 47
 - PCI SCSI-2 fast/wide single-ended 48
 - PCI SCSI-2 single-ended Ultra SCSI 48
 - PCI SSA Multi-Initiator/RAID EL 119
 - POWER GXT1000 164
 - POWER GXT110P 159
 - POWER GXT120P 159
 - POWER GXT250P 160
 - POWER GXT255P 160
 - POWER GXT500P 161
 - POWER GXT550P 161
 - POWER GXT800P 162
 - SCSI 100
 - token-ring 147

- adapters (*continued*)
 - X.25 149
- AIX operating system
 - bootlist, command 232
 - debug mode, entering 232
 - electronic key-switch function 231
 - key-switches at boot time, function keys 221
- AMD SUPERNET3 processor 151
- ANSI standard
 - X3.131-1986 91
 - X3T10.1 115
 - X3T12 151
 - X3T9.2/375R Revision 10K 91
- APAR Ix60857 48
- Apple Inc., regarding CHRP 23
- architecture
 - ISA 28
 - PCI bus 26
 - SCSI-III 93
 - SSA 115
 - system 25
- ASCII interface, SMS 222
- async. adapters, configuration 183
- asynchronous communication, SCSI 97
- audio adapter, F40 57
- automatic configuration, PCI adapters 172
- auxiliary power supply
 - F40 61
 - H50 88

B

- backplanes, internal SSA 126
- backup
 - cloning 234
 - configuration information 241
 - restoring 234
 - restoring a system backup 234
- bay configuration
 - E30 42
 - F50 and F40 73
 - H50 81
- benchmark descriptions 13
- bibliography 253
- boot sequence, firmware 214
- boot, function keys 221

- bootinfo -p, command 233
- booting
 - from dump 222
 - special modes 221
- bootlist -m service, command
- bootlist, command 222
- BOS, installation log files 240
- bus
 - 60X 29
 - 6XX 29
 - ISA 28
 - PCI 28
 - SCSI 91
- business graphics 159

C

- cables
 - 7131 SSA (details) 123
 - 7133 SSA (details) 122
 - for 7131 and 7133 SSA disk drive
 - subsystems 120
 - SSA 120, 122
 - SSA loop to internal and external SSA
 - drives 125
- cables, SCSI (various types) 108
- cache
 - coherency 34
 - controller, X5 68
 - F40 54
 - hit 34
 - miss 34
 - PCI 27
- cache controller
 - X5 30, 33
- CD-ROM drives
 - Models F40, F50, and H50 99
 - RS/6000 Model E30, internal 46
- CEPT, telephone standard 157
- cfgmgr, command 176, 177
- CFP95, test suite 15
- channel service unit 157
- characteristics, RS/6000 2
- chdisp, command 170
- chpv, command 144
- CHRP
 - description of 23

- CHRP (*continued*)
 - goals of 24
 - history 23
- CINT95, test suite 15
- cloning
 - AIX 4.3 238
 - backup related 234
- cloning, AIX 4.1 235
- coherency, cache 34
- command
 - /usp -r 226
 - bootinfo -p 233
 - bootlist 222, 232
 - cfgmgr 176, 177
 - chdisp 170
 - chpv 144
 - errclear 246
 - errpt 246
 - importvg 142
 - importvg -y 141
 - lsattr 176, 179, 180
 - lscfg 176
 - lsdisp 170
 - lspv 141
 - lsresource 179, 182
 - mkdev 181
 - mksysb 234
 - rmdev 139, 142
 - rmlvcopy 145
 - shutdown 213
 - xinit 171
- commercial performance
 - RS/6000 Model E30 17
 - RS/6000 Model F40 17
 - RS/6000 Model F50 17
 - RS/6000 Model H50 17
- common command set, SCSI 92
- communication adapters 147
 - ATM 153
 - configuring ISA 199
 - Ethernet 148
 - F40 64
 - FDDI 150
 - H50 86
 - ISDN 150
 - multiport 154
 - RS/6000 Model E30 49

- communication adapters (*continued*)
 - token-ring 147
 - X.25 149
- Compaq, regarding PCI involvement 26
- concurrent mode
 - diagnostics 211
- configuration information, system backup 241
- configurations
 - ISA from diagnostics 208
 - RS/6000 Model E30 51
 - RS/6000 Model F40 66
 - RS/6000 Model F50 79
 - RS/6000 Model H50 89
- considerations, stand-alone diagnostics 207
- console mirroring, SP console 227
- controller, memory I/O 69
- cross invalidate 35
- Ctrl-\ 233
- Ctrl-alt-numpad4 232
- cycle redundancy check 215, 218

D

- DASD 91, 98
- data service unit 157
- debug mode, entering 232
- default boot list
- device configuration
 - database 173
 - database for storage of 173
 - defining ISA adapters 181
 - device driver installation 180
 - device types 172
 - ISA adapters, DIP switch setting 182
 - ISA adapters, general discussion 177
 - ISA adapters, recording settings of already configured 178
 - ISA adapters, selecting hardware parameters 180
 - location codes 174
 - making ISA adapters available 182
 - non-graphic PCI 176
 - on PCI-based RS/6000 servers 172
 - PCI adapter 176
 - SCSI devices 177
- devices, at boot time 232

- devices.scsi.disk filesset 235
- diag, command 213
- diagnostics
 - booting stand-alone from CD-ROM 207
 - diagnostic support 205
 - ISA adapters, configuring from diagnostics 208
 - online 210
 - online in concurrent mode 211
 - online in maintenance mode 213
 - online in service mode 212
 - operating considerations 205
 - running stand-alone in concurrent mode 210
 - Service Aid 226
 - stand-alone 207
 - stand-alone considerations 207
- differential Ultra SCSI adapter, PCI 48
- differential, SCSI 97
- Digital (DEC), regarding PCI involvement 26
- digital trunk quad adapter 157
- DIMM
 - RS/6000 Model E30 51
 - RS/6000 Model F40 54
 - RS/6000 Model F50 71
 - RS/6000 Model H50 84
- disconnect/reconnect, SCSI 92
- disk bays, F40 58
- disk drives 98
 - E30 45
 - F40 58
 - F50 76
 - H50 84
 - supported in hot swap 130
- dual-monitor 158
 - changing default display 170
 - E30 46
 - F40 53, 63
 - F50 77
 - feature comparison 166
 - multiple adapter support 169
 - POWER GXT1000 164
 - POWER GXT110P 159
 - POWER GXT120P 159
 - POWER GXT250P 160
 - POWER GXT255P 160
 - POWER GXT500P 161
 - POWER GXT550P 161

dual-monitor (*continued*)
POWER GXT800P 162
second display 171
second PCI adapter 169
dual-monitor, graphics adapter 158
DVD 20

E

E30 44
edition notices, publication ii
electronic keyswitch 231
entering debug mode 232
EPROM update, SP flash 228
errclear, command 246
error log, accessing through SMS 225
error, power-off alert 218
errpt, command 246
Ethernet adapters 148
external subsystems, RS/6000 Model E30 50

F

F1-key, boot function 222
F40
communication adapters 64
configuration notes 66
detailed description 52
graphics adapters 63
performance 12
SCSI adapters 59
standard features 54
F5-key, boot function 222
F50
configuration notes 79
graphics adapters 77
highlights 67
performance, TPC-C 17
standard features 71, 81
system architecture 68
TPC-D performance results 17
F6-key, boot function 222
F8-key, boot function 222
false sharing 39
fast
SCSI 96

fast-20 narrow, SCSI 93
fast-20 wide
SCSI 93, 96
fast-20, SCSI 96
fast/wide, regarding SCSI-2 96
FDDI adapters 150
feature comparison, graphics adapters 166
feature number
2443 6-pack SCSI to DASD cable and terminator 60
2444 6-pack SCSI to DASD cable 60, 61, 85
2446 PCI Ultra-SCSI to 6-pack cable 61
2447 6-pack PCI SCSI SE to DASD cable 61, 85
2450 6-pack SSA to PCI SSA adapter cable 63
2452 SSA blank disk module 63, 118
2453 SSA internal to SSA external cable 63
2493 PCI RAID adapter 47, 59, 61, 77, 85, 86, 100, 102, 114
2618 8x CD-ROM 99
2619 20x max. CD-ROM 46, 99
2630 1/4 inch internal tape 66
2638 PCI video capture card 64
2701 4-port multiprotocol ISA adapter 49, 64, 78, 87, 155, 191
2705 4-port multiprotocol cable 155, 191
2708 Eicon ISDN DIVA PRO 150
2741 SysKonnnect FDDI-LP adapter (SAS) 50, 65, 78, 87, 150, 151, 154
2742 SysKonnnect FDDI-LP adapter (DAS) 50, 65, 78, 87, 150, 151, 154
2743 SysKonnnect FDDI-UP adapter (SAS) 50, 65, 78, 87, 150, 152, 154
2837 MVP multi-monitor adapter 158, 166
2838 GXT120P graphics adapter 86, 159, 166
2839 GXT110P graphics adapter 159, 166
2851 GXT250P graphics adapter 160, 166
2852 GXT255P graphics adapter 160, 166
2853 GXT800P graphics adapter 162, 168
2854 GXT500P graphics adapter 161, 166
2855 GXT550P graphics adapter 168
2859 GXT800P adapter with texture 162, 168
2900 4.5 GB Ultra SCSI hot swap drive 45, 61
2908 9.1 GB Ultra SCSI hot swap drive 45
2913 9.1 GB Ultra SCSI hot swap drive 84

feature number (*continued*)

2919 9.1 GB Ultra SCSI hot swap drive
select 84
2920 PCI token ring adapter 64, 78, 86, 148
2931 8-port asynchronous adapter 155, 156,
183
2932 8-port async. adapter EIA
232/RS-422 156
2933 128-port asynchronous controller 156,
186
2943 8-port asynchronous adapter PCI 156
2944 128-port asynchronous adapter PCI 157
2961 X.25 ISA co-processor adapter 49, 64,
78, 87, 149, 197
2962 2-port multiprotocol PCI adapter 49, 64,
78, 87, 157
2963 TURBOWAYS 155 PCI UTP ATM
adapter 49, 64, 78, 87, 153
2968 10/100 Mbps Ethernet PCI adapter 49,
64, 78, 86, 148
2979 PCI token ring adapter 49, 64, 78, 86,
147
2985 PCI Ethernet BNC/RJ-45 adapter 49,
64, 78, 86, 149
2986 3Com 10/100 Mbps PCI Ethernet
adapter 49, 64, 78, 86, 148, 150, 154
2987 PCI Ethernet AUI/RJ-45 adapter 49, 64,
78, 86, 149
2988 TURBOWAYS 155 PCI MMF ATM
adapter 49, 64, 78, 87, 150, 153, 154
2998 TURBOWAYS 25 ATM PCI adapter 49,
64, 78, 87, 154
3005 4.5 GB Ultra-SCSI drive select 61, 62,
77
3019 9.1 GB F/W Ultra-SCSI drive select 84
3033 2.2 GB SCSI F/W hot swap drive 45
3071 4.5 GB SSA hot swap drive 62, 77, 118
3072 9.1 GB SSA hot swap drive 62, 63, 77,
84, 118
3095 2.1 GB SCSI hard drive 45
5500 fiber optic extender 121
6137 4 mm internal tape 66
6142 4/8 GB internal tape drive 46, 100
6147 5/10 GB 8 mm internal tape 100
6154 20/40 GB 8 mm internal tape 100
6159 12/24 GB internal tape drive 46, 100
6206 PCI single-ended Ultra SCSI
adapter 47, 48, 59, 61, 62, 76, 85, 86, 100,
102, 106, 108

feature number (*continued*)

6207 PCI differential Ultra SCSI adapter 47,
48, 59, 62, 77, 85, 86, 100, 102, 111
6208 PCI SCSI-2 F/W SE adapter 47, 48, 59,
60, 76, 85, 100, 101, 102, 106, 108, 109
6209 PCI SCSI-2 F/W differential adapter 47,
48, 59, 77, 85, 100, 101, 102, 111
6215 PCI SSA multi-initiator RAID
adapter 62, 77, 86, 119, 120
6218 PCI SSA 4-port RAID adapter 63
6222 SSA fast write cache option card 119,
120
6296 redundant A/C power supply 88
6309 digital trunk quad PCI adapter 86, 157
6424 boot drive mounting kit (H50) 85
6427 boot drive mounting kit (H50) 86
6509 media-bay mounting kit 61
6520 first six SCSI hot swap backplane 60
6523 SCSI 6-pack non hot swap
backplane 60, 61
6527 F40 SSA base backplane 58
6537 SCSI 6-pack hot swap backplane 60, 61
6538 SSA select backplane 62, 77
6539 SSA backplane power cable 62, 77, 86
6549 additional power supply for SSA 61, 62
7252 GXT1000 internal graphics
accelerator 164, 168
7253 GXT1000 advanced graphics
option 165, 168
7254 GXT1000 video output option 165, 168
8130 remote 16-port async. adapter 156
8131 4.5 m 128-port async. cable 156
8132 9 inch 128-port async. cable 156
8134 non-US remote async 16-port
EIA-232 156
8137 enhanced remote 16-port EIA-232
node 157
8138 enhanced remote 16-port RS-422
node 157
9291 single digital trunk processor 157
9295 multiple digital trunk processor 157
9398 base 2.2 GB SCSI-2 hot swap drive 60,
61
9995 internal RAID indicator 61
DASD 99

- feature summary
 - communication adapters 147
 - E30 3
 - F40 5
 - F40, optional 53, 58
 - F50 8
 - graphics adapters 158
 - H50 11
 - H50 RAS 81
 - internal storage 91
 - PCI 27
 - SCSI 91
 - SSA 115
- Fiber Channel Physical and Signaling Interface 95
- fiber optic extender (SSA) 121
- filesets, SCSI devices 235
- FireWire (IEEE 1394) 19
- firmware
 - boot sequence 214
 - flash update 217
 - major sections 213
 - passwords 216
 - power-on password 216
 - privileged access password 217
 - recovery 219
 - understanding it 213
- five resources for ISA adapters 178
- FixDist, APARS 48
- flash EPROM
 - update 217, 228
- function keys, boot time 221

G

- Genlock 165
- GL 3.2, supporting adapters 163
- graphics adapter 170
 - E30 51
 - F40 63
 - F50 77
 - H50 86
- graphics support matrix 158
- GXT1000
 - in F40 64
 - internal graphics adapter 164

- GXT110P
 - graphics adapter 159, 166, 167
 - RS/6000 Model E30 46
- GXT120P
 - graphics adapter 166, 167
 - screen resolutions 160
- GXT250P, graphics adapter 160, 166, 167
- GXT255P, graphics adapter 160, 166, 167
- GXT500P, graphics adapter 161, 166, 167
- GXT550P, graphics adapter 161
- GXT800P, graphics adapter 162

H

- H50
 - Bays 82
 - communication adapters 86
 - disk drives 84
 - highlights 80
 - performance, TPC-C 17
 - RAS features 81
 - RS/6000 rack server 79
 - SCSI adapters 85
 - TPC-D performance results 17
- HA50 special offering 88
- HA50, RS/6000 highly available server 88
- hard disks 91
- hardware
 - design 25
 - overview 23
 - packaging 41
- hints and tips
 - accessing System Management Services with an ASCII terminal 225
 - error logging facility, using 246
 - ISA adapter configuration 177
 - ISA adapter configuration information 176
 - logging facility, using 246
 - power-on passwords 216
- hit ratio 34
- hit ratio, multiprocessing concepts 34
- hot swap
 - adding new drive to live system 141
 - mirroring 144
 - removing a drive with its own volume group 142
 - replacing previously defined drive 142

hot swap (*continued*)
subsystem 128

I

I/O ports
E30 44
F40 56
F50 74
H50 83
ISA setting 182
IBM, regarding PCI involvement 26
IEEE 1275, CHRP Standard 24
IEEE 1394 High Performance Serial Bus 19, 95
illustration - picture
E30 42
F40 52
F50 67
H50 79
importvg -y, command 141
importvg, command 142
installation
BOS log files 240
hot swap 133
integrated devices 172
integrated SCSI-2 adapter
E30 44
F40 56
F50 74
H50 82
Intel, regarding PCI involvement 26
internal
CD-ROM drives, E30 46
CD-ROM drives, F40, F50, H50 99
graphics adapter, GXT1000 164
SSA backplanes 126
storage features 91
tape drives 99
E30 45
F40 59
F50 76
H50 85
introduction, general 1
IrDA 24
IRQ, ISA setting 182
ISA
attribute selection 209

ISA (*continued*)
description of 28
ISA adapter
128-port EIA-232 asynchronous 186
4-port multiprotocol 191
8-port EIA-232 asynchronous 183
communication 147
five resources for 178
testing 210
X.25 interface co-processor 197
ISA IBM 8-Port Async EIA-232 adapter 185
ISA X.25 Coprocessor/1 Adapter 180
ISDN adapters 150
ISO 9241 160
Ix60857, APAR 48

K

key attributes
RS/6000 Model E30 2, 3
RS/6000 Model F40 5
RS/6000 Model F50 7, 11
RS/6000 Model H50 10
key-switches at boot time, function keys 221
keyswitch, electronic 231

L

L1/L2 caches
LRU (Least Recent Used) 39
RWNITM 38
store policy 38, 39
languages, setting console 223
lanstreamer adapter 147
LINPACK DP description 13
LINPACK TPP description 13
LocalTalk 24
Logical Block Diagram 25
logical key-switch 231
LRU (Least Recent Used) 39
lsattr command 176, 179, 180
lscfg, command 176
lsdisp command 170
lspv, command 141
lsresource command 179, 182

M

- main purposes, service processor 226
- maintenance mode
 - diagnostics 211
 - diags 213
- memory
 - E30 42
 - F40 58
 - F50 71
 - H50 81
 - hierarchy 32
 - I/O controller 69
 - SDRAM 71
- menus, service processor 226
- MESI protocol 35, 70
- metrics, benchmarking 16
- mirroring, hardware console 227
- mkdev, command 181
- mksysb, command 234
- Model E30 2, 41
- Model F40 52, 53
- Model F50 66, 67, 68
- Model H50 79
- models, naming scheme 2
- Motorola, regarding CHRP 23
- mu-law 57
- multi-monitor, graphics adapter 158
- multi-port adapters, configuration 183
- multiple port adapters 154
- multiprocessing
 - cache coherency problem 34
 - cache hit versus cache miss 34
 - concepts 32, 34, 35, 38, 39
 - cross invalidate 35
 - false sharing 39
 - hardware characteristics 32
 - L1/L2 caches and store policy 38
 - memory hierarchy 32
 - MESI protocol 35
 - sharing resources 32
 - snooping 35
- multiprotocol adapters 157
- MVP Power Multi-Monitor Graphics Accelerator 158

N

- naming convention, RS/6000 2
- NCR, regarding PCI involvement 26
- non-uniform memory access 30
- normal, keyswitch 231
- notes
 - RS/6000 Model E30 configuration 51
 - RS/6000 Model F40 configuration 66
 - RS/6000 Model F50 configuration 79
 - RS/6000 Model H50 configuration 89
- NTSC output 165
- NUMA 30
- NURBs 169

O

- Object Data Manager 234
- ODM 234
- online diagnostics 205
- Open Firmware, CHRP Standard 24
- OpenGL, supporting adapters 163
- optional features
 - E30 44
 - F40 53, 58
 - F50 75
 - H50 84
- overview
 - hardware 23
 - RS/6000 Model E30 2, 41
 - RS/6000 Model F40 4
 - RS/6000 Model F50 7
 - RS/6000 Model H50 10

P

- packaging, hardware 41
- passwords
 - firmware 216
 - hardware related 216
- PCI
 - bus architecture 26
 - features and benefits 27
 - general discussion 24
 - Special Interest Group 26
 - Version 2.0 27
 - X-Bus 30

- PCI adapters, communication 147
- PCI differential Ultra SCSI adapter 48
- PCI SCSI-2 adapters, RS/6000 Model E30 47
- PCI SCSI-2 fast/wide differential adapter 48
- PCI SCSI-2 fast/wide RAID adapter 47
- PCI SCSI-2 fast/wide single-ended adapter 48
- PCI SCSI-2 single-ended Ultra SCSI adapter 48
- PCI SCSI, F50 76
- PCI SSA multi-initiator/RAID EL adapter 119, 120
- PCI-based RS/6000 servers,
 - troubleshooting 246
 - accessing diagnostics with NIM
 - booting from CD-ROM 247
 - booting hangs during boot process 247
 - booting problems 246
 - dealing with power failures
 - hints and tips
 - accessing System Management Services with an ASCII terminal 225
 - ISA adapter configuration 177
 - ISA adapter configuration information 176
 - power-on passwords 216
 - installation problems
 - boot device order list
 - microcode problems 218
 - data loss 218
 - language problems 219
 - miscellaneous
 - power-up problems 218
 - system recovery with no bootable media 248
- performance summary 12
 - RS/6000 Model E30 16, 17
 - RS/6000 Model F40 16
 - RS/6000 Model F50 16
 - RS/6000 Model H50 16
- permanent storage 91
- PHIGS, supporting adapters 163
- Plug and Play, PCI 28
- PNP descriptors 178
- POST 221
- POST, processor 220
- POWER GXT1000, internal graphics adapter 164

- POWER GXT110P
 - graphics adapter 159, 166, 167
- POWER GXT120P, graphics adapter 159, 166, 167
- POWER GXT250P
 - graphics adapter 160, 166, 167
- POWER GXT255P
 - graphics adapter 160, 166, 167
- POWER GXT500P
 - graphics adapter 161, 166, 167
- POWER GXT550P, graphics adapter 161
- POWER GXT800P, graphics adapter 162
- power-on password 216
- power-on self test
 - DIMM and L2 cache POST 220
 - graphics adapter POST 221
 - keyboard and mouse controller POST 221
 - POST 220
 - subsystem POST 221
- PowerPC 630 processor 18
- PowerPC futures 18
- PowerPC Microprocessor Common Hardware Reference Platform 23
- privileged Access password 216
- processor
 - 80C186 155
 - AMD SUPERNET3 151
 - POST 220
 - PowerPC 630 18
 - RS64 18
- processor independence, PCI 28
- processor subsystem 31
- processor subsystem, RS/6000 Model E30 42
- publication
 - abstract xiii
 - bibliography 253
 - target audience 249

R

- RAID, SCSI, F40 61
- RAS
 - diagnostics 205
 - H50 81, 88
- recovery, procedure 219
- redundant AC power supply 88

- related publications 253
- repeaters, SCSI 98
- resetting, terminal 206
- restoring a system backup 234
- ring indicate power on 226
- RJ-45 connector 147
- rmdev
 - command 139, 142
- rmlvcopy, command 145
- rmt, as boot device 232
- RS/6000
 - general characteristics 2
 - Model E30, standard features 42
 - performance results 12
 - performance, Model E30 16
 - performance, Model F40 16
 - performance, Model F50 16
 - performance, Model H50 16
 - product description 41
 - product line 1
 - server offerings, picture 1
 - servers, general discussion 2
 - trends and directions 18
- RS/6000 Model E30 2
 - commercial performance 17
 - communication adapters 49
 - configuration notes 51
 - external subsystems 50
 - graphics adapter 46
 - internal CD-ROM drives 46
 - internal disk drives 45
 - internal tape drives 45
 - key features 2, 3
 - limitations 51
 - PCI SCSI-2 adapters 47
 - performance summary 16
 - processor subsystem 42
 - product description 41
 - supported optional features 44
 - System dimensions 4
- RS/6000 Model F40 4
 - commercial performance 17
 - key features 5
 - overview 4
 - performance summary 16
 - standard features 6
 - upgrade 65
- RS/6000 Model F50 7
 - commercial performance 17
 - feature summary 8
 - graphics adapters 77
 - highlights 67
 - key features 7, 11
 - optional features 75
 - overview 7
 - performance summary 16
 - standard features 71, 81
 - system architecture 68
 - TPC-D performance results 17
- RS/6000 Model H50 10
 - bays 82
 - commercial performance 17
 - general information 79
 - key attributes 10
 - overview 10
 - performance summary 16
 - TPC-D performance results 17
- RS64 processor 18
- RWNITM (Read With No Intent to Modify) 38

S

- S1 serial port, service processor 226
- scdisk, as boot device 232
- screen resolutions, GXT120P 160
- SCSI
 - adapters 100
 - adapters, F40 59
 - asynchronous 97
 - common command set 92
 - devices, configuration 177
 - differential 97
 - disconnect/reconnect 92
 - IDs, hot swap 132
 - introduction 91
 - physical interface layers 93
 - RAID, F40 61
 - repeaters 98
 - specifications summary 98
 - summary of specifications 98
 - synchronous 97
 - table of modes/rates 98
 - tagged command queuing 92
 - terminology 96

- SCSI (*continued*)
 - Ultra 96
 - Ultra-2 20
- SCSI-2 91
- SCSI-2 adapters, RS/6000 Model E30, PCI 47
- SCSI-2 fast/wide differential adapter, PCI 48
- SCSI-2 fast/wide RAID adapter, PCI 47
- SCSI-2 fast/wide single-ended adapter, PCI 48
- SCSI-2 single-ended Ultra SCSI adapter, PCI 48
- SCSI-I 91
- SCSI-III 93
 - Fiber Channel Protocol 93
 - physical interface layers 93
 - protocol standards 93
 - SSA definition within SCSI III specification 93
 - standards 93
 - transfer rate 93
- SCSI-III interconnect technologies 93
- SCSI-III Parallel Interface 95
- SCSI-III Parallel Interface (SPI) 95
- SCSI, location codes 174
- SDRAM memory 71
- second display, accessing 171
- secure, keyswitch setting 231
- self tests 220
- serial storage architecture 115
- Serial Storage Architecture, in SCSI-III 95
- Service Aid
 - diagnostics 226
 - menus 228
- service mode
 - diagnostics 211
 - diags 212
- service processor 226
 - F50 72
 - main purposes 226
 - menus of 226
 - surveillance 228
- service, keyswitch 231
- shutdown, command 213
- slots
 - adapter 56
 - E30 44
 - F40 56
 - F50 74
 - H50 80
- Small Computer System Interface 91
- SMP hardware characteristics 32
- SMS (see System Management Service)
- SMT 7.3 standard 151
- snooping 35
- Softgraphics 169
- SP Menus 228
- SPEC95 software benchmark 15
- SPECfp_base_rate95 description 13
- SPECfp_base95 description 13
- SPECfp95 description 13
- Special Interest Group, PCI 26
- SPECint_base_rate95 description 13
- SPECint_base95 description 13
- SPECint95, description 13
- SPI (SCSI Parallel Interface) 95
- SSA 115
 - architecture 115
 - cables 120, 122, 125
 - internal and external 127
 - loop to internal and external SSA drives 127
 - support, RS/6000 models 86
 - topology 116
- stand-alone diagnostics 205
 - considerations 207
 - running 207
 - support 207
- standard features
 - E30 42
 - F40 6, 54
 - F50 71
 - H50 81
- standards
 - ANSI X3.131-1986 91
 - ANSI X3T12 151
 - CEPT 157
 - CHRP 23
 - PREP 23
 - SCSI-III 93
 - SMT 7.3 151
 - X3T9.2/375R Revision 10K 91
- start mode, unattended 217
- start-up menu, SMS 224
- storage options, F40 53
- store policy, L1/L2 caches and store policy 38

- string topology, SSA 116
- subsystems, initializing 215
- supported optional features
 - RS/6000 Model E30 44
 - RS/6000 Model F40 58
 - RS/6000 Model F50 75
 - RS/6000 Model H50 84
- surveillance, service processor 228
- symmetric multiprocessing, F40 4
- synchronous, SCSI 97
- system backup
 - cloning 234
 - configuration information 241
- System dimensions, RS/6000 Model E30 4
- System Management Service
 - accessing error log 225
 - ASCII interface 222
 - ASCII main menu 223
 - graphical main menu 223
 - overview of 222
 - start-up menu 224
 - test menu 224
 - tools menu 225
 - using 245
- system recovery, procedure 219
- system, architecture 25

T

- tagged command queuing, SCSI 92
- tape drives
 - internal, F40 F50 H50 99
 - internal, RS/6000 Model E30 45, 59, 76, 85
- target audience, publication 249
- terminal, tty 206
- terminology, SCSI 96
- testing ISA adapters 210
- tests, machine
- Texture mapping 169
- token-ring adapters 147
- TPC-C 14
 - F50 17
 - H50 17
- TPC-D 14, 17
- trends and directions, RS/6000 18
- troubleshooting 246
 - accessing diagnostics with NIM

- troubleshooting (*continued*)
 - booting from CD-ROM 247
 - booting problems 246
 - dealing with power failures
 - hints and tips
 - accessing System Management Services
 - with an ASCII terminal 225
 - ISA adapter configuration 177
 - ISA adapter configuration
 - information 176
 - power-on passwords 216
 - installation problems
 - boot device order list
 - microcode problems 218
 - data loss 218
 - language problems 219
 - miscellaneous
 - power-up problems 218
 - system hangs during boot process 247
 - system recovery with no bootable
 - media 248
- tty terminal, diagnostics 206
- twin-monitor, graphics adapter 158

U

- Ultra SCSI 96, 98
- Ultra SCSI adapter, PCI differential 48
- Ultra SCSI Wide vs Ultra SCSI 96
- Ultra SCSI, defined 96
- Ultra-2 SCSI 20
- UMS/6000, GXT120P 160
- unattended start mode 216, 217
- understanding firmware 213
- Universal Serial Bus 18
- update SP flash EPROM 228
- upgrade, F40 to F50 65
- USB 18

V

- VDP, read through SP 229
- virtual disk file system 215
- vital product data 218
- VOO feature 165
- vpd 218

W

WAN 157

wide

 regarding SCSI 96

X

X-Bus, PCI 30

X.25 adapter

 configuring 199

 Coprocesor/1 Adapter, ISA 180

 DIP-switch setting 198

 overview 149

X5 cache controller 30, 33, 68

xinit, command 171

Z

Z-buffer, 3D adapters 163

ITSO Redbook Evaluation

RS/6000 Models E30, F40, F50, and H50 Handbook
SG24-5143-00

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

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!)

